Line Capacity vs. Speed for Same Speed Railways

Summary

This article explains how line capacity is determined, for Same Speed railways, of which High Speed railways are an important category. It is written for an intelligent but non-specialist and perhaps non-technical audience. It deals with an important subject, which is not widely understood. Indeed, insofar as it is known at all, it is widely misunderstood.

The article was originally called simply ‘Line Capacity vs. Speed’. Then it changed to ‘Line Capacity vs. Speed for High Speed Railways’ to stress that it wasn’t concerned at all with conventional, mixed-traffic railways. Now it has changed again, to stress that, except for a very few aspects, (which are highlighted,) the contents apply to **all** categories of **Same Speed** railways, which constitute a whole new world of transport. The individual sections are now described.

Introduction and Background explains what Same Speed Railways are, and how they differ from conventional, mixed-traffic railways.

Line Capacity vs. Speed: the Results contains precisely those. The results are presented in graphical form, with an explanation of why the graph has that precise shape, and what it signifies.

Elucidation of the results is, again, self-explanatory. Most readers will find the capacity values quoted as incredibly high, (and highly incredible). This is frankly admitted, if one’s familiarity with railways is confined to the traditional, conventional, mixed-traffic railway, (as, for the vast majority of people, it inevitably and entirely reasonably will be). The promise is given that later sections will **credibilise** the values.
The section ends with the admission that High Speed is antithetic to High Capacity, and that any attempt to argue that High Speed **promotes** high capacity is completely mistaken if not downright mendacious.

Consequences of the Results This section makes a number of arguments for High Speed which are not primarily concerned with capacity. It even states that there are **technical** reasons for (particular) high speeds, which will be made clear in a later section.
The section further points out that there is a (fairly low) limit to the speed at which divergence from the main line can take place at points. It explains the **basic** and **extended** Train Separation Distance, TSD(b) and TSD(e), standards, how these work and what effect, if any, these have on capacity. (TSD(b) is the normal situation, so has no additional effect, but TSD(e) does have a significant, albeit still small, effect.)
Having introduced these standards, a clear distinction is then made between the different categories – High, Medium and Low – of Same Speed Railways.

Through Stations The effect of stations is now explained, through stations (at which overtaking is possible) first of all.
The **Capacity-Slot Model** is described in superficial detail, but sufficient for present purposes, in particular in the next four sections, where its importance is fundamental and foundational.

The preceding sections, (and the later one on Terminal Stations,) have been part of the article for some time, but have in most cases been significantly updated. The remaining sections are new, at v3.0 and later, the first appearing in late October 2020, though the work that they document began in early 2020, and is still (August 2021) continuing. The content of these sections is presented in considerably more detail than

that of previous sections, because it is highly unlikely that it will be known to anyone. Indeed, the content of these sections is, so far as I can tell, completely new knowledge. What I mean by this is that it is not published anywhere that I have been able to find by web-searching. These are the results of an analysis of the Timetabling of Same Speed Railways, and describe, inter alia, how the fantastic-seeming capacities claimed for Same Speed Railways are actually to be achieved. (The penultimate section is on Resilience, which is a completely different topic.)

Timetabling Considerations and Sweet-Speeds explains how timetables are constructed for Same Speed Railways, which enable the promised capacities actually to be achieved. It further explains how a traffic mix of non-stop and stopping trains can be scheduled, such that the non-stop trains overtake the stoppers, **without any sacrifice whatever of line capacity**.
The section contains several other surprises.
It explains how the fundamental, underlying requirement, which is how to be able to remove a train from the capacity slot stream on the main line, (invariably so as to be able to call at a station,) and later to return it to the main line, coinciding precisely in location and speed with a suitable, empty slot, can actually be achieved. Note that this is possible for **any** line capacity and thus line speed, but that only a very, very few speeds – I call them the **Sweet-Speeds** – enable a viable, usable timetable.
A new feature is added to this section at v4.2, the Station-Calling Diagram. This illustrates and, hopefully, illuminates the station-calling and overtaking process. (This has always been illustrated by a detailed description, which is now, I believe, much enhanced by the accompanying diagrams.)

Timetable Diagrams is a new section at v4.2. It introduces the technique of representing a timetable graphically, providing an example showing how a station is represented, (hopefully) to clarify the Slot Stream Advance concept introduced in the previous section.

Stations on the Main Line: HS-Metros, Pure Metros and Semi-Metros introduces, describes and defines metros. Any Same Speed line (or section thereof) where all trains stop at all stations **is a (pure) metro, irrespective of line speed**. The term HS-Metro is, nonetheless, still used for High Speed lines which operate in this manner, because their appearance, as regards line speeds and station wait times, is so very different from conventional metros, but their underlying properties, in particular the maths involved, are identical. Above all, the underlying principle, described above, of removing a train from the slot stream on the main line, and putting it back again later, applies to **all** station stops, whether or not overtaking is involved. All that overtaking additionally requires is the physical availability of a through track, bypassing the station platforms.
The capacity slot stream is a purely virtual construct, and does not require an actual, physical main line. It continues unaffected, while the physical main line divides and subdivides into a number of platform tracks, and all trains stop there. (The main line re-constitutes itself and re-joins the slot stream later, you may like to imagine it.)
Semi-metros are simply metros which allow (some) overtaking.

Change of Line Speed, for a Semi-Metro Metro routes typically consist of a number of sections, with a central core of pure metro, where all trains stop at all stations, then Low Speed semi-metro sections beyond the core in both directions, over which some trains travel non-stop, to become stopping trains again in the outer sections, which are Medium Speed semi-metros. The line speed actually changes at a station where all trains stop, the stopping trains of the inner, Low Speed section often terminate at this station, and the continued service is taken over by the former non-stop trains.
This section explains how the change of line speed is scheduled.

Change of Line Speed for a High (or Medium) Speed Railway this section has been removed, as it seriously unbalanced the present article. It now appears in the new article 'Line Capacity vs. Speed for Same Speed Railways – Volume 2'.

Terminal Stations argues that terminal stations (like, for example, Euston,) are a **very bad idea** for High Speed Railways, except at a natural terminus, (like at the coast, where if you go any further you're in the sea). It is predictably and justifiably scathing about HS2 Ltd.'s insane proposals for Euston.
It proposes a long-term, mixed-function role for existing termini, the future, rail-only needs for which will be restricted to rush-hours.

Resilience is a very important topic, which is usually forgotten about. The resilience problems of metro and semi-metro systems are very different, in their manifestations and also in the problems for which solutions can realistically be attempted. Proposed solutions are offered for each type.

Conclusions wraps it all up, summarising what has been achieved already, and what needs to be done next.

Contents — *Page*

Introduction and Background

The Same Speed Railway is a theoretically ideal model of a railway with certain properties and characteristics such that its behaviour and performance can be predicted and mathematically analysed. Various types of physical railway do or can be made to approximate to the ideal, with very desirable consequences. I recognise three categories of Same Speed Railway:

- High Speed Railways
- Medium Speed Railways
- Low Speed Railways

An alternative classification is between (pure) Metros (where all trains stop at all stations) and Semi-Metros, which allow a mixture of stopping and non-stop trains, with overtaking at stations. Metros and semi-metros are available in all three speed ranges.

Metro systems have been around for a long time, and their performance is well understood. High Speed railways are now very familiar, but it is my contention that their performance is **not** well understood. The Medium Speed range is foreseen as for conversion of classic routes to the Same Speed model, though there are none as yet, and also for the outer ranges of metro systems. The model would also be suitable for dedicated freight lines, though, again, there are none as yet.

The fundamental property of a Same Speed Railway is that the traffic is **homogeneous** in performance, consisting of trains whose dynamic performance, specifically their speed range and acceleration and deceleration rates, are identical, within a very narrow range. These trains all travel on the line at the same speed**, (**hence the name,) the **line speed**. (Strictly speaking, the trains do not need to be identical, but their dynamic characteristics must be such that they are all able to **perform** identically, as regards acceleration, deceleration and speed. This allows new trains, of improved performance, to be introduced. Initially they will run at the same performance as the existing trains, until these have been progressively withdrawn, when the improved characteristics of the new stock can be taken advantage of.)

This homogeneity of performance has a most important consequence: Same Speed Railways **do** have a precise value of line capacity, for a particular line speed: this value varies as the line speed itself varies, and this can be analysed mathematically and its characteristics deduced.

The precise meaning of line capacity is how many trains the line can accommodate, specifically how many trains can pass a particular point in a particular time (in a particular direction!), generally stated as trains per hour (tph). The faster they're going, the more get past, right? Wrong! What matters is not how fast you can go, but how fast you can (in a controlled manner) stop. Also, while it is true that line capacity increases linearly with speed at very low values, this quickly ceases to be the case as the speed increases further. A maximum value of line capacity is reached, at a surprisingly low speed, and thereafter, further increase in speed causes causes line capacity to decrease, albeit slowly.

Line capacity is a rather amorphous concept for conventional, mixed-traffic lines. The idea is clear enough, but such a line has no fixed value of line capacity. The value is heavily dependent on the traffic mix, different types of train, of different levels of performance, travelling at different and varying speeds. This can change several times every day, as the traffic-mix changes. Railwaymen over many generations have developed reliable but essentially rule-of-thumb methods for determining what traffic can be scheduled on the line. But the principles underlying Same Speed Railways ensure that line capacity is a **property**, with a specific, definite value, dependent only on the line speed. (There **are** other influences, most particularly acceleration and deceleration rates, but these are not readily varied, and may be strictly limited by passenger tolerance.)

While the line speed (and therefore the line capacity) of a Same Speed Railway has been stated above to be constant, strictly speaking, this is true only for individual sections. It is perfectly possible for adjacent sections, (separated by some feature, typically a station at which at least some, possibly all, trains stop, but it could equally well be a junction, or a permanent speed restriction, such as across the Forth Bridge,) to have different values of line speed. The important point is that line speed is fixed within a section, but the value may undergo a step-change between sections. The overall line capacity is that of the lowest capacity section, which, for High and Medium Speed Railways, will be that or those with the highest value of line speed.

The model assumes constant acceleration and deceleration rates. Representing the real-world acceleration and deceleration by an equivalent straight-line approximation seems to be a standard technique in railway engineering. It certainly simplifies the calculations! I go along with it, since the whole point of a model (in any field) is that it focuses on the essential elements, and allows clear results to be derived readily, without being complicated by inessentials.

A same speed, non-stop train accelerates, at a prescribed rate, from its originating station up to the line speed, then travels the entire journey at that speed, until it decelerates, at a prescribed rate, to a stop at its destination. This behaviour applies similarly for the sections between intermediate stations, for stopping trains. All trains take exactly the same time for a specific journey, and for all sections thereof. This mode of operation has the coincidental benefit of making it very easy to calculate journey times. Note that I say ‘calculate’ rather than ‘estimate’. An estimate is a simplified calculation, which omits certain less important features – a quick and dirty calculation in fact. But these are precise and exact results for the Same Speed Railway model. They become estimates only when quoted for (usually) a High Speed line, whose realisation of that ideal is at best only approximate (but they are still **very good** estimates nonetheless).

The main reason for analysing line capacity is to find out how to optimise it, for some prescribed objective. Line capacities of Same Speed lines are very much higher than for mixed traffic lines of the same (maximum) line speed; this has nothing **directly** to do with the speed, but is purely due to the traffic homogeneity. However the actual line capacity value of a given Same Speed line is very much dependent on the line speed, and in a rather surprising, (at least, it surprised me,) perhaps counter-intuitive, way. The present article gives those results, with a little background and some further elucidation, but all the details, the assumptions, the calculations and the real-world numerical values, are contained in Appendix B of the article ‘Same Speed Railways’, published on my website – www.croal.uk – and freely downloadable. The intelligent but non-specialist reader does not need this stuff to understand the results, but the techno-freaks, the people who just **have to have** the really hard stuff, should love it. The information on which my calculations are based comes from web articles published by Piers Connor, of PRC Rail Consulting Ltd., which are gratefully acknowledged; the above article gives the references.

The (basic) Train Separation Distance – TSD(b) – is the minimum distance which must be maintained between adjacent trains. It is equal to the stopping distance, the distance required to decelerate from line speed to a standstill, (which is of course the same for all trains, under the fundamental assumption of homogeneity), plus a safety margin. Note that this is deceleration under normal service conditions, not an emergency stop. The idea is that a train has adequate distance to come to a complete standstill, without crashing into the preceding train, should that preceding train suffer a catastrophic accident, and come to a stop effectively instantaneously.

The reader may reasonably ask how the following train would know. This is of course an idealisation. It is assumed that every train has continuous knowledge of its own instantaneous position and distance behind the preceding train. This is the province of Automatic Train Control and the Digital Railway, still in its early stages, but global positioning systems (or a track circuit approach) can already keep track of the location of each train, and it is simply(!) a matter of calculating the distances, and communicating them to all the trains in real time. The technology is either already available or soon will be. However, rest assured that the results derived, even in such a theoretical context, have real-world, current relevance.

(What I am envisaging here is the fully automatic railway – driverless trains and all. To those who throw up their hands in exasperation at such airy-fairy foolishness, I would just point out that, mere months ago, the idea of driverless cars would have been regarded as swivel-eyed lunacy, yet now the idea is main-line and scarcely excites comment. Serious resources in research and development are being devoted to it, and it will happen, in some form (though whether it will fully realise the enthusiasms of its protagonists is yet to be seen). Driverless trains is a much simpler concept: they don't have to find their way anywhere; their routes are fully prescribed and rigidly enforced. The present article could be viewed as the Railway's answer to driverless cars – high speeds and huge capacities.)

TSD(b) has a dynamic and a static component. The dynamic component is readily calculated, knowing the average deceleration rate. The composition of the static component will be defined in the section 'Consequences of the Results', since it is intimately bound up with the subject of High Speed switches / point-work. Suffice it for now that there is nothing arbitrary about it; it is **not** a rule-of-thumb.

The calculation of line capacity is simple: it is equal to line speed (m/s) divided by the Train Separation Distance (m) and the result then multiplied by 3600 to convert from trains per second to the usual trains per hour.

Note that in all measurements involving one (or two) train(s), the reference point is always the **front** of the train. The capacity is readily calculated (I use Microsoft Excel©) for a series of line speeds, and the results plotted automatically. (Further clarifications are added with M/S Paint©.) The graph of capacity against line speed follows.

Line Capacity vs. Speed: the Results

[A word of explanation for readers who may not be familiar with Excel© charts:
There are many types of Excel© chart, suitable for different types of data distribution. That used here is a **line chart**, which is similar to the familiar line graph, except for one, very noticeable difference. The normal line graph displays the independent variable along the x-axis, together with its range of numerical values. The dependent variable(s) are displayed as continuous, smoothed lines on the graph, with their (common) range of values along the y-axis. But on a line chart, **all** variables **including the independent variable** are displayed as lines on the chart. The annotation along the x-axis denotes **sets** of values, one of which (in each set) will be for the independent variable, and all the others are derived from this. In other words the numbers along the x-axis are not, themselves, numerical values but labels (of sets of numerical values). In the following charts, the very prominent straight blue line is the graph of line speed, which is the independent variable, and its range of values is along the y-axis, the same as for all the others. (Note that the graphs of all the **dependent** variables have exactly the same shape as they would have in the traditional line graph; only the way in which the independent variable is shown differs. Unsurprisingly,

the charts are commonly – by me too – also referred to as graphs, although, strictly speaking, they aren't. You quickly get used to it.)

The results are calculated for a series of line speeds, with a constant increment, of 2.5m/s in the present case. The numbers along the x-axis are simply the sets of values in the sequence, thus 1 is the set of values for line speed 0m/s, 2 for 2.5m/s, 3 for 5m/s and so on. All the results are on the ordinate (y-) axis. So, to read the chart, use a ruler, parallel to the y axis, and take the values, from the y-axis, where the ruler intersects the two lines. The process has (I hope) been clarified by the addition of three very important capacity / speed correlations.

(Note that I frequently deliberately adopt rather esoteric units, particularly for the independent variable. Thus, for the Capacity / Line Speed graph, the unit of speed is 2m/s. This is purely to get an improved display for the other values, allowing them to spread themselves, rather than being cramped at the bottom of the display.)]

Persons with an engineering background will be familiar with graphs of this type, but most non-engineers will probably never have encountered them. So, with absolutely no condescension, some further elucidation is offered which I hope will be helpful, and explain precisely why the graph has that shape.

The line capacity depends directly on speed and inversely on the Train Separation Distance, TSD, (the several variants of which, basic, extended and mixed will be explained shortly) The separation has, as explained above, a fixed component and a variable component, which latter itself depends on the square on the speed. At low speeds, the fixed component is dominant, so the capacity increases linearly with speed. But as speed increases, the variable component becomes dominant, and thereafter the capacity becomes progressively inversely proportional to speed. The speed of maximum capacity is where the two influences are in balance

A further graph will elucidate:

This second graph adds two further quantities, the capacity value determined only by the fixed component ('line capacity 830' – the odd name alludes to the fixed component of the separation, in the current treatment, having the numerical value 830m), and the capacity value determined only by the variable component ('line capacity var').

This clearly demonstrates that at the lowest speeds capacity increases linearly with speed. But the effect of the variable component quickly becomes evident, and, at high speeds, completely dominant. It is clear that, at the highest values shown, the capacity with basic separation is merging with the inverse graph. (In the graph, the point of intersection of the '830' and 'var' graphs is, clearly, vertically above the net capacity maximum, but if the vertical line through the two points were omitted, the point of maximum capacity would appear well to the left of the point of intersection – an optical illusion caused, I imagine, by the different gradients each side of the maximum, making the point of maximum seem further to the left than it actually is.)

(A further personal style-point is that the two graphs are displayed on even / odd pages. This is so that, in a printed version of the article, they can both be viewed together.)

I hope that is helpful. (It was certainly fun to produce!)

