

Track geometry standard deviation calculator

Part 1

AUTHOR:

David Marriott
Consultant/developer
Freelance

In this article David will describe the way in which the standard deviation values produced by Network Rail's Track Recording Cars are interpreted and sometimes misinterpreted. He will also explain the reasons why the Track Geometry SD Calculator was developed. In doing so, David will look in some detail at the way the track data is processed to compute SDs and in particular, the effect of the filtering which is employed.

TRACK RECORDING CAR

The track recording cars (TRC) operated on behalf of Network Rail use an inertial reference system and displacement sensors mounted on the recording vehicle to determine the position of the track in space. From this positional data, profiles oriented vertically (top) and horizontally (alignment) are derived.

Each profile can be considered to be a sum of two parts, the first part being the profile corresponding to the designed geometrical shape of the track, the second being the irregularities or deviations from that designed shape.

Ideally, the profile corresponding to the designed shape of the track should be subtracted from the recorded profile, leaving the irregularities or deviation from design. It is the deviation from design which, when quantified, can then be used to monitor the geometrical quality of the track.

At the time the TRC was first developed by British Rail in the 1970s, other than simple and predominantly paper records of gradient and curvature, there was no network-wide database of the designed shape of track which was defined accurately enough to create such profiles. Consequently, the deviation from design had to be isolated by an alternative method.

FILTERED PROFILES

A standard way of isolating one part of an electrical signal which is a sum of parts is to use a filter. For example, hi-fi equipment may use a filter to suppress interference of a mains power supply from an audio signal. The filter is an electronic device which is selective in terms of frequency and the principle of operation in the hi-fi is that the noise from the mains is at a specific frequency (50Hz). The filter attenuates only the 50Hz component of the audio signal, removing the interference from the mains power supply, leaving only that part of the signal corresponding to the wanted music or speech.

The same idea can be used to remove the profile corresponding to the designed shape of the track from the profile recorded by the TRC, leaving only the deviations from design. The main difference being that while the audio signal varies with time, the recorded profile of the track varies with position along the track. So instead of being selective in terms of temporal frequency in cycles per second (Hz), the filter used to remove the design profile is selective in terms of spatial frequency in cycles per metre (m⁻¹). However, instead of spatial frequency, we normally refer to its inverse, namely: wavelength (m). The principle of operation in the TRC is that the profile corresponding to the design shape of the track is predominantly of long wavelength, that of the irregularities in the track are of short wavelength.

To isolate only the short wavelength part of the profile, we need to use a "high pass" filter, named because only high spatial frequencies (short wavelengths) will pass through the filtering process, longer wavelengths will be removed. A "cut off" wavelength defines where the boundary lies between the long and short wavelengths.

However, it's not possible to create a filter which has a sharp cut-off between the long and short wavelengths. A practical filter is an approximation to this ideal and many types of filter have been produced, each with different approximations to that ideal. The type of filter used in the TRC is a 4-pole, high-pass, Butterworth filter.

In fact, the TRC uses two different Butterworth filters, one with a 70m cut off wavelength and one with a 35m cut off wavelength. These produce two different versions of recorded irregularity for both top and alignment, the longer 70m wavelength irregularities being relevant for higher line speeds.

The choice of the cut-off wavelengths and the choice of the 4-pole Butterworth for the filter type, was made by the International Union of Railways (UIC). 35m (25 m for other UIC members) being representative of irregularities which could typically be corrected by automatic tamping, and 70m the irregularities which could typically be corrected by design tamping.

Isolating irregularities at wavelengths much shorter than 35m is possible, but irregularities at wavelengths as short as 2-3m are associated with rail shape which could only be corrected by bending, grinding or planing of the rails.

STANDARD DEVIATION

Having isolated the irregularities in top and alignment, the next step is to quantify them. The TRC produces 1/8 mile standard deviations (SD) for top and alignment for both 35m and 70m filtered irregularities. So, four values for each 1/8 mile of recorded track.

The process is quite straight forward with the filtered profiles being sampled at 200mm longitudinal spacing along the track.

The sampled data is then separated into 1/8 mile lengths aligning with milepost mileage, producing just over 1000 data values for each filtered profile for each 1/8 mile.