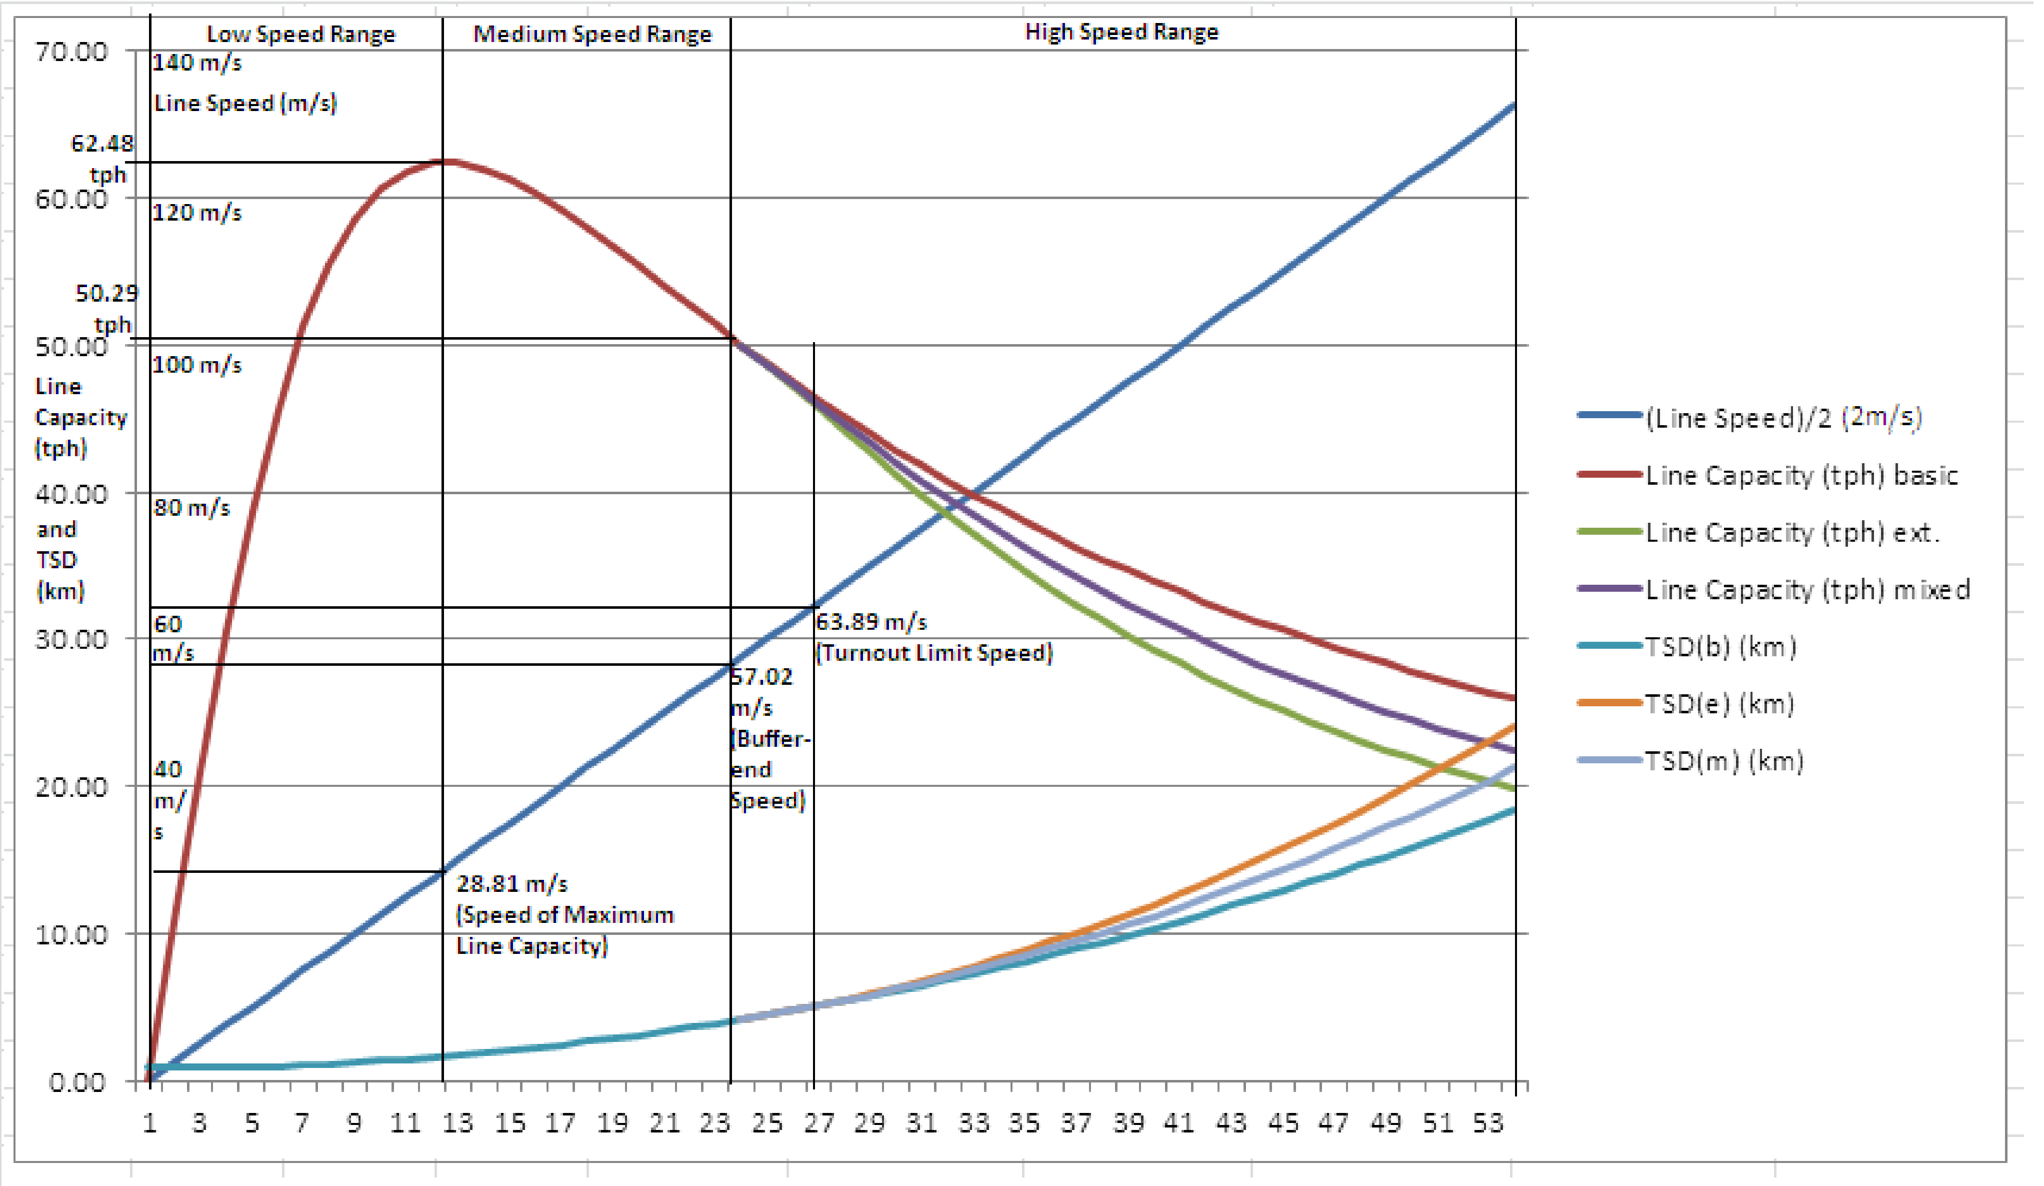

Line Capacity vs Line Speed, together with TSDs, indicating the several Speed Ranges and highlighting certain Speeds of Particular Significance

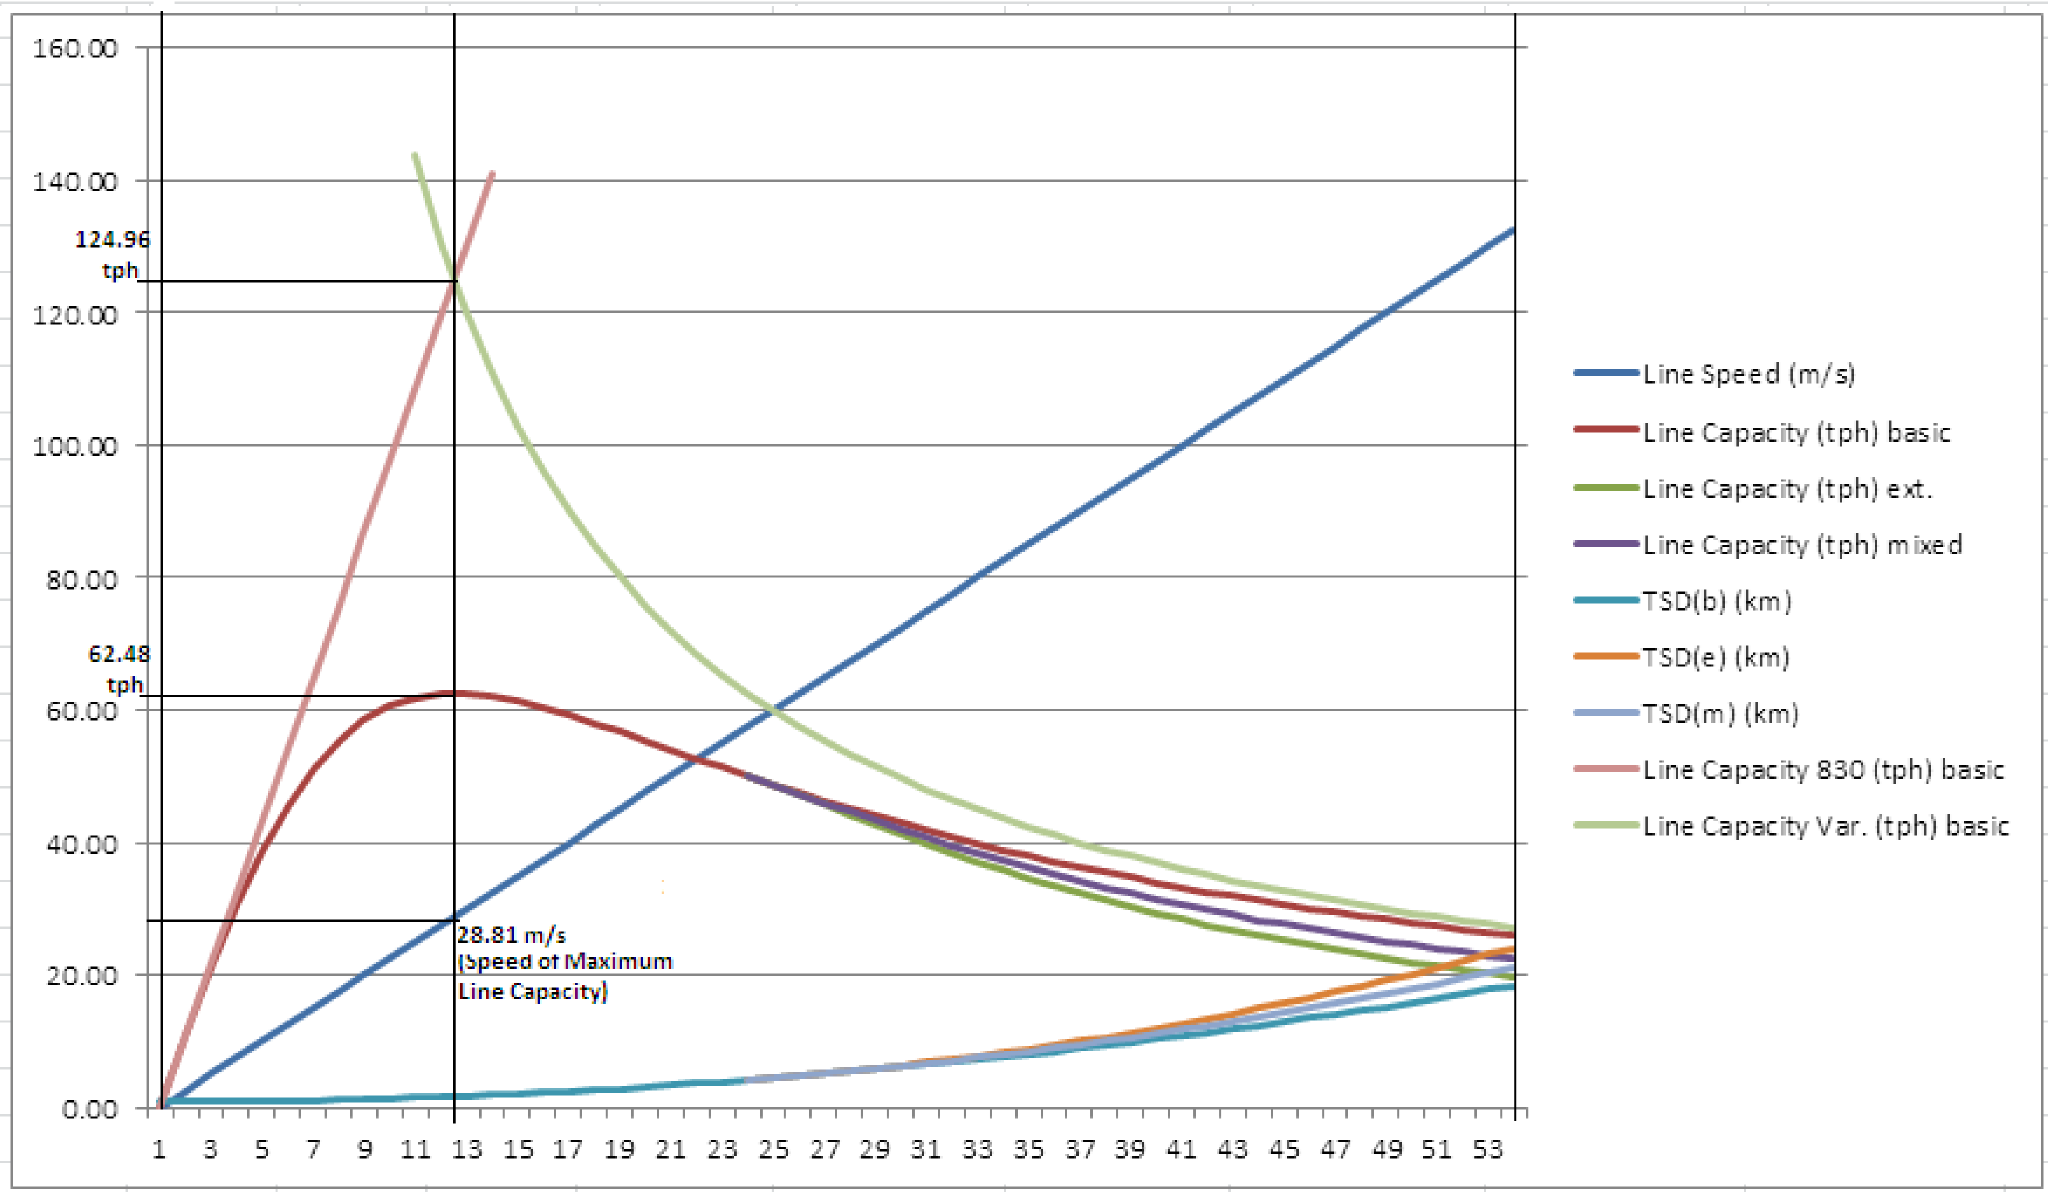

Line Capacity vs Line Speed, with added Constant and Variable Components of Line Capacity displayed Separately.

Elucidation of the Results

The immediate reaction of most readers will be that the capacities indicated are way above any occurring in reality. That is true, at present. It is important to appreciate exactly what is being demonstrated. These capacities are theoretical values, determined **only** by the TSD(b), and represent ideal maxima, thus if TSD(b) were **all** that determined capacity, then these are the maximum values possible, there is no way to get more. They are an ideal, assuming, inter alia, absolute precision in timekeeping and total adherence to the timetable. This could only be achieved by completely automatic train control; there is no way that human control could provide the necessary precision. These capacities are real values, but they assume a degree of perfection in operation which is not available – yet. They thus represent an ideal we should aim for, even if it is as yet beyond our abilities. While such a level of perfection is not yet attainable, values of around 50% of these ideal capacities **are** a reasonable goal to aim for, right now. The best modern metro systems already achieve 50% of the theoretical capacity; the Victoria line, for example, has been performing at a peak of 34tph for the past few years. Crossrail is designed to deliver 24tph through the central core initially, later rising to 32tph. It is not unreasonable to aim for capacities of around 30tph for HS lines also.

The above expresses my earlier opinions, but I now regard it as quite unnecessarily defensive. Yes, the capacities shown do appear preposterous; that is perfectly correct, if the reality that you have in mind is that of the traditional, mixed traffic railway, which is, of course, what everyone is familiar with. Same Speed Railways are something else entirely; a new paradigm. The present article explains what they are, derives their properties and explains how these results are in fact to be achieved. I assure everyone that the capacity values just presented are real, and achievable, once we have got train operation fully automated. (That’s all!)

In practice, all sorts of extraneous factors also affect capacity, some reducing it markedly. Junctions are one such, and will be dealt with shortly, likewise through stations. Terminal stations are the most destructive of all, imposing a limit of 2tph per available platform face, on the assumption, as in the infamous plans for the redevelopment of Euston, ridiculously as a terminus, of 20 minutes to unload, service and reload a train, with 10 minutes contingency.

The capacity graph has a very familiar shape with, at low speeds, capacity increasing rapidly with speed, until it reaches its theoretical maximum, 62.48tph at the astonishingly low (to me!) speed of only 28.81m/s / 103.72kph / 64.48mph. (This may not be so surprising to metro operators, whose trains perform in precisely this speed range.) The maximum **practical** line capacity is of course 60tph, at 38.37m/s / 138.12kph / 85.83mph, which is still pretty good, Thereafter the capacity gradually decreases as line speed increases, the rate of decrease itself decreasing as speed increases.

The results may be theoretical in their numerical values, but not in their qualitative properties; the shape of the graph will remain the same, stretched or compressed, shifted up or down or sideways a bit, but still essentially the same graph.

And here is the takeaway: justification of **High Speed** lines (as opposed to Same Speed lines in general,) by arguments based on line capacity is completely wrong. They are high-capacity only in comparison with mixed-traffic lines of the same (maximum) line speed (whose capacity levels, as is well known, are rubbish). Above the low speed of maximum capacity, the faster you go, the less capacity you have. A line speed of 100mph, 45m/s, has **twice** the (theoretical) capacity, 57tph, of a line speed of 225mph, 100m/s, which is 28tph. You may argue over the precise values, but not over their relative magnitudes.

It may possibly be suggested that, since HS2 is being built to the GC standard loading gauge (for historical reasons which have never been properly challenged or justified), it will allow much larger trains, including double deckers, to run, so that its **passenger-carrying capacity** will be much larger. This is true, but a complete red herring: since **twice** as many GC-gauge double deckers (or any other type of train of similar dynamic characteristics) could be accommodated with a line speed of 100mph as with one of 225mph. **Line capacity** – the number of **trains** per hour – is what matters, and here high speed has nothing to offer; it is in fact detrimental.

The justification for high speed lines is therefore just that, that they enable you to go faster, and arrive at your destination sooner, and that's it.

Consequences of the Results

(It's strictly nothing to do with the present article, but an intriguing possibility suggests itself. Should some national emergency require an immediate increase in transport capacity, a high speed line could provide it – immediately – simply by reducing speed.)

Of course, no-one is going to be happy with travelling long distances at just under 86 mph, and will certainly not be consoled by the thought that by so doing, another 59 trains per hour are able to share the track with theirs. There are sound justifications for high(er) speeds, but they are business and commercial reasons, not technical ones. It is necessary to strike a balance between the benefits passengers perceive from high speed and what they are prepared to pay for it (remembering that, since power consumption varies essentially as the square on speed, a speed of 225mph has **five and a half times** the power consumption of 100mph (and 250mph over six times); it is thus **very** much more expensive to provide, even after all the new infrastructure is in place).

The above paragraph expresses what were my beliefs until very recently, and these are not wholly correct, in that I have now discovered actual **technical** reasons for (particular) high speeds. This will all be explained in the section 'Timetabling Considerations and Sweet-Speeds'.

We have considered so far only the basic separation, TSD(b), which is the absolute minimum separation which must, no matter what, be maintained between trains. On the assumption, as before, that a train 'knows', at all times, how far ahead of it the preceding train is, (should be TSD(b), of course,) then, should that train begin to decelerate for some reason, then our train is immediately aware that it too must decelerate, to maintain the TSD(b) (which itself decreases as speed decreases, of course, so the trains get closer together). Likewise the following train, and the one after that and … This is clearly ridiculous.

For high speed lines, their Achilles heel is switches or point-work. There are no switches available which allow a diverging train to do so at line speed. The fastest currently available allow for a (maximum) **turnout limit speed** (i.e. diverging; converging is presumably the same though I have never seen this discussed,) of 230kph, 143.8mph. A diverging train must therefore decelerate down to 230kph **on the main line**, by the time it reaches the switch. (Trains continuing straight ahead simply continue at line speed.) In order that a diverging train does not delay the following, straight-ahead train, an **Extended Train Separation Distance, TSD(e)** is proposed, such that, as the diverging train decelerates, the following train gets nearer to it, but only reaches the TSD(b), (the irreducible minimum, remember,) at the point where the diverging train has just completely diverged at the junction, so it is no longer in the path of the following train.

‘Completely diverged’ means three things:

- The back-end of the train has cleared the junction.
- The back-end of the train has, in addition, cleared all moving parts of the switch.
- The switch has been reset to the main line, i.e. directly ahead.

These three items together constitute the buffer length, i.e. the constant component of both TSD(b) and TSD(e). (TSD(e) has an extra, speed-dependent, component to cover the in-line deceleration of diverging trains.)

My information on Very High Speed Switches comes from Vossloh Cogifer, who have over 40 years experience in developing and manufacturing this equipment; they were the suppliers for the original Paris-Lyon line in 1980, and have since then supplied most of Europe, including the UK, and elsewhere. (In addition, one of their switches featured in the world rail speed record, of 574.8kph on the LGV Est on 3rd April 2007, when it was crossed at 560kph – on the direct track, presumably!). Their latest offering, the Swing Nose Crossing with Manganese Cradle, allows for a maximum of 230kph on the diverging track, and a normal 350kph operating speed on the direct track.

I have spoken by telephone and email with Vossloh Cogifer UK, and wish to record my appreciation of how helpful they were, in particular David Walters, in supplying me with the information I needed.

Values for the above three items are:

- Train length 400m (16 or 2*8 vehicles of length 25m each).
- Moving switch parts length 194.5m.
- Time to reset switch 4sec. This is an engineering standard for swing-nose switches, and so may be regarded as a maximum. It results in a distance travelled by the train, while the switch is operating behind it, of, at most, 232.2m.

– giving a total of 826.7m, which I round up to **buffer length b = 830m**.

That is the (minimum) distance that (the front end of) a diverging train, decelerating at the standard, constant rate throughout, must travel beyond the switch points to be completely out of the path of a following, straight-ahead train. It must of course have already decelerated down to the turnout limit speed of 230kph on reaching the point of divergence, i.e. the switch points.

In the normal case, for stopping at an intermediate station, the diverging train decelerates at constant rate from line speed to stopping at the station platform. In particular, it continues its steady deceleration across the switch. The (instantaneous) speed reached at the end of the buffer length beyond the switch points, i.e. at the point where it is exactly TSD(b) ahead of the following, non-stopping train, and at the instant when the connection is broken by the switch re-engaging with the main line, is 57.02m/s / 205.29kph / 127.48mph. Like the Turnout Limit Speed, this value is constant, irrespective of the line speed. I call it the **Buffer-end Speed**. It is indicated on the Capacity / Line Speed graph on p.8, and is the boundary between the High and Medium Speed ranges, **not** the TLS, as you might, (and as I originally did,) expect.

To clarify: TSD(e) specifies the performance when a diverging train performs **any** of its deceleration in the path of the following non-diverging train. This connection is broken **only** when the switch is reset to main line, specifically it persists while the diverging train travels the initial buffer length along the diverging track. In the limiting case, when the actual line speed is precisely the same as the buffer-end speed, both trains travel in lock-step at that speed, precisely TSD(b) apart, until the connection is broken by the switch, and only at that point does the diverging (actually, diverged) train begin its deceleration. At this speed, TSD(e) and TSD(b) are equal. (Note that both these values always refer to the line speed, since it's the distance of the non-diverging train behind the diverging one that is at issue, and that is travelling at line speed throughout.)

As stated above, this relates to the normal case, where a train diverges to call at an intermediate station. But there is another possibility, where a train diverges permanently, onto another route. Here, it decelerates on the main line down to the TLS, as before, but then diverges **at that speed**, since there's no need to decelerate further, indeed, it begins its re-acceleration back up to line speed on the new route at the earliest possible moment, which may even be as soon as the back end of the train has cleared the moving parts of the switch. It may be pointed out that this (traversing the switch at constant speed) could also apply to the stopping case. That's true, but there would be no point – it would save less than 1sec crossing the switch, but increase the length of the station loop by 830m, since the deceleration would still have to be performed.

I've dealt with this matter exhaustively because of its fundamental importance, and because it can be seriously confusing.

The graph of capacity splits into three strands at the higher speeds. The buffer-end speed is the value at which the split takes place, The lowest of the three describes the behaviour when TSD(e) is the distance maintained between each pair of adjacent trains, when they are both travelling at full line speed. (Note that the capacity quoted earlier for line speed 225mph, 28tph, is for TSD(e), as that would be the standard used at that speed.)

The final refinement is to recognise that TSD(e) is slightly pessimistic (no bad thing, of course). It is only really required when a diverging train is followed by a straight-ahead train. When two trains are both straight-ahead, or the first is straight ahead and the second diverging then the distance between them only actually needs to be TSD(b). The actual worst case, as far as capacity is concerned, is when trains are alternately diverging and straight ahead. This actual worst case capacity, ('worst' because it requires the maximum proportion – 50% – of TSD(e) separations; any other traffic mix would require fewer, and so have a slightly higher net capacity,) is depicted in the middle strand of three on the graph, termed 'mixed' separation. This is actually of purely theoretical interest, since, in practice, we would always choose TSD(e) throughout, because there's no point introducing massive complications for marginal gains. We're always content with slightly pessimistic standards; it's the slightly optimistic ones that tend to bring unpleasant surprises. These variations in train separation naturally only apply at line speeds exceeding the buffer-end speed of 57.02m/s, and are certainly significant, but not huge. All this stuff is expounded at length in the 'Same Speed Railways' article, which also demonstrates that the case of trains **joining** the main line at a junction is also completely covered by TSD(e), which is thus a very good, conservative standard.

(Really-wide-awake readers might wonder what would happen if a diverging train were followed by a second diverging train. That is never allowed to happen. Indeed it cannot, since the whole point of the TSD(e) standard is that the diverging train gets out of the path of the straight-ahead train in a timely fashion; that's why it works. The first, diverging train has decelerated to the turnout limit speed when it reaches the junction, and has decelerated a little further when it has completely diverged at the junction, at which point it is precisely TSD(b), the absolute minimum, ahead of the following train, which is still travelling at line speed. If this following train were also diverging, then the first train would still be in its path, and it would continue to get closer to it, which cannot be allowed. In the unlikely event that it were required to schedule two diverging trains in succession, then they must be separated by at least one empty capacity slot – see the next section on Through Stations for an overview of the Capacity Slot model. In effect this fits a phantom straight-ahead train between them.)

As far as the current article is concerned, the only point of describing most of the above effects, (some of which are decidedly esoteric,) is to reassure the reader that they have been taken into account. (And also to pre-empt questions if readers wish to perform their own checking.)

Now that the Extended Train Separation Distance standard has been introduced, we can define the distinction between the categories of High Speed and Medium Speed railways (as was alluded to when the Buffer-end speed was introduced).. High Speed Railways are those for which TSD(e) applies. Their line speeds are thus above the Buffer-end Speed, (57.02m/s / 205.29kph / 127.48.8mph in the present treatment), and thus diverging trains, generally to stop at a station, must decelerate down to the Turnout Lmit Speed, (63.89m/s / 230kph / 143.75mph) on the main line by the time they reach the switch-points. Medium Speed Railways are those whose line speeds are in the range between 28.81m/s / 103.71kph / 64.82mph – the speed at which line capacity is at its (theoretical) maximum of 62.48tph – and the buffer-end speed.

For Medium Speed Railways, the basic Trains Separation Distance, TSD(b), standard applies. Note that the prohibition of adjacent diverging trains applies not just to High Speed Trains, but for Medium Speed Trains also, and for the same reason as stated above. The second train diverging would be too near to the previous one and also travelling too fast, irrespective of whether or not it had done any prior deceleration on the main line. Note also that, for High Speed Railways, the length of the station loop is the same in all cases, since it depends on the turnout limit speed. But for Medium and Low Speed (see below) railways, it is determined by the **line** speed, and thus varies (and it is always less than for the High Speed case).

Just for completeness, Low Speed Railways are those with line speeds below the speed of maximum line capacity. For this category, adjacent, diverging trains **are** permissible. For this speed range, the deceleration time from line speed down to zero on the (physical – same as virtual for Medium and Low Speed) station loop is less than the capacity slot time, (see next section,) which is the time taken to travel a distance of the capacity slot size, which is equal to TSD(b), and thus the distance behind of the following train at the time that the diverging train diverged. It is easily demonstrated that this is true for all line speeds less than that of maximum line capacity, and this of course coincides with the Low Speed range. This is an important result, as will be explained in the section ‘Stations on the Main Line: HS-Metros, Pure Metros and Semi-Metros’, but is mentioned here to give a full account of the distinction between the categories of Same Speed Railway.

Referring back to the capacity graph in the section ‘Line Capacity vs. Speed: the Results’, the correspondences are clearly indicated.

We next consider the effect of stations, which are very different for through and for terminal stations.

Through Stations, with Overtaking

The treatment of through stations on HS lines depends on whether all trains stop there, or only some of them. If only some of them stop, then there must be provision for non-stop trains to overtake stoppers. In this situation, the platform lines are long loops off the main line, and the stations have no effect on trains which do not stop there, since a diverging train (to stop at the station) has no effect on a following non-stopping train, as it gets out of its way in a timely fashion, as described above. Any train in steady motion occupies one capacity slot. If a train stops at an intermediate station, it gives up that capacity slot, and requires another one to be available for it to occupy when it restarts. These are, of course, capacity slots **on the main line**.

The Capacity-Slot model is explained in detail in the 'Same Speed Railways' article. Very briefly, it envisages a continuous stream of slots, each of length TSD(e), travelling along the main line at constant line speed. Each train, when travelling on the main line, occupies a single capacity slot. A train which travels non-stop between origin and destination occupies the same slot throughout, and requires only that one slot for the entire journey. A train which stops at intermediate stations gives up its slot when it diverges from the main line onto the station loop, and obtains a new slot when it re-joins the main line after calling at the station. Thus if it makes n intermediate station stops, it uses $n+1$ slots in total, albeit only one at a time. The slot given up when diverging for a station stop immediately becomes available for re-use by another train, either joining the main line, (from another route,) or re-joining the main line after calling at a later (or possibly even the same) station. **It is always possible, in principle, for a released slot to be re-used later**, potentially several times. (The full slot theory considers precisely how trains join and leave slots, how their position within the slot varies, and how slots come into being at the start and disappear at the final destination. It also explains the **Slot Window**, which is the time / distance range behind the train in the preceding slot during which a train must join its new slot, to be able to reach its prescribed position within the slot. All terrific stuff, and necessary to demonstrate the rigour of the theory, but absolutely not needed to understand its implications in the present context.)

The problem here is that, at the time a train wishes to restart from an intermediate station, a free slot may not immediately be available for it, and it must therefore wait (i.e. delay its departure from the station) for the next free slot. It may well be, if the main-line loading is high, that several capacity slots in a row are occupied, before the next free slot occurs. Given a slot time of c.2 minutes, that could impose a severe time penalty on a station stop, in addition to the unavoidable c.7 minutes (this is the time penalty for decelerating to a standstill at the station, a standard wait time of 3 minutes, then accelerating back up to line speed, as compared with travelling the same distance at full line speed). So, while this model **will always work** – the capacity is still there, though the dynamic distribution of it may not be optimal – for **good** performance, it requires some very neat scheduling, and this may not always be practicable. This scheduling has two aspects:

1. to draw up the optimum timetable, so that the (dynamic) slot distribution in normal service minimises the (probably unavoidable) extra time penalties, and
2. to perform dynamic re-scheduling in real time, in particular, when a train, through lax operating performance or following an unavoidable incident, misses its scheduled slot.

These matters are addressed rigorously in the next section. The present section was written long before the theory required to deal with them was available. The above introduction to the Capacity-Slot Model is retained, as that is still valid and very relevant to what follows, but the remainder has been omitted, as outdated..

Timetabling Considerations and Sweet-Speeds.

This has been a seriously difficult subject to elucidate, but, having finally solved it, the results are astonishingly simple and easy to explain and understand. That is indeed remarkable, since I believe that this has been the most intellectually challenging problem that I have ever attempted, requiring clarity of thinking to an extreme degree. I will share with you what I believe was its most intractable difficulty.

The treatment here is essentially the same as that in the 'Same Speed Railways' article, except that that article contains the details of the calculations and spreadsheets of the numerical results, whereas the present article contains extra commentary and some personal remarks. I believe that the argument is sufficiently straightforward for a non-specialist audience / readership, whose appreciation will be enhanced by the extra detail.

The fundamental aspect of the problem is that it obliges consideration of timetabling, of how the various services are to be scheduled. So far, all calculations have started with the line speed as the independent variable (which means it's the one to which values are assigned, and all the results are derived from that, simply by inserting the relevant speeds into formulae). But timetabling requires us to look at things from the other end, starting from the capacity, strictly the inter-train times, but these relate directly to capacity. Although line speed is still the independent variable, it is now itself also the desired result. Solution of this problem has to be numeric, by an iterative process. Effectively this involves trying different values of line speed until we get the desired capacity value, to any desired degree of precision. (This is a well-known but tedious and time-consuming process, but spreadsheets are a terrific help in performing it.)

Having spent a lot of time and effort deriving the results in this manner, I then realised that I could indeed start from the other end, with line capacity as the independent variable, and obtain the corresponding line speed from that, simply by solving a quadratic equation (admittedly a very complicated and nasty one).

The idea occurred to me some time ago that it could, in theory, be possible to fill the capacity slot left vacant when a stopping train diverges at a station, by another train which had stopped earlier at that station, when that empty slot reached the other end of the station loop. I say 'station loop', but, strictly speaking, (for the High Speed category only,) it is that section of track on the main line from the location where a stopping train starts to decelerate, before it physically diverges onto the station loop, to that location where it finally reaches line speed, and its correct position within its new capacity slot, after physically re-joining the main line from the station loop (into its new slot but not in its final position within the slot) and completes its acceleration up to line speed, on the main line. (All of this is explained in exhaustive detail in the 'Same Speed Railways' article, where the Capacity Slot model is defined, but that amount of detail is quite unnecessary to understand the results.) From now on, I shall call this the **virtual** station loop, and the physical station loop, which diverges from the main line, and on which the station platforms are physically located, I shall call the physical station loop, explicitly. So there can no longer be any ambiguity about which I mean. (For the Medium and Low Speed categories, the virtual and physical station loops are the same.) The main implication of this idea is that it should, theoretically at least, be possible to operate a mixture of non-stop and stopping trains without any loss of line capacity, a very desirable outcome. Note that the distance is the same, going via the physical station loop or staying on the main line throughout.

It is, in fact, entirely possible. More than that, it is the **only** sensible and indeed the only practicable way to operate a Same Speed Railway (even a metro – see the section 'Stations on the Main Line: HS-Metros, Pure Metros and Semi-Metros'). Indeed, the original limited and restricted, and even rather boring topic of investigation, seemed to take control of me, and widen itself spontaneously into the entire, fundamental subject of the Timetabling of Same Speed Railways, and became, as I have said, the most intellectually challenging problem I have ever attempted

– to my entire surprise. I believe it's called serendipity.

In order to be able to merge the timings of non-stop and stopping services, we need to have an integer multiple of capacity slots per hour to make the construction of a usable timetable possible at all. If this were not so, there would be no regular framework constant over time; the environment would change from hour to hour. The hourly number of capacity slots is of course simply the line capacity in trains per hour. So all that this is really saying is that the line capacity in tph must be an integer. For **any** line capacity, it is possible to determine the corresponding capacity slot time in seconds, simply by dividing 3600 by the line capacity in tph. So this condition, while clearly necessary, is nowhere near sufficient. But before determining what would make it sufficient, it is necessary to consider the details of the process by which a stopping train is overtaken.

The fundamental principle involved with stopping trains is:
Accelerating from zero, with an arbitrary but uniform rate of acceleration, up to an arbitrary speed, and then immediately decelerating back down to zero, with an arbitrary but uniform rate of deceleration, takes precisely twice the time required to travel the same distance at that arbitrary but constant speed. The same obviously applies to deceleration followed by acceleration.

By numerical demonstration, this is clearly always true, (that indeed is how I discovered it, by pure chance – serendipity again – I had never knowingly encountered it previously). In fact the underlying reality applies to both acceleration and deceleration portions individually. The formal proof is straightforward:

> An initially stationary object accelerates with uniform acceleration rate a up to speed v in time t, such that $v = at$
> in the same time, it travels a distance $s = at^2/2$.
> In the same amount of time, travelling at constant speed v, it would cover the distance
> $s' = vt = at^2, = 2s$.
> In other words, in the time it takes to accelerate to v, it could travel twice as far at constant speed v. So it takes twice the time it would take to cover the same distance at constant speed. Q.E.D.
> (I am obliged to my good friend, Dr. David Sutherland, for the above neat exposition.)

What this means is that by the time the stopping train reaches the end of the (virtual) station loop, having decelerated to zero and re-accelerated back up to line speed, but without any wait time at the station, the capacity slot which it gave up on entering the (physical) station loop has travelled in the slot stream on the main line, at constant line speed, twice the distance of the (virtual) station loop length. In other words, it is now a distance equal to the (virtual) station loop length ahead of its former occupying train. I call this distance (or time, since the speed is constant,) expressed in slots, the **Slot Stream Advance**. This will not, automatically, be an integer multiple of the slot (length or time). For the train to be able to re-join the slot stream, and thus the main line, the slot stream advance must be **made to be** an integer multiple of slots. This ensures that the train, on reaching the end of the (virtual) station loop, coincides precisely in location and speed with the next (empty) slot, previously occupied and given up by the **next** (stopping) train. To achieve this, the train is held at the station for an equalisation or basic wait time equal to that fraction of a time slot which must be added to the slot stream advance, to make the (Corrected) Slot Stream Advance an integer. Incidentally, that determines the repeat frequency: a stopping train occurs every n slots where n = the slot stream advance.

Note that no numerical values at all have been mentioned. The above argument applies to **any** value of line capacity whatever. (It still applies if the capacity isn't even an integer.)

I was surprised, I admit it.

The value just described, the basic wait time, being a fraction of the slot time, is too small to be a usable station wait time; it is too low for any but a very small number of passengers to leave and join the train. An arbitrary same number of slot times may be added to the wait time and the slot stream advance, while still maintaining the ability of the train to re-join the slot stream. The slot stream divides logically into a number of virtual sub-streams, the same number as the slot stream advance. At least one sub-stream, but possibly more, will be a stopping sub-stream, and the rest will all be non-stop. An individual stopping sub-stream is associated with a particular set of stations, and a particular platform face at each station (though this last requirement is readily varied operationally). The actual traffic pattern on the main line consists of a train from each sub-stream in turn, the pattern repeating indefinitely.