To compute the SD of each 1/8 mile length of sampled profile, we use the formula (1), where the z_i are simply the sampled values of the filtered track profile within the 1/8 mile.

$$SD = \sqrt{\frac{\sum_{i=1}^n (z_i - \frac{\sum_{i=1}^n z_i}{n})^2}{n - 1}} \quad (1)$$

This is essentially how the 70m and 35m Alignment SDs are produced by the TRC, using the filtered horizontal profile of the track position.

There is a similar process for the 70m Mean Top SD, using a 70m filtered vertical profile of the track centreline position.

For the 35m Top, 1/8 mile SDs are computed for both left and right hand rails, using separate 35m filtered vertical profiles of the positions of each rail. The SD which is output for 35m Worst Top SD is just the larger of those two values.

INTERPRETATION OF THE SD VALUES

As standard deviation is a measure of variation from the mean, a larger SD value means the more the track varies from its designed position along that 1/8 mile.

When the TRC was first developed, the SDs were used to indicate the general level of “roughness” of the track, with bands defined as “good”, “satisfactory”, “poor” and “very poor” for each track category, with increasing levels of SD for each of the four parameters. The percentage of 1/8 miles in each of the bands being an indicator, where an increasing percentage in the poor or very poor band from one recording run to the next, indicated a general deterioration of a route, e.g Rugby to Crewe.

For a particular 1/8 mile, an SD value increasing from one recording run to the next indicated a section of track with deteriorating “roughness”. Here “roughness” refers to the deviation from the designed track shape rather than the surface roughness of the rails.

NEW INTERPRETATIONS

Over time, the need arose to quantify track quality to manage interfaces, e.g. between the infrastructure operator and a tamping contractor. As TRC SDs were readily available, the SD values began to be used in new ways.

For example, the fact that a recorded SD of a particular parameter for a particular 1/8 mile was in the “poor” band was taken to indicate that track was “poor” with respect to e.g. 70m Alignment for that 1/8 mile.

Decisions about which track to tamp and the performance of the tamping contractor was then based on the SDs of individual 1/8 miles and their values recorded before and after tamping.

As we shall see shortly, it is misleading to use SDs in this way because the value of the 1/8 mile SD is not an absolute measure of track quality.

NEW CONSTRUCTION STANDARDS

Around 2005, Network Rail inserted a requirement into its specification for track construction which was that when recorded by the TRC, the “as built” track should have SDs which were in the “good” band for the relevant track category. This caused a problem.

In cases where the track had been installed and adjusted to its designed geometry per the specification, the TRC SD values were nevertheless recording some 1/8 miles as “satisfactory” or even “poor” rather than “good”. This non-compliance seemed to occur where the track design included significant changes of curvature and cant.

EFFECT OF GEOMETRICAL DESIGN ON FILTERED PROFILE

To see why this is the case, Figure 1 shows a typical section of track having a designed vertical profile comprising gradients, vertical curves, cant transitions, etc, but with no irregularities present.

The shape of the designed profile contains predominantly long wavelengths and as there are no short wavelength irregularities we might expect that when vertical profile of the track centreline is filtered with the 70m filter, the result would be a zero-valued profile corresponding to “perfect” track.

But that isn't what happens.

The filtered profile in Figure 1 is not zero at all. There are strange blips and things which look like some kind of damped oscillation, which if we look closer appear to be related to the points in the original profile where there is an instantaneous change, either of gradient (at the ends of the cant ramps) or curvature (at the ends of the vertical curves)

As shown in Figure 1, because the filtered profile is not zero valued everywhere, when we separate this filtered profile into 1/8 mile sections and compute the MT70 SDs, not all are zero.

Reminding ourselves that this is a vertical profile of track having no defects or irregularities at all, the SDs we see might be large enough to initiate some kind of intervention, such as tamping, to fix the non-existent defects. So what is going on?