The further condition which, together with the condition that the line capacity must be an integer number of trains per hour, ensures a viable, usable timetable, is two-fold, and is simply stated. The slot stream advance, and thus the number of capacity sub-streams, must be an integer sub-multiple of the line capacity, so that it repeats an integer number of times every hour. Moreover, the **other** integer submultiple, (the number of repeats per hour,) must divide up the hour, 60 minutes, in a useful and convenient manner. It is hard to formulate this **precisely** (like **defining** an elephant, even though we all recognise it without difficulty when we see one); what is means in effect is that you can perceive it immediately, without having to perform a calculation. Furthermore it produces a clock-face timetable: the trains stop at the same times every hour, and the time interval between adjacent trains is always the same (but not necessarily an integer number of minutes, or even of seconds). Thus, supposing the line capacity is 32tph, if there are 4 sub-streams, consisting of 8tph each (the other integer sub-multiple), then the stop times at a station repeat every 7½ minutes, likewise a line capacity of 30tph with 5 sub-streams each of 6tph gives station stop times of every 10 minutes.

There are only a very few possible values of line speed which satisfy the above two (three?) necessary and, together, sufficient conditions. The table below contains the results, for the range of line speeds of interest, (with my selection of best choices in red).

Line Cap-acity (tph)	Slot time (sec)	Line Speed (m/s)	Line Speed (kph)	Line Speed (mph)	Minimum Inter Station Distance (km / miles)	Slot Stream Advance (integer Slots)	Station Wait Time (sec)	Clock-Face Timetable (every ↓ min)
60	60	38.37	138.12	85.83	3.92 / 2.44	4 / 5 / 6	138 / 198 / 258	4 / 5 / 6
50	72	57.58	207.29	128.81	8.84 / 5.49	5	206	6
48	75	61.15	220.13	136.79	9.97 / 6.19	4 / 6 / 8	212 / 287 / 362	5 / 7½ / 10
45	80	69.01	248.44	154.38	11.68 / 7.26	5	216	6⅔
40	90	74.71	268.94	167.12	14.88 / 9.24	5 / 8	251 / 521	7½ / 12
36	100	82.19	295.90	183.87	18.02 / 11.19	4 / 6	181 / 381	6⅔ / 10
32	112.5	90.80	326.87	203.11	21.98 / 13.65	4 / 8	208 / 658	7½ / 15
30	120	95.70	344.51	214.08	24.42 / 15.17	5 / 6	345 / 465	10 / 12
25	144	110.57	398.04	247.34	32.60 / 20.24	5	425	12
24	150	114.14	410.92	255.34	34.74 / 21.57	6 / 8	596 / 896	15 / 20

These are the only ones worth worrying about. In fact, if there is more than one option on offer, then, generally, only one is worth considering: the others offer too little – or even too much – wait time at the station. There are a few more, but at line speeds beyond even the dreams of HS2 Ltd. (The next one is 20tph, with a line speed of 294mph!)

Note that the values in the above table **are the only speeds at which it is sensible or even possible to operate a High (or Medium) Speed Railway**. I call them the **Sweet-Speeds**. In other words, **there is no unfettered free choice of line speed for a HS railway**, only those few Sweet-Speeds are available. Overtaking is possible for **any** line speed whatever, but very, very few speeds give a viable timetable.

This is an astounding outcome. Encountering it for the first time is an epiphany, a 'Eureka' moment, experiencing the wonder of 'On First Looking into Chapman's Homer'. Like all such experiences, once seen and understood, as the above explanation seeks to enable, it is obviously true; it could not be otherwise.

Of the nine available Sweet-Speeds in the High Speed range, the first two (capacities 50 and 48tph are a bit too slow. Their line speeds are in the narrow range between the turnout limit and buffer-end speeds. They are genuine High Speed cases, in that they perform the first part of their deceleration while still in the path of the following train, But in both cases, they start their deceleration only when at least the front of the train has moved on to the physical station loop. (in the 50tph case, it starts decelerating only when the switch is already in the process of resetting back to the main line!). Irrespective of formal category, these two would in practice be used for Medium Speed applications. (the only genuine Medium-Sweet-Speed case is for capacity 60tph.)

The two highest speeds (capacities 25 and 24tph) give station wait times of 7 minutes and above, which are far too long. (They are also, in my opinion, rather too fast.). The remaining 5 all give very acceptable options, as I have highlighted. Of these, the one I personally like best is 32tph for the (to me) decisive reason that it offers the perfect clock-face timetable of a train every 7½ minutes, or the equally good one, (by leaving alternate slots empty – see the section on Resilience,) of every quarter of an hour. It also offers a very decent line speed of 203mph, not too fast and not too slow, and the best station wait time, of 3min28sec. My second choice would be 40tph, for its equally good timetable, though it may be considered, (by HS2 Ltd.,) as, at (only!) 167mph, a bit slow.

The technique alluded to above, of running a timetable with half the slots empty, is in fact a very useful, indeed essential, operational method. It is essential if the Same Speed Line divides into two routes, and the services divide equally, trains taking alternate routes in turn. The branching routes continue to be scheduled exactly as before; half the trains are phantoms, but their dynamic behaviour is identical to the real ones. Travelling in the opposite direction, all the (real) trains are perfectly aligned to merge the two branches' traffic.

Exactly the same considerations apply if the services divide but the line itself doesn't, e.g. if half the trains terminate at an intermediate station, while the rest continue further.

In fact, it makes very good sense to schedule an entire Same Speed route, (or even several, interconnected, cooperating routes, scheduled as a group,) to the same capacity-based timetable throughout. This serves the fundamental purpose of **imposing a uniform, capacity-slot-based time standard throughout**. Over much of the route, (at least) half of the (stopping) trains will be phantoms, but so what? The remaining real trains will be delivering exactly the service required at a particular location.

From the numerical results, it is seen that, for High Speed lines over the entire speed range of interest (and beyond), 108 – 264mph (55 – 23tph), the minimum slot stream advance is 3 slots. Outside that range, at both ends, it is 2. For Medium Speed lines, over their entire rather short speed range, between 64 and 127mph, (62 – 51+tph), the minimum slot stream advance is 2 slots.

Before considering an actual numerical example, something else is worth illustrating in general. More than one train will be on the (virtual) station loop at one time. Is there any possibility they could get in each other's way?

For uniform rates of acceleration and deceleration, from zero to and to zero from a given speed, both the times taken and the distances travelled are inversely proportional to those rates. This is very easily demonstrated:

$$v = a_a t_a = a_d t_d \quad \text{so } t_a/t_d = a_d/a_a$$

For the same times $\quad s_a = a_a t_a^2/2 \qquad s_d = a_d t_d^2/2$

so $\quad s_a/s_d = a_a t_a^2/\, a_d t_d^2 = (a_a/a_d)*(t_a/t_d)^2 \quad = (a_a/a_d)*(a_d/a_a)^2 \quad = a_d/a_a \qquad$ Q.E.D.

Therefore, provided **only** that the deceleration rate is greater than the acceleration rate, (as, in practice, it always is,) the acceleration time and distance are greater than the deceleration time and distance. In the present context, the train reaches the station in a shorter time and distance than it then requires to re-accelerate back up to line speed. The empty slot given up by the train reaches (via the main line) the far end of the (virtual) station loop at the same time as the train has stopped at the station and waited there for a certain time. The empty slot must travel a certain distance further until it is an integer number of slots (time or distance) from the start of the (virtual) station loop, at which point the slot containing the next stopping train reaches the start of station loop. The train has waited at the station for the same amount of time, but since it did part of its waiting before its empty slot reached the end of station loop, by the time that the slot containing the next stopping train reaches the start of station loop, the current stopping train has already departed from the station. In other words, the train leaves the station before the next stopping train reaches the start of station loop. (In the impractical situation of deceleration and acceleration rates being equal, the train would depart the station precisely as the next stopping train entered the station loop. No time values have been specified, so this is always true. The argument is a bit cumbersome, but the result is clear and definitive.

This stuff is not easy to visualise, so here is an actual numerical example, using my favourite line capacity of 32tph.

Line capacity = 32tph. Slot time = 112.5sec. Slot length = 10.2146km.
Line speed = 90.80m/s =326.87kph = 203.11mph.
Deceleration time = 181.59sec = 1.6142 slots Deceleration distance = 8.2441km = 0.8071 slots
Acceleration time = 302.66sec = 2.6902 slots Acceleration distance = 13.7401km = 1.3451 slots
(Virtual) Station Loop Travelling time = 484.25sec = 4.3044 slots
Station Loop distance = 21.9842km = 2.1522 slots

By the time the stopping train has travelled the length of the station loop, its empty slot, which it gave up on entering the loop, has travelled 4.3044 slots (time or distance) on the main line. It is thus 2.1522 slots (time or distance) beyond the end of the (virtual) station loop. In order to make this 3 slots exactly, it must travel a further 0.8478 slots (time or distance). This implies that the train must wait for 0.8478 **time** slots = 95.38sec at the station.

We now follow the progress of slot stream and stopping train.

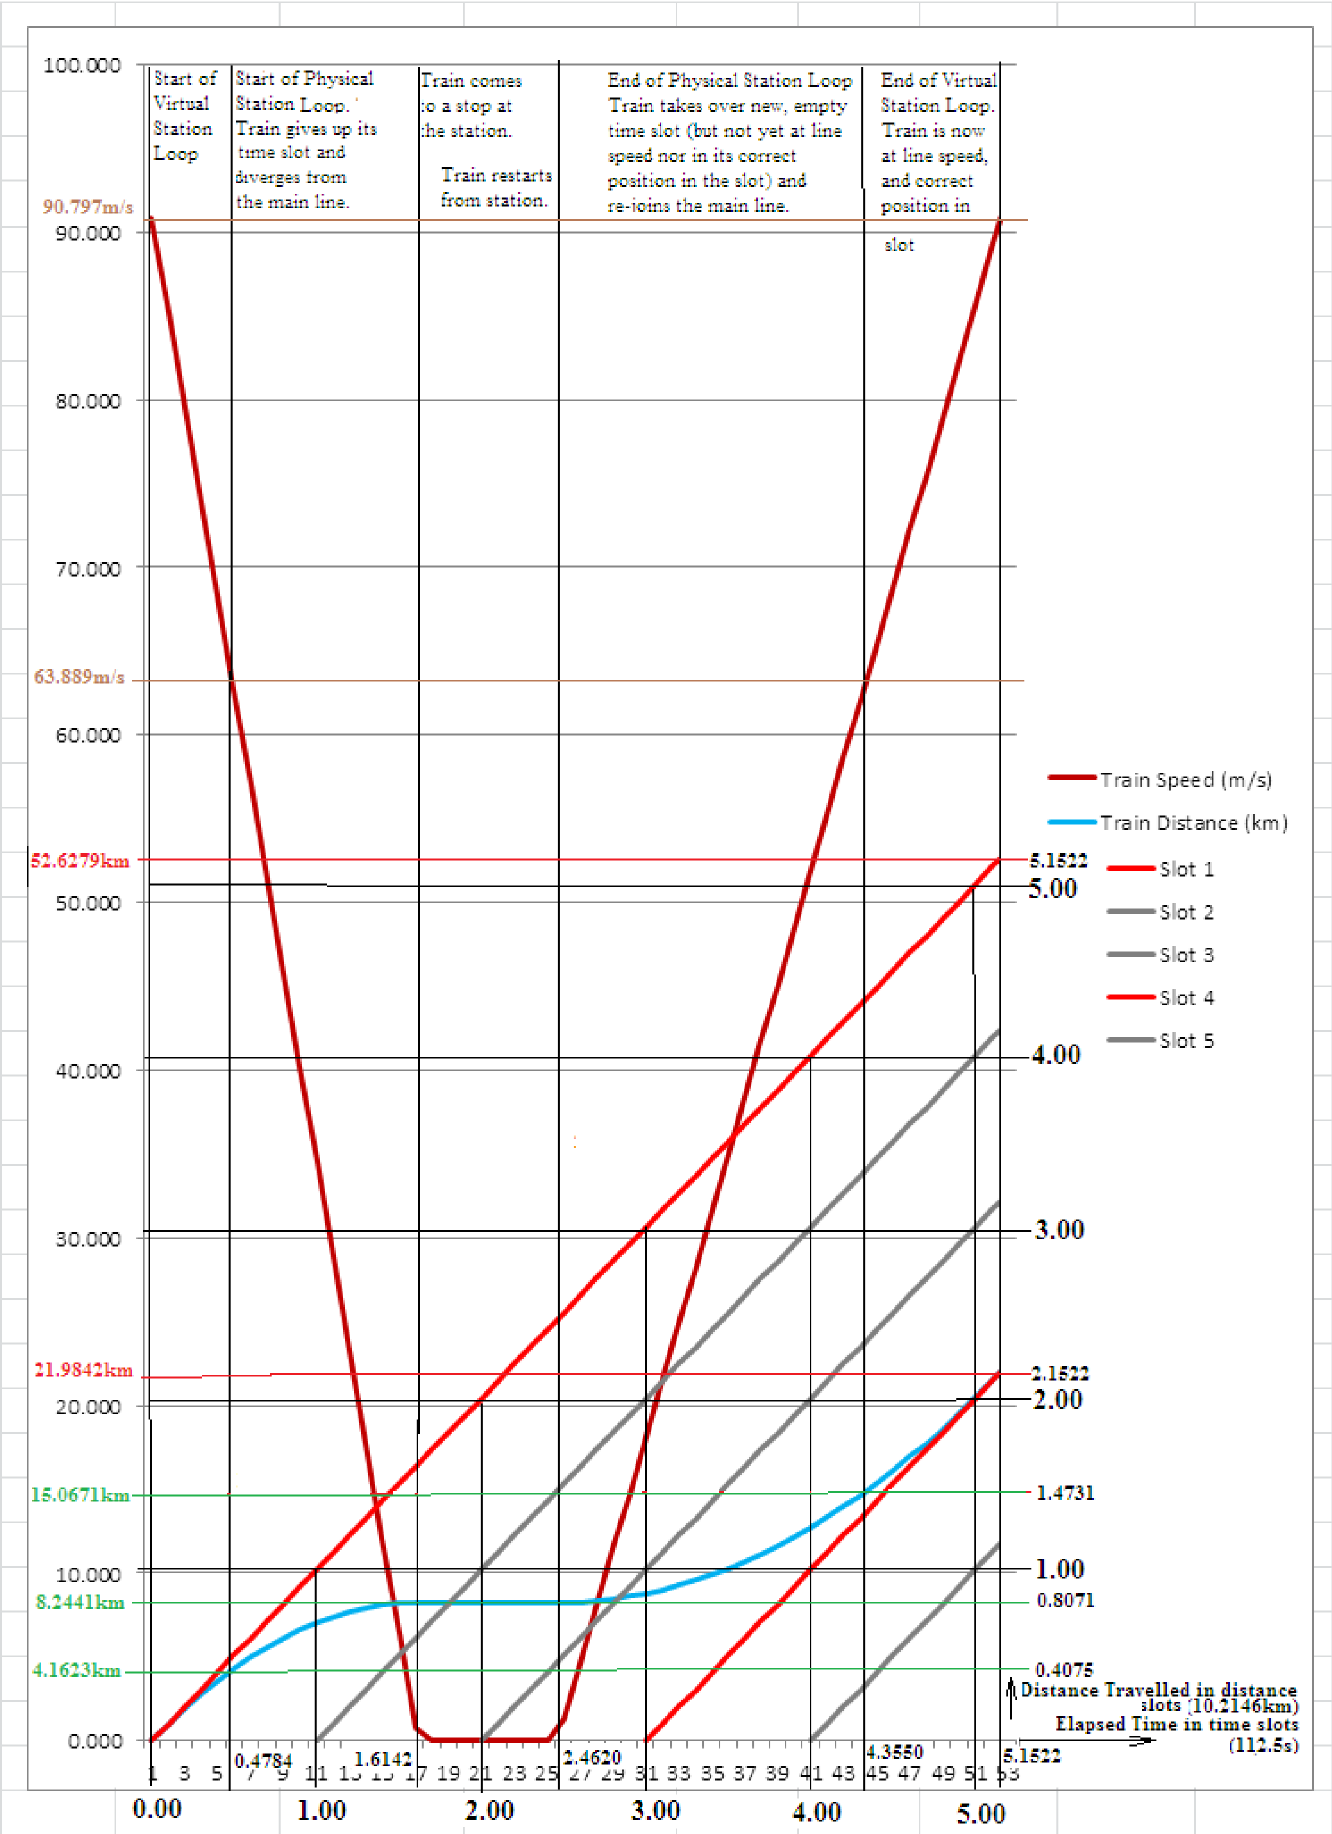
Start of Virtual Station Loop
Start of Physical Station Loop. Train gives up its time slot and diverges from the main line.
Train comes to a stop at the station.
Train restarts from station.
End of Physical Station Loop Train takes over new, empty time slot (but not yet at line speed nor in its correct position in the slot) and re-joins the main line.
End of Virtual Station Loop. Train is now at line speed, and correct position in slot
100.000
90.797m/s
90.000
80.000
70.000
63.889m/s
60.000
52.6279km
50.000
40.000
30.000
21.9842km
20.000
15.0671km
10.000
8.2441km
4.1623km
0.000
Train Speed (m/s)
Train Distance (km)
Slot 1
Slot 2
Slot 3
Slot 4
Slot 5
5.1522
5.00
4.00
3.00
2.1522
2.00
1.4731
1.00
0.8071
0.4075
Distance Travelled in distance slots (10.2146km)
Elapsed Time in time slots (112.5s)
0.4784
1.6142
2.4620
4.3550
5.1522
0.00
1.00
2.00
3.00
4.00
5.00

(The following description was written at the same time as the rest of the section, published in October 2020, v3.0 of this article. It now has an accompanying Station-Calling Diagram, above, produced much later (18 months or so), based on techniques I developed to illustrate the process by which the line speed of Same Speed railways can be varied; an astonishingly complicated and very difficult subject, published in Volume 2 of this article – the **really** hard stuff!)

The above diagram illustrates the speed and the distance travelled by the train, as it calls at a station, distance measured from the start of the virtual station loop. The various inclined parallel straight lines are the paths of the capacity slots, travelling at constant line speed, each one either containing a non-stop train (the grey lines) or (the red lines) being empty, either (slot 1) given up by the present train, when it diverges onto the physical station loop or (slot 4) given up by the **next** stopping train (which does **not** appear in the diagram, though it certainly appears in the accompanying description), when it in turn diverges onto the physical station loop. Slot 4 is then taken over and occupied by the present train, when it re-joins the slot stream on the main line at the end of the physical station loop, completing its acceleration back up to line speed and moving to its correct position in the slot by the time it reaches the end of the virtual station loop – at the very rightmost side of the diagram, where the green line of the train's trajectory becomes enveloped in the red of the slot (a fortuitous but perfect outcome of the order in which the lines are drawn, since the train is then actually inside the slot).

The numerical values almost all appear in the description, but the following aspects pertain to the diagram only, so are elucidated in advance.

The row of small numbers along the x-axis are **not** numerical values as such, but **labels** for the **sets** of values appearing vertically above. They are numbers starting at 1. (This is simply how Excel presents it – no intrinsic significance.) The numbers I have added manually are the elapsed times, and these are given for every time slot, (112.5s for capacity 32tph,) starting at zero. Far from being a strange unit, the time slot is absolutely the natural unit to use for this type of application, dealing with timetables.

The ordinate (y-axis) is either speed (m/s) or distance travelled (km). The unit is clearly stated together with the value. Most unusually, the same value range is convenient for both. (I usually have to use strange special units – tenth of a kilometre for example – to get different units to display legibly on the same graph.) Note that, at the right hand side of the graph, the results for distance travelled are also given in **distance slots** (10.2146km for line speed 90.797m/s), both the standard values 1, 2, 3, 4 and 5, but also the explicit values for any km values specified explicitly.

The virtual station loop is the whole width of the graph, between time slots zero and 5.1522. The physical station loop is between the next-inward pair of full-height vertical lines, at time slots 0.4784 and 4.3550, where the (front of the) train begins to cross the switch, diverging/converging, at the turnout limit speed (63.889m/s). The wait time st the station is between the innermost pair of vertical lines, at slot times 1.6142 and 2.4620, thus 0.8478 time slots = 95.35s = basic wait time. These aspects don't really feature in the following description as they don't affect the calculation, but they're certainly of interest.

Finally, the distance between the train and its original containing slot (1), when the train reaches the end of the virtual station loop is 52.6279km (slot 1) – 21.9841km (train inside slot 4), = 30.6438km. This, the (corrected) Slot Stream Advance, is **exactly** 3 × 10.2146km, the slot length, as it (an integer) should be. (Note that this is **not** an integer sub-multiple of the line capacity, so does not give a useful timetable. The minimum value for that is 4. Immediately following the present exposition, the whole process is repeated (starting at p.24,) with the Slot Stream Advance set to 4 ab initio. An updated diagram is included. This may seem too much of a good thing, but the subject is of such fundamental importance that I consider it worth repeating, in a form which is actually used for real.

So, finally, the description:

Time Slot 1:
The empty slot advances 1 slot along the main line.
The train decelerates for 1 time slot. It thus has 0.6142 time slots of deceleration still to do.

Time slot 2:
The empty slot advances a further 1 slot along the main line. It is now 2 slots beyond start of (virtual) loop.
The train completes its deceleration in 0.6142 time slots, and reaches the station. It waits there for 0.3858 time slots. It thus has 0.4620 time slots still to wait.

Time Slot 3:
The empty slot advances a further 1 slot along the main line. It is now 3 slots beyond start of loop. It has also passed the end of the (virtual) station loop; it is 0.8478 slots beyond end of loop. (Refer to the distances along the RH axis. The end of the virtual station loop is at 2.1522 distance slots; the empty slot, at 3 distance slots exactly, is 0.8478 distance slots beyond it.) The slot containing the next stopping train slot 4) has arrived at start of loop.
The train waits at the station for a further 0.4620 time slots. It then departs, performing the first 0.5380 time slots of its acceleration. It thus has 2.1522 time slots of acceleration still to do.

Time Slot 4:
The empty slot advances a further 1 slot along the main line. It is now 4 slots beyond start of loop, and 1.8478 slots beyond end of loop.
The second empty slot advances 1 slot along the main line.
The train performs 1 time slot of acceleration. It thus has 1.1522 slots of acceleration still to do.
The second stopping train decelerates for 1 time slot. It thus has 0.6142 time slots of deceleration still to do.

Time Slot 5:
The empty slot advances a further 1 slot along the main line. It is now 5 slots beyond start of loop, and 2.8478slots beyond end of loop.
The second empty slot advances a further 1 slot along the main line. It is now 2 slots beyond start of loop.
The train performs 1 time slot of acceleration. It thus has 0.1522 time slots of acceleration still to do.
The second train completes its deceleration in 0.6142 time slots, and reaches the station. It waits there for 0.3858 time slots. It thus has 0.4620 time slots still to wait.

During Time Slot 6:
The empty slot advances a further 0.1522 slots along the main line, to 5.1522 slots from start of loop and the (corrected) slot stream advance is **3 slots exactly beyond end of loop.**
The second empty slot advances a further 0.1522 slots along the main line. It is now 2.1522 slots beyond start of loop. **It therefore coincides precisely with end of loop.**
The train performs its final 0.1522 time slots of acceleration. **It therefore coincides precisely with end of loop and is travelling at line speed.**
(The second stopping train also accelerates for 0.1522 time slots – but so what?)

In practice, this line speed requires a slot stream advance value of 4, thus also 4 slot sub-streams. An extra time slot must therefore be added to the station wait time also, giving 1.7813 time slots = 207.88sec.

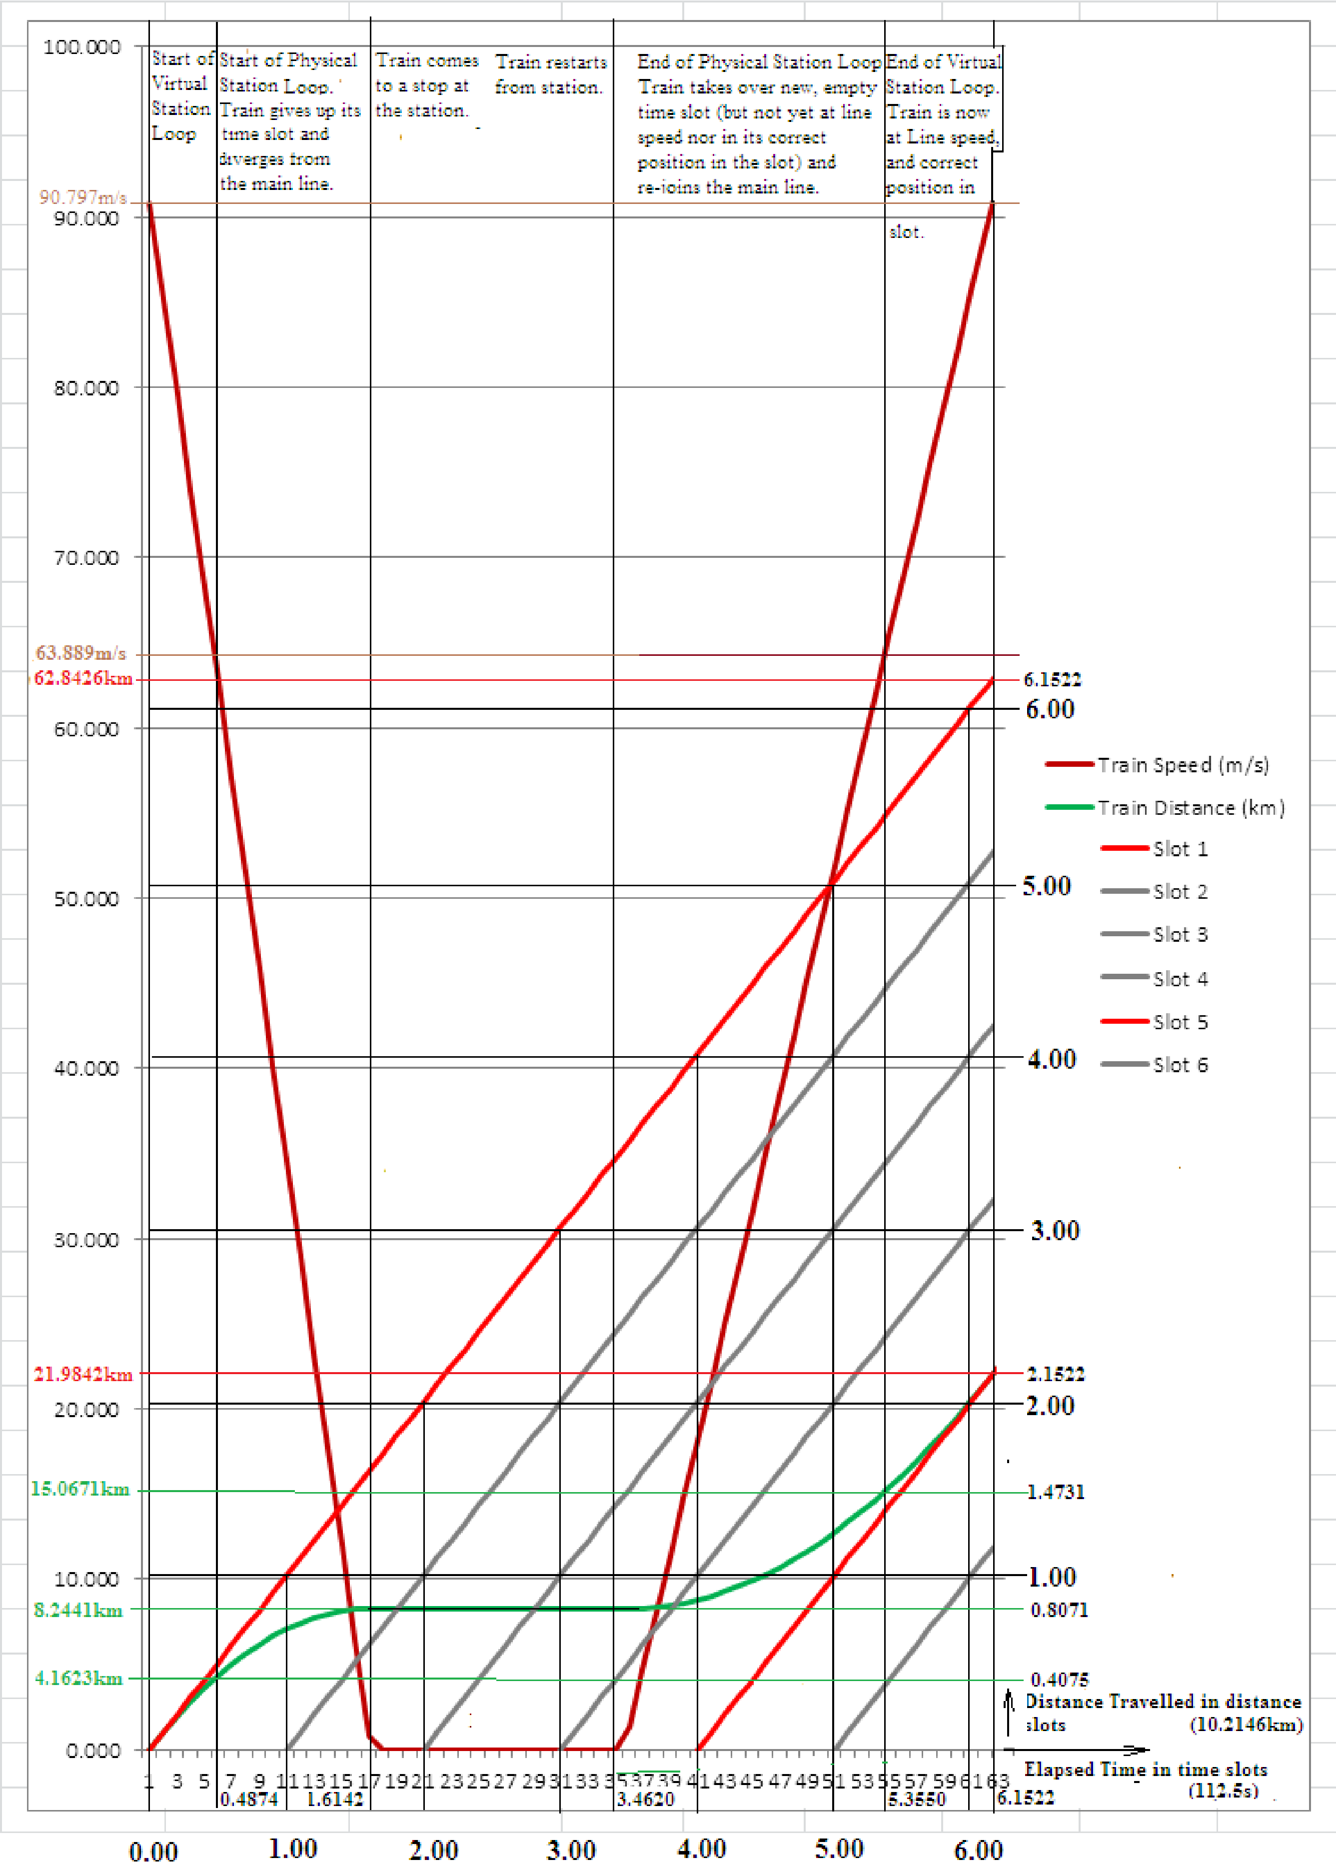

A new Station-Calling Diagram (above) has been added, reflecting this configuration:

Taking the capacity value of 32tph as usual:

Line capacity = 32tph. Slot time = 112.5sec. Slot length = 10.2146km.
Line speed = 90.80m/s =326.87kph = 203.11mph.
Deceleration time = 181.59sec = 1.6142 slots Deceleration distance = 8.2441km = 0.8071 slots
Acceleration time = 302.66sec = 2.6902 slots Acceleration distance = 13.7401km = 1.3451 slots
Station Calling Section travelling time = 484.25sec = 4.3044 slots
Station Calling Section distance = 21.9842km = 2.1522 slots

By the time the stopping train has travelled the length of the station loop, its empty slot, which it gave up on entering the loop, has travelled 4.3044 slots (time or distance) on the main line. It is thus 2.1522 slots (time or distance) beyond the end of the (virtual) station loop. In order to make this 3 slots exactly, it must travel a further 0.8478 slots (time or distance). This implies that the train must wait for 0.8478 **time** slots = 95.38sec at the station. But we know in advance that this would give a slot stream advance of 3, when the value required is actually 4. So take that value ab initio, thus station wait time = 1.8478 time slots, = 207.88sec.

We now follow the progress of slot stream and stopping train.

Time Slot 1:
The empty slot advances 1 slot along the main line.
The train decelerates for 1 time slot. It thus has 0.6142 time slots of deceleration still to do.

Time slot 2:
The empty slot advances a further 1 slot along the main line. It is now 2 slots beyond start of (virtual) loop..
The train completes its deceleration in 0.6142 time slots, and reaches the station. It waits there for 0.3858 time slots. It thus has 1.4620 time slots still to wait.

Time Slot 3:
The empty slot advances a further 1 slot along the main line. It is now 3 slots beyond start of loop. It has also passed the end of the (virtual) station loop; it is 0.8478 slots beyond end of loop.
The train waits at the station for 1 time slot. It thus has 0.4620 time slots still to wait

Time Slot 4:
The empty slot advances a further 1 slot along the main line. It is now 4 slots beyond start of loop, and 1.8478 slots beyond end of loop.
The slot containing the next stopping train has arrived at start of loop.
The train waits at the station for a further 0.4620 time slots. It then departs, performing the first 0.5380 time slots of its acceleration. It thus has 2.1522 time slots of acceleration still to do.

Time Slot 5:
The empty slot advances a further 1 slot along the main line. It is now 5 slots beyond start of loop, and 2.8478slots beyond end of loop.
The second empty slot advances 1 slot.
The train performs 1 time slot of acceleration. It thus has 1.1522 slots of acceleration still to do.
The second train decelerates for 1 time slot. It thus has 0.6142 time slots of deceleration still to do.

Time Slot 6:
The empty slot advances a further 1 slot along the main line. It is now 6 slots beyond start of loop, and 3.8478slots beyond end of loop.
The second empty slot advances a further 1 slot. It is now 2 slots beyond start of loop.
The train performs 1 time slot of acceleration. It thus has 0.1522 time slots of acceleration still to do.
The second train completes its deceleration in 0.6142 time slots, and reaches the station. It waits there for 0.3858 time slots. It thus has 1.4620 time slots still to wait.

During Time Slot 7:
The empty slot advances a further 0.1522 slots along the main line, to 6.1522 slots from start of section and the (corrected) slot stream advance is **4 slots exactly beyond end of loop.**
The second empty slot advances a further 0.1522 slots along the main line. It is now 2.1522 slots beyond start of loop. **It therefore coincides precisely with end of loop.**
The train performs its final 0.1522 time slots of acceleration. **It therefore coincides precisely with end of loop and is travelling at line speed.**
(The second train also waits at the station for 0.1522 time slots – but so what?)

There follows yet another Station-Calling Diagram, illustrating the interaction between three adjacent stopping trains.

It takes quite a bit of brain-pounding to get completely at ease with the idea of a stopping train giving up its capacity slot to the **previous** train, (the one in front of it, in other words,) and later taking over the slot given up by the **next** train, (the one following it, in other words). This is a quirky consequence of the stopping train being overtaken, not so much by individual trains, but by the (constantly moving at line speed) slot stream, together with any non-stopping trains carried along with it.

Nonetheless, a picture being worth a thousand words, (I've often wondered who performed that calculation,) here is a picture, to illustrate the overtaking process. This is my best shot at a graphical aid; I hope you like it.

Three (stopping) trains participate in the diagram, the **previous** train, (orange line,) on the left, showing just the end of its trajectory, the **present** train, (dark green line,) showing its complete trajectory, and the **next** train, (light green line,) on the right, showing just the beginning of its trajectory.

The present train is decelerating. At it reaches the switch onto the physical station loop, and also the turnout limit speed, it diverges from the main line onto the physical station loop, giving up its capacity slot on the main line and diverging onto the physical station loop. This occurs at elapsed time 0.4784time slots, = 52.820s. 98.618s later, at 1.3550time slots = 152.438s the previous train takes over this slot, as it reaches the switch at the end of the physical station loop, its accelerating speed reaches the turnout limit speed, and it joins the main line and the slot stream. The trajectory of the empty slot between these times is the straight red line of slot 1. Note that both trains are moving, relative to the slot. At time 2.1522 time slots = 242.123s, the previous train has reached line speed, and its correct position in the slot.

As the present train waits at the station two non-stop trains pass it on the main line, as does a third, when it hss just departed from the station. At elapsed time 4 time slots = 450s, the next stopping train reaches the start of the virtual = station loop and begins its deceleration. In a further 0.4784 time slots = 52.820s it has reached the switch, at the turnout limit speed, and gives up its slot on the main line, diverging onto the physical station loop. The present train reaches the switch at the end of the physical station loop, and the turnout limit speed, at the same time as the empty slot, takins it over and re-joining the slot stream on the main line, at 6.3550 time slots = 714.938s. It reaches line speed, and its correct position within the slot, at 6.1522 = 692.123s.

And that's it.