Image 1: Design cant and vertical profile and their filtered counterpart

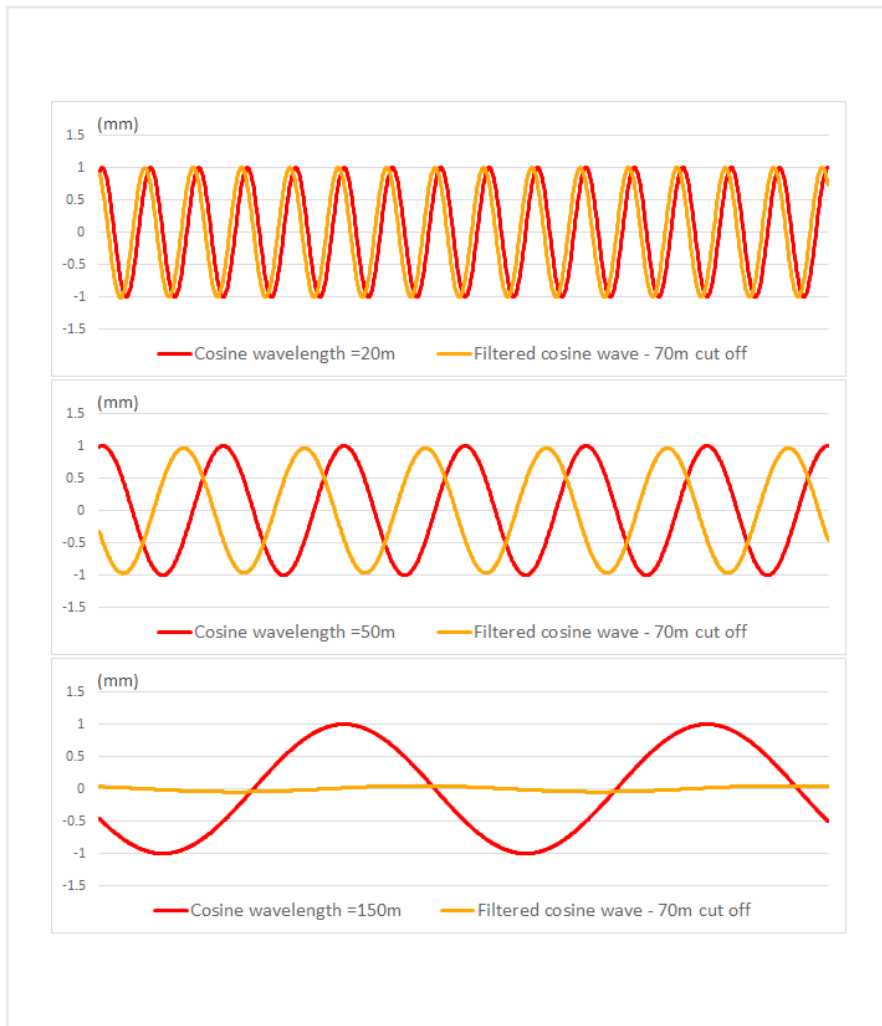


Figure 2. Cosine waves and their filtered counterpart

FILTER BEHAVIOUR

Let's look in more detail about how our filter behaves.

If the input is a periodic function, like a cosine wave, the output is nicely predictable. The graphs in Figure 2 shows how cosine waves of 3 different wavelengths are changed by filtering with a 70m cut-off. In each case the filtered output is a cosine wave of the same wavelength, but with amplitude changed and phase (the location of the peaks) shifted.

The gain of the filter is the ratio of the amplitudes of output to input at a particular wavelength. Looking at the graphs, the gain of the 70m filter is virtually 1 at (short) 20m wavelength, just less than 1 at 50m and virtually zero at (long) 150m wavelength.

The way the filter treats different wavelength components of a periodic signal is known as the transfer function of the filter. It is the transfer function having a gain of 1 at short wavelength and a gain of zero at long wavelength which determines that our 4-pole Butterworth is a high-pass filter.

If our recorded profile is a summation of many cosine waves of different wavelength, the filtered profile will also be periodic with the same wavelengths present, but each with their amplitude and phase shifted according to the transfer function of the filter.

But what about the functions we use in our track designs?

DESIGN ELEMENTS

Figure 3 shows a profile comprising a straight gradient. When it's filtered, the output is just a zero-valued profile. The designed gradient has been "filtered out", as intended. Figure 4 shows a profile comprised of a parabolic curve which is the form of our designed vertical curves. When its filtered, once again we see a zero-valued profile.