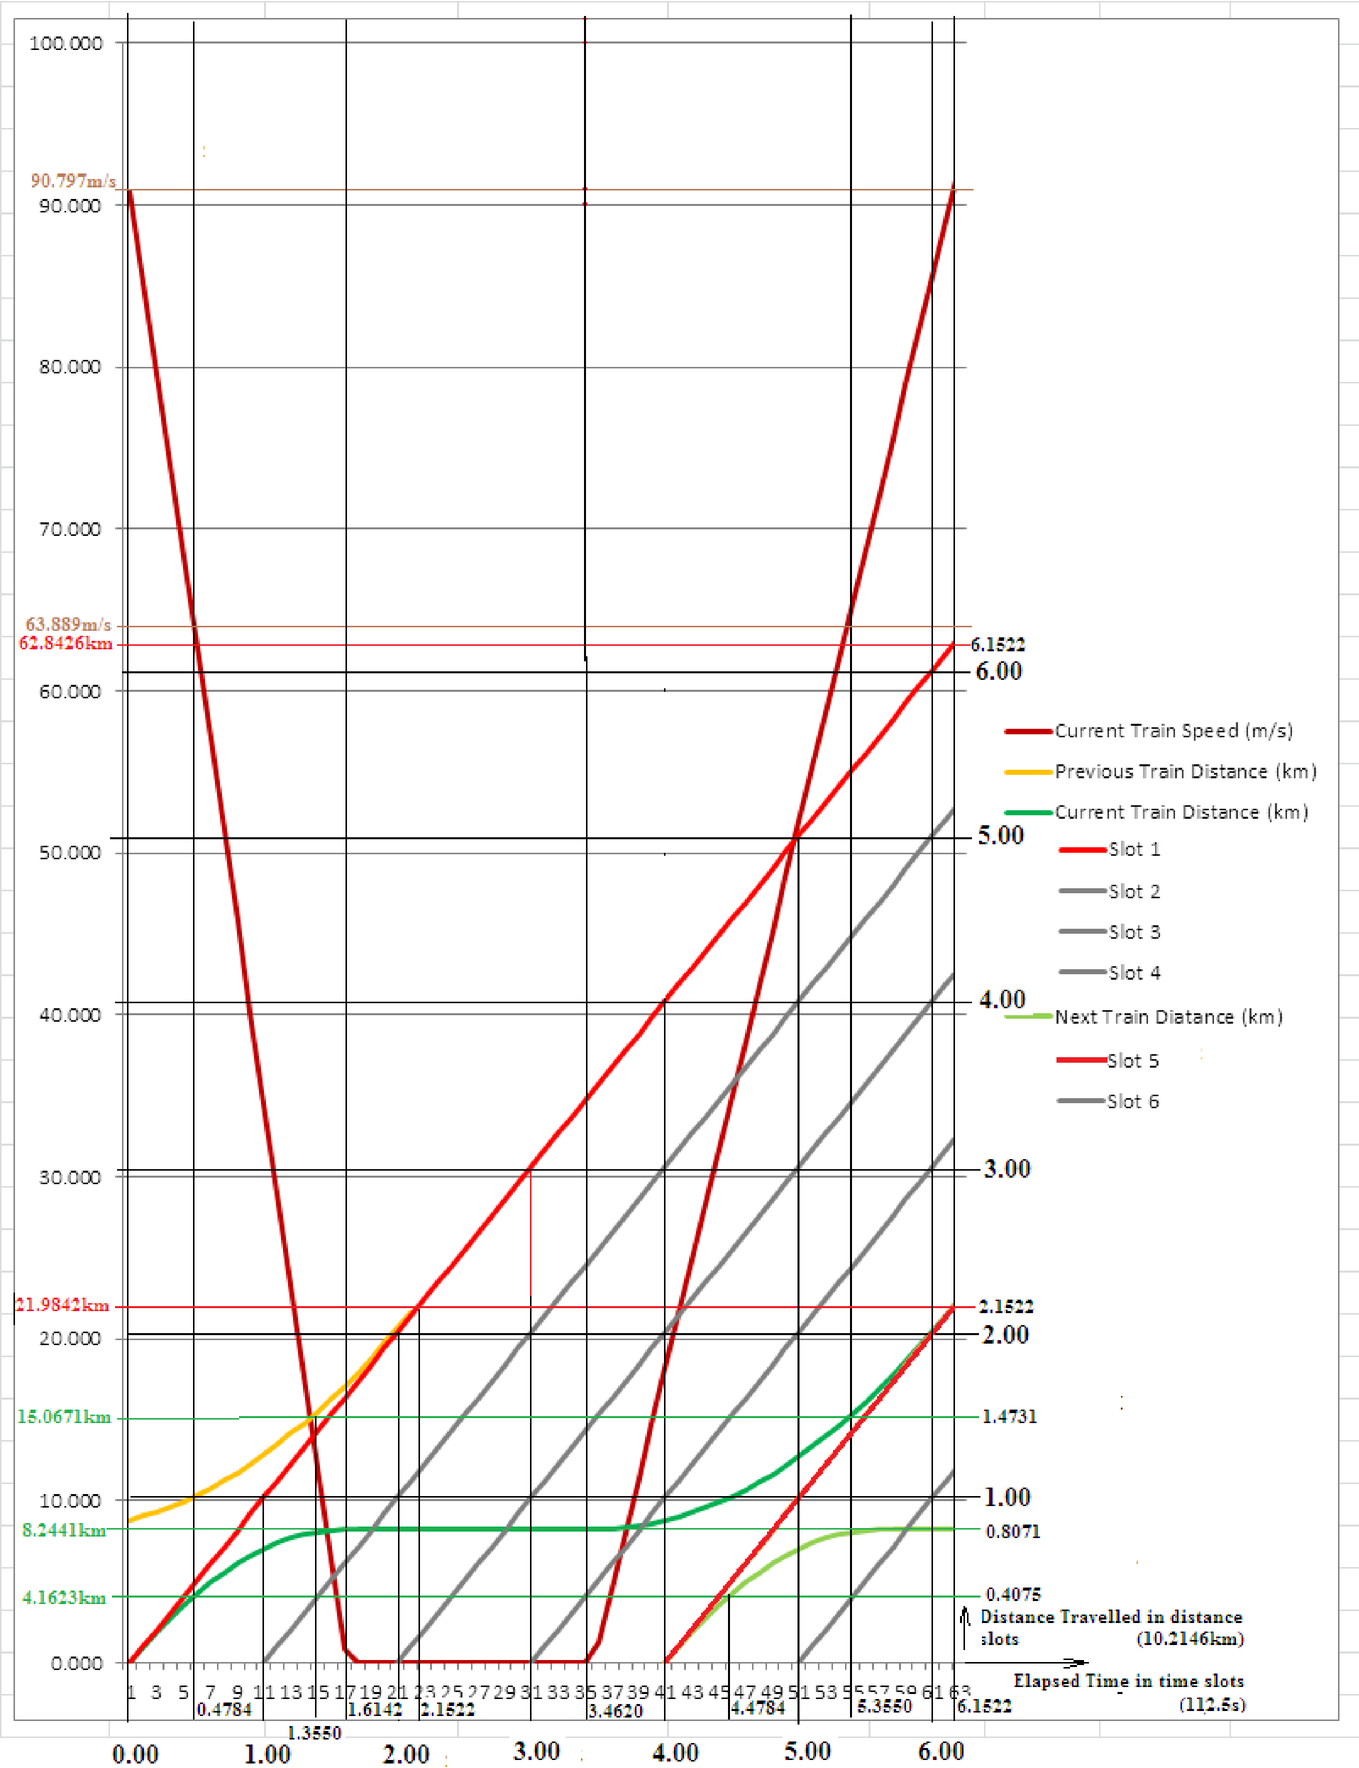

Readers may legitimately wonder why I am so careful to specify **time** slots on every occasion for the train, but only 'slots' for the slot stream. This is because the slot stream moves at constant speed – the line speed – throughout, so time and distance slots are in permanent 1:1 correspondence. The **only** constant speed experienced by the train is zero, as it waits at the station. But its motion **(or lack of it!)** is still accurately measured in time slots.

(I wasted an awful lot of time in this investigation, in believing that the way to ensure that stopping trains were able to re-join the slot stream and thus the main line was to make the station stop time, being the sum of deceleration, waiting and acceleration times, an integer multiple of the slot time. The trouble was, this gave very plausible results, not so very different from those given earlier in the summary table. It was only by performing a detailed traverse of the station loop with actual numerical values, as in the example above, (which was in fact the one I did use,) that I discovered that the end points just didn't coincide. Having made that discovery, a little **clearer** thinking identified the correct value to use – there weren't many candidates.)

It may be wondered: what is the point of the Medium Speed category? The basic idea is that it may be suitable for conversion of existing classic routes, or portions thereof, to Same Speed standards. Imagine running the West Coast Main Line at 48, 50 or even 60tph! 48tph would, in my opinion, be the best choice, since it offers what is, in classic terms, a high line speed, (137mph,) and the perfect timetable with 6 slot sub-streams each of 8tph – a clock-face of every 7½ minutes.

A point I haven't stressed at all, but which is certainly relevant here, is that Same Speed Railways only need two tracks. So converting a classic route to Same Speed requires **absolutely no widening**. All the work required is concentrated at stations, with station loops and at least two platform faces in each direction – a significant amount of work, for sure, but trivial compared with quadrupling. Another possible application of Medium Speed is as the semi-metro outer reaches of a metro line. More on this in the next section but one.

This section has, in a sense, addressed a superficial problem, that of overtaking, but, in order to solve it, had to address the fundamental, underlying problem, which is:
how to remove individual trains from the slot stream, on the main line, and then put them back again, into the slot stream, on the main line, later. (This is, invariably, in order to stop at a station.)

If the main line bypasses the station platforms, then non-stop trains are able to overtake. But this does not have to be the case. It could be that the main line simply divides into a number of tracks, each one of which serves a station platform, so **all** the trains stop there. This is the definition of a pure metro. The slot stream simply continues on its way, and the various platform branches merge to re-form the main line, which re-joins the slot stream, (**very** much that way round!).

The treatment given in the present section exactly follows the way I originally discovered it. The notion of a station loop, on which a stopping train physically calls at a station, while all the overtaking trains, still in their capacity slots in the slot stream on the main line (i.e. the non-platform track) overtake it, reflects what actually, physically happens, and forms a clear picture in everyone's mind. But it just won't do for the next but one section, which deals with metros.

So I now introduce the concept of the **station calling section**. This is the section of track on and over which a train calls at a station. Its composition is formally identical to the (virtual) station loop, dealt with previously, for a train which is to be overtaken. There will in all cases be a unique portion of track associated with a specific station platform, but the rest of the section, covering the deceleration and acceleration tracks, **may** be unique, (for a physical station loop,) or **may not be** unique, (for the divisions and subdivisions of the main line into platform tracks, at a station where all trains stop, and there are no through, non-platform tracks). In the former case, the slot stream has an imaginable reality, but in the latter, it doesn't. It just passes through the station area without any perceptible presence or location, but when the platform tracks, on leaving the station merge and coalesce to re-form the main line, it is re-associated with it. It has thus become what it truly is, a purely virtual time standard, to which the behaviours of all trains participating in the service refer. The only point that the reader really needs to take away from this is that the calculations are, in all essentials, identical for both cases.

Timetable Diagrams.

As something of an aside, but directly following on from the Timetabling Considerations of the previous section, there now follows a graphical illustration of the overtaking process. I call this a Timetable Diagram. A graphical technique similar in appearance is already in widespread use in the timetabling of conventional railways, and doubtless has a name, probably something else entirely. But my technique has certain individual characteristics, unique to Same Speed railways.

Firstly and fundamentally, it is based on the Capacity Slot Model. For elapsed time as the independent variable, it illustrates the distance travelled by all the trains in all the capacity sub-streams departing from the origin location in a period of one clock-face hour, the origin time and distance being both taken as zero. It pertains only to one specific Same Speed route, so there will in general be graph lines which come to a sudden end, when the train in question diverges onto another route. Other lines will suddenly appear, when the train in question converges onto the route, from some other route. When a line is actually absent, for some time, (as opposed to the very short break at stations, see below,) this should be considered as a phantom.

For all periods of constant speed, the graphs are straight lines, horizontal distance apart 1 time slot (of course!) and vertical distance apart 1 distance slot (of course!).

The specimen diagram following is for the by-now-familiar HS capacity of 32tph, time slot 112.5s, distance slot 10.2146km, Sweet-Speed 90.80m/s / 326.87kph / 203.11mph, with 4 sub-streams of 8tph, (so only four colours are needed,) thus a train in each sub-stream every 450s / 7½mins and station wait time of 207.88s / 3min28s. It give just the first two trains in each sub-stream; the full treatment is four times the width. A specimen station, served by the trains in sub-stream 1, is included at a distance of precisely 96.6504km from the starting point. (The start of deceleration is at the elapsed time of 10 slots – I didn't actually set out to locate the station at that particular distance!).

(Note that I have omitted the chart line for elapsed time, the independent variable. It's not essential that this always be present, particularly in an example like this one, where I can easily add the elapsed times explicitly, by hand, which looks a lot tidier.)

he overtaking process is clearly illustrated; a stopping train is overtaken by three non-stopping trains, and so falls back 4 slots in the overall slot stream, i.e. the slot stream advance is 4. Expressing this another way, the **penalty** of an intermediate station stop is 4 time slots – 450s / 7½ minutes, and 4 distance slots – 40.8586km. By this is meant that stopping at an intermediate station causes the train to fall back in the overall slot stream 7½ minutes and 40.8586km.

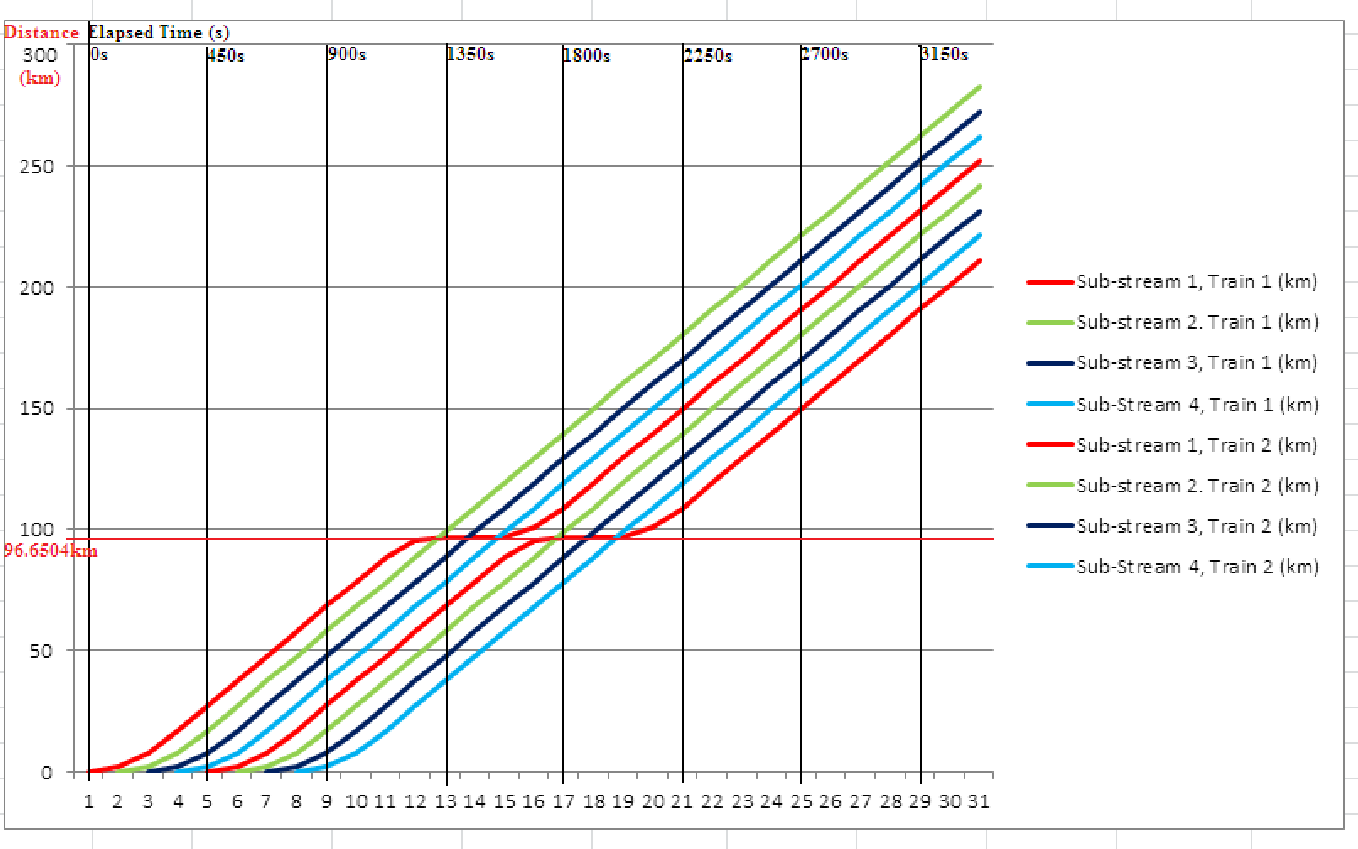

The practical utility of this technique is to be able to see at a glance the impact of a particular stopping pattern on a given service. Note that the following are not **estimates**, but exact calculations, based on the Same Speed model. They become estimates only when quoted in the context of actual services, whose adherence to the model can never in practice be exact. The aim is to get as close as practical.

Consider London – Edinburgh, by HS3. The distance is 631km. Allow 13.7401km for initial acceleration to line speed (in 302.66s) and 8.2441km for final deceleration (in 181.59s). So the distance travelled at full line speed is 609.0158km, taking 6707.49s, thus 7191.74s, and so 119min51.74s in total, including the acceleration and deceleration. Thus we have a shade under 2 hours for London – Edinburgh non-stop (cf ~75min flying time). To this must be added 7½ minutes for each intermediate stop. I doubt if the traffic on offer between London and Edinburgh would be sufficient to justify a purely non-stop service, (or even a 1 stop, at York, service, though I have made provision for such a service, splitting the train at York to serve the 8 intermediate stops between York and Edinburgh by the second portion,) but I am certainly open to persuasion, from people who know the travel markets better than I do.

Considering London – York by HS3 (323km) in the same way gives 3799.52s, 63min19.52s non-stop, a different sort of prospect entirely.

I have provided journey time estimates, (actually calculations, as explained above,) in the Route and Service Plans articles for several years now, but these have always been done the hard way, by individual calculations for the several distinct route sections The Timetable Diagram approach will allow them to be read off directly from the graphs (in practice from the spread-sheet data underlying them).

Stations on the Main Line: HS-Metros, Pure Metros and Semi-Metros

A route, or section thereof, where all trains call at all stations, **is a metro**, a **pure** metro, in fact. This is the case, irrespective of the line speed.

The traditional metro has been around for a long time, developing its characteristic properties and operational methods over a lengthy period. It is, in fact, the earliest form of a Same Speed Railway. But it was only by considering the characteristics of High Speed Railways that I originally developed the **concept** of Same Speed Railways. The realisation that metros shared many of the same properties led to the idea of the Same Speed Railway as an underlying paradigm, of which High Speed Railways and (Low Speed) metros were manifestations or categories. I do claim this as an original perception, but cannot possibly, and do not suppose that no-one else has ever thought of it. But if they have, they don't seem to have made anything of it. I first documented the idea in the initial version of the 'Same Speed Railways' article, which goes back possibly as far as late 2013, but was first published (at v2.1) in the earliest version of my website, in May 2015. (A previous version was published earlier as a magazine article, in July 2014.)

The concept of the High Speed Metro was originally developed and proposed as an alternative way of constructing and operating High Speed Railways. Unfortunately it was based on a misperception that, since no overtaking was required, and therefore no station loops (true), then no high-speed point-work was required at stations. This last is false. A train cannot simply begin decelerating at an arbitrary point on the main line, or its non-decelerating or, rather, not-yet-decelerating successor would begin to close the already-minimum separation (once that had dccreased to TSD(b), if necessary). So the main line must **divide into two** at the beginning of the station calling section, (see final paragraph of previous section,) and alternate trains take alternate tracks in approaching the station. The fact of two approach tracks determines that the number of platform faces (in each direction) at the station must be even – this is not absolutely essential but it is more convenient and operationally sanitary to have each platform face correspond to a particular approach track, (though it must be possible to override this, in non-standard operation).

As noted above, the all-trains-stop-at-all-stations, or no-overtaking model is a defining characteristic of a pure metro. A conventional, Low Speed metro may appear very different from a High or Medium Speed railway, but appearances are deceptive; the same theory applies to **all** Same Speed Railways. But it is convenient to retain the name HS-Metro, in recognition of the vast difference in perception. And there are indeed differences of detail, if not of essence.

A **semi-metro** is a metro which allows (some) overtaking. So how does that differ from a High Speed Railway with overtaking, as detailed in the previous section? Stated simply, there's less overtaking in a semi-metro. In practice, it's easy to tell the difference (easier than to define it, in fact, like defining an elephant). For the sort of intermediate stations along a High Speed Railway, a stopping service of 4tph is likely to be entirely adequate (and highly satisfactory to the populations served). The stopping service will typically be by a single, 8tph sub-stream, half of them phantoms. The other sub-streams (typically three) will all be non-stop, and may or may not contain phantoms. A semi-metro in the Low Speed range will typically be of 2 sub-streams, one stopping at every station and the other non-stop. The stopping sub-stream will itself very likely be 50% phantom.

Metro routes typically consist of a number of sections, with a central core of pure metro, where all trains stop at all stations, then Low Speed semi-metro sections beyond the core in both directions, over which some trains travel non-stop, to become stopping trains again in the outer sections, which are Medium

Speed semi-metros. The line speed actually changes at a station where all trains stop, the stopping trains of the inner, Low Speed section often terminate at this station, and the continued service is taken over by the former non-stop trains. By this means significantly shorter journey times between the outer reaches and the central core are made possible, avoiding such disagreeable present journeys as, for example, between West Ruislip and Central London, or, even more so, between Epping and Central London.

One genuine difference, and a defining one, between Low Speed metros or semi metros and those in other speed ranges is that, in the Low Speed range, line capacity **increases** with line speed. This means that we adopt the highest appropriate line capacity and speed, and, if the capacity is more than actually needed, introduce a proportion of phantoms. (If reducing the actual trains by half seems excessive, we can be more subtle than that. But note that, by adopting a proportion or phantoms of less than 50%, the clock-face timetable is retained, but the regular time distance between trains is lost.)

Another genuine difference is that, for a semi-metro although not for a pure metro, the line speed can (and usually will) change. But before explaining this, the actual performance needs to be analysed,

It can happen that the station stop distance, being the sum of deceleration and acceleration distance, (and equal to the station calling section length,) actually exceeds the distance between adjacent stations. This is quite usual, albeit not frequent, with High Speed Railways, where this value can be quite large (14.9km = 9.2miles for line speed 269kph = 167mph, 40tph, and 22km = 13.7miles for line speed 327kph = 203mph, 32kph). In such cases the values have to be calculated explicitly and input manually (in calculating journey times, for instance).

But the station stop distance has reached only 2.21km = 1.37miles by the time line capacity reaches its maximum at 103.71kph = 64.45mph, so, provided that there are no stations closer together than this distance, the calculation process for metros described above can be used with confidence. For the capacity value we’re really interested in, 60tph, the minimum inter-station distance is only 1.25km = 0.78 miles. (If there are pairs of stations closer that the above distance, then the calculations are only valid for lower line speeds, which yield the lower inter-station distance.) I don’t think this is likely to be a problem. Referring back to the earlier table, for the line speed corresponding to 60tph in the Medium Speed range, the minimum inter-station distance is 3.92km = 2.44 miles. This might be more of a problem.

For pure metros, every train is a stopping train. Alternatively, every capacity slot sub-stream is a stopping one. There are really only three capacity values worth considering, 48, 50 and 60tph. (As usual, my best choices are in red.)

Line Cap-acity (tph)	Slot time (sec)	Line Speed (m/s)	Line Speed (kph)	Line Speed (mph)	Minimum Inter Station Distance (km / miles)	Slot Stream Advance (integer Slots)	Station Wait Time (sec)	Clock-Face Timetable (every ↓ min)
60	60	21.63	77.88	48.39	1.25 / 0.78	2 / 3 / 4 / 5 / 6	62 / 122 / 182 / 242 / 302	2 / 3 / 4 / 5 / 6
50	72	14.41	51.89	32.24	0.55 / 0.34	2	106	2.4 = 2m24s
48	75	13.49	48.58	30.19	0.48 / 0.30	2 / 3 / 4	114 / 189 / 264	2½ / 3¾ / 5

For capacity 50tph, the value chosen does actually give what is technically a clock-face timetable, of 2 sub-streams of 25tph.

For a pure metro, it is possible to vary the station wait times, by varying the number of platforms. It may well be that the anticipated passenger volumes at one particular station make a longer wait time desirable at that station. An integer number of time slots can be added to the slot stream advance (and therefore to the number of sub-streams) and to the station wait time. But this number is not arbitrary: the slot stream advance must still be an integer sub-multiple of the line capacity, since this is one of the necessary conditions for a viable timetable. So this is not a wonderful facility, but may be of use in special cases. (60tph has possible values of 2,3,4,5 and 6!).

For metros and semi-metros, station wait times of around 1 minute are generally acceptable; anything over say 2½ minutes is almost certainly too long.

Readers may fairly wonder why anyone should care about clock-face, with frequencies of around one every minute. But the **line** frequency is spread over several platform faces, each of which enjoys only a fraction of the overall value, and a clock-face timetable applies to a particular platform face. So they remain valuable, even at metro level.

Note that the values (apart from the slot time, of course,) are all strikingly different from those encountered in the previous section for those capacities. Naturally so, for these line capacities are points on the ascending side of the capacity graph, whereas the previous ones are all from the other side of the hump.

The technique of scheduling more capacity than in fact is going to be used, and leaving half or more of it as phantom slots, was introduced in the previous section. This is more common in metros, which frequently need to split their services. Almost all metros show a wide variation in the capacities actually used, between the central core and the outer ends.

One point needs careful clarification. As has been explained previously, adjacent trains, from adjacent sub-streams, cannot both be diverging, except for the Low Speed range, for which the deceleration time is less than the slot time, so that a stopping train has already reached the station before the following train reaches the start of (physical) station loop. This is perfectly correct, but applies only to (High and Medium Speed) **station loops**, which consist of a **single track** diverging from the main line, allowing overtaking. This is the situation with all semi-metros, but it does not apply to **pure** metros, **in any speed range**. The way trains on a pure metro approach a station has been described in the third paragraph of the present section. The critical point is that alternate trains diverge **along alternate tracks**; a particular train has therefore always reached the station well before the next train due to use the same **approach track** (not the same **platform face**) has even reached the dividing points.

Metro stuff is effectively impossible to visualise for reasons I did my best to explain at the end of the previous section. All I can do here is to state the maths. Sorry.

Line capacity = 60tph. Slot time = 60sec. Slot length = 1.2980km.
Line speed = 21.63m/s = 77.88kph = 48.39mph.
Deceleration time = 43.27sec = 0.7211 slots Deceleration distance = 0.4680km = 0.3606 slots
Acceleration time = 72.11sec = 1.2019 slots Acceleration distance = 0.7800km = 0.6009 slots
Station Calling Section travelling time = 115.38sec = 1.9230 slots.
Station Calling Section distance = 1.2480km = 0.9615 slots

By the time the stopping train has travelled the length of the station calling section, its empty slot, which it gave up on entering the section, has travelled 1.9230 slots (time or distance). It is thus 0.9615 slots (time or distance) beyond the end of the section. In order to make this 1 slot exactly, it must travel a further 0.0385 slots (time or distance). This implies that the train must wait for 0.0385 **time** slots = 2.30sec at the station.

We now follow the progress of slot stream and stopping train.

Time Slot 1:
The empty slot advances 1 slot. It is now 1 slot beyond start of section and also 0.0385 slots beyond end of section. The slot containing the next stopping train has arrived at start of section.
The train decelerates for 0.7211 time slots and reaches the station, where it waits for 0.0385 time slots. Finally, it departs from the station and accelerates for 0.2404 time slots. It thus has 0.9615 time slots of acceleration still to do.

During Time slot 2:
The empty slot advances a further 0.9615 slots. It is now 1.9615 slots beyond start of section, and **1 slot exactly beyond end of section.** The slot stream advance is thus 1.
The second (empty) slot advances 0.9615 slots. **It therefore coincides precisely with end of section.**
The train accelerates for a further 0.9615 time slots, thus completing its acceleration. **It therefore coincides precisely with end of section and is travelling at line speed.**
(The second stopping train also decelerates, waits and accelerates for 0.7211, 0.0385 and 0.2019 time slots respectively, a total of 0.9615 time slots – but so what?)

This describes a pure metro. All the trains are in a single slot stream, using a single platform. There is no overtaking; there cannot be, as the trains are all on the same track. The capacity slot stream, being a purely virtual concept, is in no way obstructed by this, and advances by one slot (time or distance) as against the stopping train.

This is, of course, a totally unrealistic situation. While it would actually work, a station wait time of under 3sec is no use for anything. The wait time and the slot stream advance must therefore be increased by (at least) 1 slot time (or distance, for the slot stream advance). This means that the slot stream is now divided into 2 sub-streams, and both platform faces are now in use. One of the sub-streams is for stopping trains, all using the same platform, (this is not necessarily so, but it's easier to imagine that way,) with a station wait time of 62.31sec, and the other is either also stopping, with likewise a station wait time of 62.31sec, in which case we have a pure metro, or non-stop, overtaking, in which case we have a (Low Speed) semi-metro.

It is instructive to demonstrate this for a High Speed (Pure) Metro also. In fact, the values are identical with the overtaking case, pp.20-21. However, it is given again, here, but adding the extra 1 time slot to make the slot stream advance 4, from the beginning, rather than as a post-hoc correction.

Taking the capacity value of 32tph:

Line capacity = 32tph. Slot time = 112.5sec. Slot length = 10.2146km.
Line speed = 90.80m/s =326.87kph = 203.11mph.
Deceleration time = 181.59sec = 1.6142 slots Deceleration distance = 8.2441km = 0.8071 slots
Acceleration time = 302.66sec = 2.6902 slots Acceleration distance = 13.7401km = 1.3451 slots
Station Calling Section travelling time = 484.25sec = 4.3044 slots
Station Calling Section distance = 21.9842km = 2.1522 slots

By the time the stopping train has travelled the length of the station loop, its empty slot, which it gave up on entering the loop, has travelled 4.3044 slots (time or distance) on the main line. It is thus 2.1522 slots (time or distance) beyond the end of the (virtual) station loop. In order to make this 3 slots exactly, it must travel a further 0.8478 slots (time or distance). This implies that the train must wait for 0.8478 **time** slots = 95.38sec at the station. But we know in advance that this would give a slot stream advance of 3, when the value required is actually 4. So take that value ab initio, thus station wait time = 1.8478 time slots, = 207.88sec.

We now follow the progress of slot stream and stopping train.

Time Slot 1:
The empty slot advances 1 slot along the main line.
The train decelerates for 1 time slot. It thus has 0.6142 time slots of deceleration still to do.

Time slot 2:
The empty slot advances a further 1 slot along the main line. It is now 2 slots beyond start of section.
The train completes its deceleration in 0.6142 time slots, and reaches the station. It waits there for 0.3858 time slots. It thus has 1.4620 time slots still to wait.

Time Slot 3:
The empty slot advances a further 1 slot along the main line. It is now 3 slots beyond start of section. It has also passed the end of the (virtual) station section; it is 0.8478 slots beyond end of section.
The train waits at the station for 1 time slot. It thus has 0.4620 time slots still to wait

Time Slot 4:
The empty slot advances a further 1 slot along the main line. It is now 4 slots beyond start of section, and 1.8478 slots beyond end of section.
The slot containing the next stopping train **for that particular platform** has arrived at start of section.
The train waits at the station for a further 0.4620 time slots. It then departs, performing the first 0.5380 time slots of its acceleration. It thus has 2.1522 time slots of acceleration still to do.

Time Slot 5:
The empty slot advances a further 1 slot along the main line. It is now 5 slots beyond start of section, and 2.8478slots beyond end of section.
The second empty slot advances 1 slot.
The train performs 1 time slot of acceleration. It thus has 1.1522 slots of acceleration still to do.
The second train decelerates for 1 time slot. It thus has 0.6142 time slots of deceleration still to do.

Time Slot 6:
The empty slot advances a further 1 slot along the main line. It is now 6 slots beyond start of section, and 3.8478slots beyond end of section.
The second empty slot advances a further 1 slot. It is now 2 slots beyond start of section.
The train performs 1 time slot of acceleration. It thus has 0.1522 time slots of acceleration still to do.
The second train completes its deceleration in 0.6142 time slots, and reaches the station. It waits there for 0.3858 time slots. It thus has 1.4620 time slots still to wait.

During Time Slot 7:
The empty slot advances a further 0.1522 slots along the main line, to 6.1522 slots from start of section and the (corrected) slot stream advance is **4 slots exactly beyond end of section.**
The second empty slot advances a further 0.1522 slots along the main line. It is now 2.1522 slots beyond start of section. **It therefore coincides precisely with end of section.**
The train performs its final 0.1522 time slots of acceleration. **It therefore coincides precisely with end of section and is travelling at line speed.**
(The second train also waits at the station for 0.1522 time slots – but so what?)

Summarising: Each train takes 6.1522 time slots to call at the station and re-join the main line, thus:

deceleration	1.6142 time slots,
station wait	1.8478 time slots,
acceleration	2.6902 time slots,
`	6.1522 time slots in total.

Note that this applies to every train in each of the 4 sub-streams. Each train occupies the capacity slot given up by the next train in its own sub-stream.

This could, theoretically at least, be operated with just 2 platform faces. Imagine that trains in sub-streams 1 and 3 both used the same platform. So the second train would decelerate for 1 time slot in Time Slot 3, above, while the first train was waiting at the station, and reach the station, 0.6640 time slots into Time Slot 4, (0.6142 – 0.4620 =) 0.1522 time slots = 17.12 sec after the first train had departed. Tight scheduling is, after all, what Same Speed is all about. I merely point out the possibility, without (as yet) a recommendation!

Change of Line Speed, for a Semi-Metro

A metro service has the same line capacity throughout, (60tph – that's the most appropriate; there's no reason to use any other value,) but the line speed changes (from 48 to 86mph) when it becomes a Medium Speed semi-metro. It may switch directly to this from pure metro, but there will usually be an intervening section of Low Speed semi-metro, (if there are inter-station distances less than the minimum for the Medium Speed case; if there aren't, there's no reason not to switch directly).

The change of line speed will normally (certainly for a metro route) take place at a station. Trains decelerate from 48mph on the Low Speed side, to stop at the station, then accelerate to 86mph on continuing beyond the station. It is possible, though unlikely, that some trains will not stop at the station. For these, non-stop trains, they approach the station at 48mph, then, on passing the station, begin the acceleration to 86mph. The capacity slot stream itself accelerates, beginning at the station, from the lower to the higher speed; slot stream and non-stop train thus accelerate in lockstep, the change in speed makes no difference whatever to their relationship. In the reverse direction, slot stream and non-stop train decelerate in lockstep, reaching the lower speed as they together pass the station, and then continue together at the lower speed. This is obvious, for non-stop trains, but in fact the equivalent effect applies to trains which call at the station also. The acceleration from lower to higher line speed takes place as the final component of the train's acceleration to the higher speed. In the other direction, the deceleration to the lower speed takes place as the first component of the train's deceleration to zero at the station. Train and slot stream no longer proceed in lockstep, as with the non-stop case, (the acceleration of the slot stream takes place as the first component, immediately on passing the station, and in the other direction the deceleration takes place as the last component, immediately before,) but the effects are identical. What this means, is that this speed-change portion can be completely ignored in the calculations, since it makes no difference whatever to the relationship between train and slot stream. In effect the line speed, and hence the speed of the slot stream, changes instantaneously on passing the station, and is thus constant, although at different values, on both sides. This further means that the line capacity and thus the slot time is constant throughout the calculation, which, strictly speaking, is not the case for the section where the line speed changes. (The slot **length**, being the distance travelled at constant line speed in the slot time, is obviously different for different line speeds.) This makes a great simplification in what would otherwise be a rather intractable problem. But note that it is still necessary to convert the **distance** values on one side of the station, so that they are all expressed relative to the same speed throughout (this is to ensure that the fundamental relationship, that the time taken by the train to travel over the station loop is twice the time taken by the empty slot to travel the same distance, at the chosen speed, still applies). Note also that the equalisation time, applied to **integralise** the slot stream advance, is different for the two directions.

Note that it is perfectly okay to speak of station loops for a semi-metro, since overtaking **is** involved.

This may well sound like a computational three-card trick, so here are the actual calculations, for the usual 60tph case.

Line capacity = 60tph. Slot time = 60sec.

Line speed (Low Speed value, V_L) = 21.63m/s = 77.88kph = 48.39mph (Slot length)$_L$ = 1.2980km
(Medium Speed value, V_M) = 38.37m/s = 138.12kph = 85.83mph (Slot length)$_M$ = 2.3020km

Low to Medium Speed (V_L => V_M):

Deceleration time (V_L => 0) = 43.27sec = 0.7211 time slots
Deceleration distance = 0.4680km = 0.3606 (distance slots)$_L$ = 0.2033 (distance slots)$_M$
Acceleration time (0 => V_M) = 127.89sec = 2.1315 time slots
Acceleration distance = 2.4533km = 1.0657 (distance slots)$_M$
Acceleration time (V_L => V_M) = 55.78sec = 0.9297 time slots
Acceleration distance = 1.6733km = 0.7269 (distance slots)$_M$

Station Loop travelling time = 171.16sec = 2.8527 time slots
Station Loop distance = 2.9233km = 1.2699 (distance slots)$_M$

By the time the train has travelled the station loop distance, the empty slot has travelled 2.8527 slots (time or (distance)$_M$). Thus the empty slot is 1.5828 slots (time or (distance)$_M$) beyond end of loop. This implies that the train must wait for 0.4172 time slots (= 25.03sec) at the station to make the slot stream advance = 2, exactly. This was calculated ignoring the speed change portion, but in following the progress of train and slot stream, there is no reason why the speed change portion should not be included

Time Slot 1:
The empty slot advances 0.2033 time slots to reach the station, (it actually advances 0.4680km, which is 0.3606 time slots at V_L or, equivalently, 0.2033 time slots at V_M,) then accelerates (from V_L) for 0.7967 time slots. It thus has a further 0.9297 – 0.7967 = 0.1330 time slots still to accelerate (to reach V_M). It is now 1 time slot from start of loop.
The train decelerates for 0.7211 time slots to reach the station, where it waits for 0.2789 time slots. It thus has a further 0.4172 – 0.2789 = 0.1383 time slots still to wait.

Time Slot 2:
The empty slot accelerates for 0. 1330 time slots and reaches V_M. It then advances a further 0.8670 time slots at V_M. It is now 2 time slots from start of loop and 2 – 1.2699 = 0.7301 time slots beyond end of loop.
The train waits for 0. 1383 time slots, then accelerates for 0.8617 time slots. It thus has 2.1315 – 0.8617 = 1.2698 time slots still to accelerate.

Time Slot 3:
the empty slot advances a further 1 time slot at V_M. It is now 3 time slots from start of loop, and 1.7301 time slots beyond end of loop.
The train accelerates for 1 time slot. It thus has 0.2698 time slots still to accelerate.

During Time Slot 4:
The empty slot advances a further 0.2699 time slots at V_M. It is now 3.2699 time slots from start of loop and 2 time slots precisely beyond end of loop. The slot stream advance is thus 2.
The train accelerates for the remaining 0.2698 time slots, achieving line speed V_M precisely as it reaches end of loop. (We will brazenly ignore the rounding error of 1 part in 2700!!)

A further 2 time slots need to be added to the slot stream advance and to the station wait time, to make the slot stream advance and thus the number of sub-streams 4, as required by this line speed. The station wait time is thus 2.4172 time slots, = 145.03sec.