When filtering any other single and continuous vertical, horizontal or cant design element the result will be the same – a zero-valued profile. Since our designed vertical profile comprises only straight gradients, cant ramps of straight cant gradient and parabolic vertical curves, we would expect the filter to remove all the design profile, right?
Wrong.

Figure 5 shows an instantaneous change of gradient, which we would find also at the start of a cant ramp. The figure illustrates the effect of the instantaneous change of gradient on the filtered output. It isn't a periodic function and neither is it zero valued.

Figure 6 shows a horizontal curvature transition and its filtered counterpart.

The filtered output is a different, non-periodic, non-zero valued shape.

Perhaps now we can see why the filtered profile of our designed vertical profile had the shape it did. It was the summation of all these effects caused by the joins between our geometrical design elements. We would see similar, but different, effects at instantaneous changes of horizontal curvature and at the start and end of vertical curves where there is a join between the parabolic shape of the curve and the straight gradients.

FILTER "ARTEFACTS"

I describe these effects - the shapes that the filter produces - as "artefacts" which is a word most commonly used to describe the crafted objects which archaeologists find when digging e.g. a Saxon burial, but I'm using the word in its other sense, to describe "something observed in a scientific investigation or experiment that is not naturally present but occurs as a result of the preparative or investigative procedure".

Do these artefacts somehow describe the effective short wavelength shape of the track at those joints?

No. They aren't real and perhaps I could illustrate this with an analogy.

Figure 7 shows the map of the world with which we are most familiar. It was probably on the classroom wall when we were at school and we also see it in our atlas and illustrating stories in the newspaper. But this is not a picture of the actual surface of the globe. It is a projection of that curved surface onto a cylinder having the same axis as the earth's rotational axis - the Mercator projection.

Then that cylinder is cut along a line of longitude, usually 180 degrees, and opened out to form the flat map. It handily shows the arrangement of continents and countries on a flat plane, but they are distorted as a consequence. In particular, the points we call the North and South poles appear as lines or circles having the same length as the circumference of the equator.

Are the North and South poles actually lines or circles with the same circumference as the equator?

Clearly not. These pole lines on the map are just "artefacts" of the Mercator projection (and the choice of axis).




Figure 3. Constant gradient and its filtered counterpart




Figure 4. Vertical parabola and its filtered counterpart




Figure 5. Change of gradient and its filtered counterpart

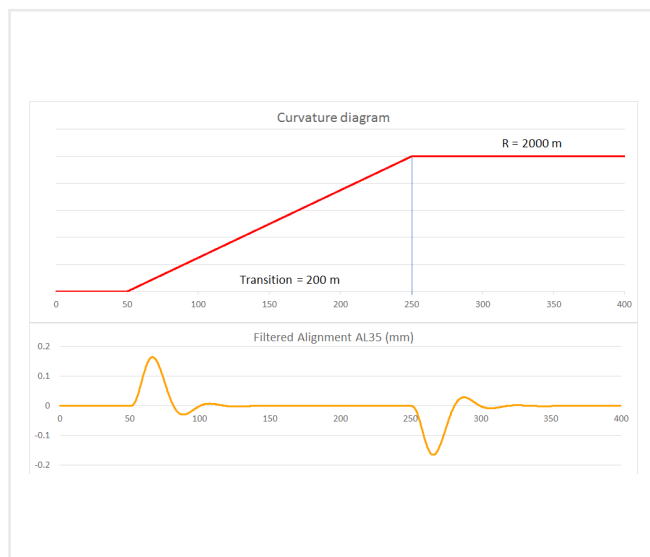


Figure 6. Curvature transition and its filtered counterpart

Similarly, the filtered profiles of the designed track shape are just “artefacts” of the filtering (and the joins between the elements of the geometrical design). The filtered shape isn’t something which is actually present in track having that geometrical design, but it is present in the filtered profile recorded by the TRC when it records such track. And their presence affects the value of SD which is calculated from those filtered profiles.

That is why the recorded SDs of the “as built” track on the renewals sites were not necessarily zero or even in the “good” band, where the site had geometrical designs comprising many of these joins.

As TRC SDs were now part of the spec for track renewals, Network Rail concluded it needed a work around for this apparent

non-compliance, due to the filter and the geometrical design of the track.

SD CALCULATOR

Network Rail asked me to produce software to compute the 1/8 mile SDs that the TRC would produce if it recorded track with a specific designed geometry.

The idea being that the actual recorded SDs would be compared with these, rather than the limits of the “good” band. The result was the Track Geometry SD Calculator (Figure 8) which most track designers will be familiar with.

Essentially the inputs to the calculator are the designed alignment in terms of curvature and the extent of curvature transitions, cant similarly and vertical profile gradients and the

extent of vertical curves, together with the location of a milepost to define the 1/8 mile sections and the direction of the recording car.

The outputs are the SDs for each 1/8 mile of each parameter, as would be produced if the real TRC had recorded track having the precise geometrical design with no irregularities.

To validate the SD Calculator, the inertial and other sensor inputs to the TRC were modelled independently for some sample track geometrical designs at typical recording speed and the resulting sensor data was processed using the actual TRC electronics and data processing software to produce SDs. The results were compared with the output from the SD Calculator when applied to the same geometrical designs and agreement was found to within 1 or 2 percentage points in SD.




Figure 7. Map of the world - Mercator projection

REVISED CONSTRUCTION STANDARDS

With the SD Calculator validated, the construction standards were amended so that instead of the recorded SDs of the “as built” track having to be within the “good” band for track quality, they had to be within a given tolerance of the SDs for the designed track, computed using the SD Calculator. This provided the required “work around”.

CONCLUSIONS AND FOLLOW-ON PAPER

The track profiles recorded by the TRC comprise both the designed shape of the track and defects or irregularities.

The filtering employed on the TRC removes most of the designed shape, but not all. The filter “artefacts” which remain mean that where track has designed changes in curvature, cant or gradient, SDs produced by the TRC will not be zero, even if no defects or irregularities are present.

Except for the special case of track having no changes of curvature, cant or gradient, TRC SDs are not an absolute measure of track quality. In the particular case of new construction, the TG SD Calculator allows engineers to determine the 1/8 mile SDs which would be produced by a TRC recording of perfectly aligned track having a particular design geometry. These values can then be used to compare with actual TRC SDs recorded on the “as built” track.

The SDs produced by the SD Calculator are simply a consequence of the filtering employed in the Track Recording Car. They are not a measure of track quality or of the quality of the design geometry. In a follow-on PWI Journal article to this, Constantin Ciobanu, Principal Track Engineer – Atkins, member of SNC-Lavalin Group, will explain why designers shouldn’t be using the Calculator SDs in this way in Part 2 of this article in the October Journal

REFERENCES

- GC/EH0038 (1993) Track Recording Handbook. British Railway Board
- NR/L2/TRK/2102 (2017) Design and construction of track. Issue 8. Network Rail
- NR/L2/TRK/001/Mod 11 (2013) Inspection and maintenance of Permanent Way. Track geometry – Inspections and minimum actions. Issue 6. Network Rail
- R. Lewis (2011) Track Geometry Recording and Usage. Notes for a lecture to Network Rail.
- D. Marriott (2007) Track Geometry SD Calculator v 1.1. User’s Guide. TGCS David C. Marriott Ltd.

Cant and Curvature			Vertical Geometry	
Distance (m)	Cant (mm)	Radius (m)	Distance (m)	Gradient (%)
0.000	0	STR	0.000	0.600000
10.000			45.000	
40.000			75.000	0.200000
230.000	30	3500.00	310.000	
270.000			350.000	0.250000
480.000	70	1500.00	630.000	
			680.000	0.400000

Track Sections Standard Deviations				
1/8 mile Section	AL35 SD (mm)	WT35 SD (mm)	AL70 SD (mm)	MT70 SD (mm)
10/4	0.230	0.430	2.347	2.558
10/5	0.218	0.537	2.614	2.199
10/6	0.175	0.479	1.620	1.197
10/7	0.188	0.522	1.984	2.309

Figure 8. The Track Geometry SD Calculator (snapshot)