Medium to Low Speed (V_M => V_L):

Deceleration time (V_M => 0) = 76.73sec = 1.2789 time slots
Deceleration distance = 1.4720km = 0.6394 (distance slots)$_M$ = 1.1341 (distance slots)$_L$
Acceleration Time (0 => V_L) = 72.11sec = 1.2019 time slots
Acceleration distance = 0.7800km = 0.6009 (distance slots)$_L$
Deceleration time (V_M => V_L) = 33.46sec = 0.5577 time slots
Deceleration distance = 1.0040km = 0.4361 (distance slots)$_M$ = 0.7735 (distance slots)$_L$

Station Loop travelling time = 148.84sec = 2.4807 time slots
Station Loop distance = 2.2520km = 2.0780 (distance slots)$_L$

By the time the train has travelled the station loop distance, the empty slot has travelled 2.4807 slots (time or (distance)$_M$). Thus the empty slot is 0.4027 slots (time or (distance)$_M$) beyond end of loop. This implies that the train must wait for 0.5973 time slots (= 35.84sec) at the station to make the slot stream advance = 1, exactly. This was calculated ignoring the speed change portion, but in following the progress of train and slot stream, there is no reason why the speed change portion should not be included

Empty slot and train both decelerate from V_M to V_L taking, obviously, the same time – 33.46sec = 0.5572 time slots. The empty slot completes this deceleration at the point when it reaches speed V_L, simultaneously with reaching the station. It continues beyond the station at constant speed V_L. The train begins its deceleration immediately on departing from start of loop, It continues to decelerate beyond V_L down to zero, which it reaches at the point when both it and the empty slot reach the station simultaneously. Therefore, the time advanced by the empty slot from leaving the start of loop to beginning its deceleration is precisely the same as that taken by the train to decelerate from V_L to zero.

Time Slot 1:
The empty slot advances 0.7211 time slots at speed V_M, then decelerates for 0.2789 time slots. It thus has 0.5577 – 0.2789 = 0.2788 time slots still to decelerate. It is now 1 time slot from start of of loop.
The train decelerates for 1 time slot. It thus also has 0.2788 time slots still to decelerate.

Time Slot 2:
The empty slot decelerates for 0.2788 time slots, and reaches tha station as its speed reaches V_L. It then advances a further 0.7212 time slots at speed V_L beyond the station. It is now 2 time slots from start of loop.
The train performs its final 0.2788 time slots of deceleration, reaching the station and stopping there. It waits at the station for 0.5973 time slots, then accelerates for 0.1239 time slots. It thus has 1.2019 – 0.1239 = 1.0780 time slots still to accelerate.

Time Slot 3:
The empty slot advances 1 time slot at speed V_L. It is now 3 time slots from start of loop and 0.9220 time slots beyond end of loop.
The train accelerates for 1 time slot. It thus has 0.0780 time slots still to accelerate.

During Time Slot 4:
The empty slot advances 0.0780 time slots at speed V_L It is now 3.0780 time slots from start of loop, and 1 time slot precisely beyond end of loop. The slot stream advance is thus 1.
The train accelerates for the remaining 0.0780 time slots, achieving line speed V_L precisely as it reaches end of loop.

An extra time slot is added to the station wait time and to the slot stream advance, since the slot stream advance and thus the number of sub-streams must be 2 for this line speed. The station wait time is thus 1.5973 time slots = 95.84sec.

In performing the above pair of calculations, it does seem to make a difference on which side the distance values are converted, and it should be the side after the station. I got this right by chance the first time, going from V_L to V_M, and left it like that for the other calculation. But this turned into a very fraught process. Switching it to go the other way completely obliterated the problem, and the above results came out pretty much by themselves.

Summarising all the values of interest for metros and semi-metros at the 60tph capacity:

Line Cap-acity (tph)	Slot time (sec)	Line Speed (m/s)	Line Speed (kph)	Line Speed (mph) _	Minimum Inter Station Distance (km / miles)	Slot Stream Advance (integer Slots)	Station Wait Time (sec)	Clock-Face Timetable (every ↓ min)
60	60	21.63	77.88	48.39	1.25 / 0.78	2	62	2
					Low to	Medium	145	
60	60	38.37	138.12	85.87	3.92 / 2.44	4	138	4
					Medium to	Low	96	

Change of Line Speed for a High (or Medium) Speed Railway

This new section was added at version 4.0. It took an enormous amount of effort to write, since the subject is quite astonishingly complicated.

On reflection, I decided that it seriously unbalanced the present article. Its size alone – 14 pages – amounted to almost 30% of the total article, and even that left out quite a lot of detail which should, really, be included. Even worse, it seemed, at first sight, to contradict quite a lot of the content of earlier sections, in particular, the concept of Sweet-Speeds.

Accordingly, I have removed the subject from the present article, and re-written it in what I believe is a more rigorous exposition, covering the topics:

- Constant capacity line over a wide range of speeds.
- The invariable requirement of two deceleration tracks when reducing line speed (though acceleration is straightforward and involves no special provision).
- Deceleration Ranges, showing how the actual separation distance differs from (exceeds) the TSD(e) standard.
- The actual length of deceleration track is constant for a given type of switch, (i.e. independent of both original line speed and target line speed. Likewise the elapsed time spent on the deceleration track depends on the target line speed but not the original line speed.
- The three separate types of deceleration:
    - From a higher to a lower speed range
    - From a higher to a lower speed within the same speed range
    - The esoteric cases when the TSD(b) standard must be used.
- Examples of each of the above types of deceleration.

This re-write is published in a new article – 'Line Capacity vs. Speed for Same Speed Railways – Volume 2, version 1.0 of which is published simultaneously with version 4.1 of the present article, providing continuing availability of the section.

Volume 2 is intended for the more advanced topics, of which change of line speed is incontrovertibly one such.

Moreover, the present article, (the de facto volume 1,) was originally written considering just High Speed railways, and although it now states, frequently, that the theory applies in all speed ranges, it still bears evidence of its origins, most particularly it deals with only one type of switch, the UHS switch with TLS 230kph. Volume 2 is rigorously speed-range neutral, and its first substantive section introduces a whole range of switch types, which extend the range of availability of the TSD(e) standard over the widest possible range of speeds.

Terminal Stations

Terminal stations are the real capacity (and other) bugbear, at least, large terminal stations in London are, like, for example, Euston. Each platform of a terminal station can handle only 2tph – 20 minutes to unload, service and reload the train, plus 10 minutes contingency. Attempting to satisfy the entire load of a HS line in a single terminal station, as HS2 Ltd., with their lunatic plans for Euston, vaingloriously assert they can do, is a catastrophe in the making. An acceptable level of capacity can be provided, in a terminal station, only by a completely unacceptable metastasis of platforms, and of station area.

But all is not lost. The correct way to design a Same Speed Railway of **any** category, but High Speed in particular, is roots – trunk – branches. Multiple services from different origins – the roots – progressively merge into a single trunk and travel the bulk of their journeys at constant high speed on the trunk. They then progressively diverge from the trunk – the branches – to reach their destinations. Each origin and destination has only one or two services, so, even at only 2tph per terminal platform, doesn't need many platforms to accommodate them. In any case, trains don't **have** to be serviced at the station platform itself. After unloading, they could be moved to a servicing area and processed at leisure, returning to the platform in good time for their next assignment, with plenty of time to reload in comfort. Such luxuries are absent from congested metropolitan termini. The roots and branches can often, at least towards the ends, be existing classic routes. HS2 Ltd. is evidently of the considered opinion that the best place to terminate a HS line is on the trunk!!!

The solution to this farrago is to do away with terminal stations, at least, big ones in London. A new, underground, through station should be built at Euston Cross (i.e. linking Euston, St. Pancras and King's Cross). With station wait times of up to 10 minutes allowed, 3 or 4 platform faces in each direction should be sufficient, with a single pair of approach tunnels. Services pass underneath London and out to the other side, fanning out to serve several terminal destinations, such as Maidstone, Gillingham/Rainham, Dover, Margate and Eastbourne, each of which, being served by only a fraction of the total, would need little if any new infrastructure.

The above statement of station capacity was my original estimate, pulled out of the air, more or less. But I can do a lot better now. The section between Old Oak Common and Stratford HL (South) simply needs to be operated as a 32tph, constant-capacity line, running at the reduced line speed of 20.11m/s, i.e. 45mph, which can be accommodated by perfectly ordinary switches – no special trackwork needed at all. The results are now given. They may look (and, indeed, are) very similar to those in the previous section, but note that that describes an overtaking situation, whereas this is a pure metro, similar to that in the section on metros, starting on p.31. The terminology is thus slightly different.

Line capacity = 32ph. Slot time = 112.5sec. Slot length = 2.2627km.
Line speed = 20.11m/s = 72.41kph = 45mph.
Deceleration time = 40.23sec = 0.3576 slots Deceleration distance = 0.4045km = 0.1788 slots
Acceleration time = 67.04sec = 0.5959 slots Acceleration distance = 0.6742km = 0.2980 slots
Station Calling Section travelling time = 107.27sec = 0.9535 slots.
Station Calling Section distance = 1.0787km = 0.4767 slots

By the time the stopping train has travelled the length of the station calling section, its empty slot, which it gave up on entering the section, has travelled 0.9535 slots (time or distance) on the main line. It is thus 0.4767 slots (time or distance) beyond the end of the section. In order to make this 1 slot exactly, it must travel a further 0.5233 slots (time or distance). This implies that the train must wait for 0.5233 **time** slots = 58.87sec at the station.

We now follow the progress of slot stream and stopping train.

Time Slot 1:
The empty slot advances 1 slot along the main line. It is now 1 slot beyond start of section, but has also (already!) passed the end of section. It is 0.5233 slots beyond end of section.
The slot containing the next stopping train arrives at start of section.
The train decelerates for 0.3576 time slots, and reaches the station, where it waits for 0.5233 time slots. It then accelerates for 0.1192 time slots. It thus has 0.4767 time slots of acceleration still to do.

During Time slot 2:
The empty slot advances a further 0.4767 slots. It is now 1.4767 slots beyond start of section, and **1 slot exactly beyond end of section.** The slot stream advance is thus 1.
The second (empty) slot advances 0.4767 slots. **It therefore coincides precisely with end of section.**
The train accelerates for a further 0.4767 time slots, thus completing its acceleration. **It therefore coincides precisely with end of section and is travelling at line speed.**
(The second stopping train also decelerates for 0.3576 time slots, and then waits at the station for 0.1191 time slots, a total of 0.4767 time slots – but so what?)

The slot stream advance is only 1; it needs to be 4 to be consistent with the original Same Speed value. (Like the slot time, the slot stream advance and thus the number of capacity sub-streams, must be constant throughput for constant capacity over different line speeds.) The station wait time must therefore be 3.5233 time slots, = 396.37sec, = 6min36.37sec.

A station wait time of 396sec would be far too long for a normal (albeit Same Speed) metro, but perfect for long distance High Speed Trains, **running under pure metro conditions** to cross London. Between Old Oak Common and Stratford, a line speed of 45mph would be perfectly adequate.

In principle, (and in actuality, for some of the routes I have considered,) we could reach a situation where all trains pass through (i.e. underneath) London, and none actually starts or terminates there. We thus face the possibility that most or all of the existing terminal stations in London, and possibly in other metropolitan areas, could become redundant, while actual rail travel went on increasing, facilitated by greatly increased available capacities. I merely flag this up as a possibility; it isn't going to happen next year or even next decade. But we should begin now to consider worthy alternative uses for some of the finest architecture in the country. (I think Liverpool Street is likely to be the first to become available.)

In practice, however, the rush hour would presumably not have gone away. There would still be a need for extra capacity at these times, and the existing terminal stations would provide this, leaving the inter-regional, cross-London connections carrying an essentially even base load throughout the day. The terminal stations would not be required for railway purposes outside the periods 7:30 – 9:30 and 16:30 – 18:30, say. So the above remark on finding worthy alternative uses for them still applies. My own initial ideas are for staging artistic and cultural events and small exhibitions in the passenger circulating areas. Most of the infrastructure required – cafes, toilets and retail units – is already there. The opening times for such events would be 10:00 – 16:00 and 19:00 onwards, Monday – Friday, (and all day at weekends,) to give time for set-up and dismantling after and prior to rail use, since the passenger circulating areas should certainly not be obstructed during rush hours.

Resilience

The focus of this article has been all about determining maximum values of line capacity, and about how such capacity levels are actually achieved. The results have all been of every last capacity slot in use. It would be very inadvisable trying to run a railway (**any** railway!) under such conditions, though, on reflection, that is the way conventional railways always have been run. **Safety** standards have been at a high levels for a long time now, but **resilience** is unknown and unprovided for. For Same Speed Railways, where everything and in particular all operating conditions are at the outer limits of tolerance, no service could reasonably be attempted under such conditions. The slightest divergence from the timetable would cause the whole system to come crashing down. There is, in other words, no built-in resilience.

As pointed out, traditional, mixed traffic railways always have operated with minimal (no!) resilience, and, every so often, suffer serious interruptions of service. The huge capacities offered by Same Speed Railways would magnify hugely the effect of a service interruption. So we don't have the luxury of ignoring resilience.

Realistically, there is only one sort of problem from which we can recover by instituting resilience – when a stopping train for some reason misses its scheduled restart time from a station. This can happen for all sorts of no-fault reasons, such as a passenger being taken suddenly ill, and the train having to be held waiting for an ambulance. This is actually very straightforward to defend against, and defended against it must be, since the train is immediately obstructing the following train and all subsequent trains in its own capacity slot sub-stream and in all other sub-streams (since the sub-streams are purely a logical abstraction – there is only a single physical track in each direction, occupied by all of them).

We consider first the overtaking categories. Assume that we have just the one stopping sub-stream, which is the most likely case.

The most direct approach is to assign a second sub-stream as a stopping sub-stream, but with no trains actually assigned to it; it serves purely as a resilience sub-stream for the actual stopping service. As soon as it is clear that there is a problem, and that a train is going to miss its restart time, the sub-stream is switched over to the other platform face, and the following trains are then no longer impeded. The capacity slot given up by the following stopping train, which ordinarily would have been taken over by the train with the problem, simply continues, empty thereafter. This is hardly a problem. The train with the difficulty simply stays in the platform until the difficulty has been resolved, and then departs, taking over the next slot (all of which are normally empty) in the resilience sub-stream, and travels in that sub-stream for the rest of its journey. There is of course no scope for it to regain any of the lost time; that is inevitable in the tightly-scheduled world of Same Speed railways, where no resources are wasted (as opposed to deliberately assigned for resilience). This is, I believe, a gratifyingly robust and efficient solution, but very expensive in line capacity, taking up, typically, 25% of the total.

A less extravagant alternative, for High Speed and Medium Speed railways, which rarely operate at full capacity in the stopping sub-stream, is to use the free capacity in that sub-stream for resilience. The stations are likely to be very adequately served by a train every 15 minutes. For my favourite line capacity of 32tph, with 4 sub-streams, that means 4 out of the 8 slots per hour in the stopping sub-stream are empty. Suppose the trains use both platforms at the stations alternately. If a train misses its restart time, the next three slots **at that platform** are empty, so it has another 2 chances of taking up a slot (the middle slot of the 3 corresponds to the train in the other platform). Then the next train arrives, uses the other platform, (because our train is blocking the one it would have used,) and departs at the correct time. Our train now has a further 2 opportunities to get moving. All the other trains continue to use the other platform. Even if only 50% of the slots are empty, that still means that the train has a free slot available every second slot in its sub-stream. Even if the train has broken down, then, provided it did actually reach a station, and provided also that a second train doesn't fail at the same station, the service continues running indefinitely. The failed train can be removed by emergency services overnight.

The slot window has been mentioned earlier. This gives a certain resilience in that it allow a train to depart any time during the ~1 minute before its scheduled departure time, and thus have more leeway in joining its new slot.

The above remarks apply to High and Medium Speed railways with overtaking; in fact they apply to semi-metros in all speed ranges. For pure metros, the situation is very different. For these, we don't bother at all with resilience, in the sense of trying to **recover** from an incident. Instead, we simply cancel the affected service. This may sound staggeringly cavalier, but the reasons are decisive.

Pure metros have very few capacity sub-streams, typically just two. Low Speed metros **invariably** have just two (in scheduled use). So sacrificing capacity for resilience is not an option.

At least one, and preferably two extra platform faces in each direction are added at each station. They are not normally in use. It a train hits a problem at a station, the slot sub-stream of which it is a member simply switches over to one of these resilience platforms, which is then uses indefinitely, for the rest of the day if need be. The train which missed its slot is cancelled, and left where it is, until the end of the day's service, if need be. The passengers on that train simply move to another platform and catch another train, an inconvenience certainly, but at the service frequencies provided, not a huge problem. The cancelled train has to sit it out because there are no scheduled free slots in a (pure) metro service.

Better still, (and the main reason for having **two** extra platform faces,) **both** slot sub-streams switch across to the reserve pair of platform faces, on opposite sides of the same island platform. Otherwise the service would be split between non-contiguous platform faces, requiring stair access to get between them. (Of course, in the disastrous but surely rare event of **two** failed trains at the same station, non-contiguous platforms would have to be used, since the second failure would be on one of the second pair of platforms, but the full service would still operate unless and until a **third** train failed there.)

In practice, it is unlikely that a metro will be operating at maximum capacity all the time, throughout the whole day, so there usually will be opportunity to recover the failed train before the end of the day's service.

Resilience is a very important subject. This section acknowledges its importance and suggests possible approaches.

Conclusions

The first five sections of this article demonstrate how line capacity relates to line speed, for Same Speed railways, as they have done for several years; the first version of this article appeared in April 2017. It summarised, for a non-specialist audience the information which had been appearing and accumulating in the ‘Same Speed Railways’ article since February 2016, when that article’s Appendix B first appeared, totally changing the character of the article and turning it into the fundamental technical authority for the totality of my proposals, and, as a consequence, distinctly user-hostile, for a non-specialist.

I have always particularly liked the present article; I probably enjoyed writing it even more than any other. But I had always, until v3.0, felt a certain uneasiness about it, that I had not satisfactorily established a sufficiently firm foundation for my arguments. The issue was, of course, that the capacity values I had produced were, at first sight, unbelievably high, as indeed they are, compared with conventional, mixed traffic railways. So I had always been chary of introducing it prematurely to hard-bitten railwaymen, and inviting (quite understandable) ridicule.

I decided that this could not continue, so, since April 2020, I have been studying timetable building for Same Speed railways, how to schedule their services. It took quite a while to get a grip on the subject, and I followed several false leads, but eventually I gained command of it, and produced the three sections starting with ‘Timetabling Considerations and Sweet-Speeds’, and also the penultimate section on resilience, and version 3.0 was published at the end of October 2020. It occurred to me that, although I had always said that, while Same Speed railways had a constant line speed, (hence the name,) it was possible that this could, in certain circumstances, vary between different sections (albeit constant within each section). This I had never really thought about properly, and, it seemed, now was a good time to try. This in fact turned out to be by no means straightforward, and it has taken a further 7 months and another three versions in January 2021 to get it to my satisfaction. This Conclusions section was added then, at version, v3.3; it seemed a good idea to tidy up the ending, and give a hint of further attractions to come.

The section on Change of Line Speeds, or, more particularly, maintaining the same line capacity, thus capacity slot sub-streams and thus, further, stopping patterns, in sections of different values of line speed, has turned out to be such a major enhancement as to justify a new version level, so quickly after the previous raise, so this is now version 4.0. In addition, following discussion with Vossloh Cogifer, the buffer component of the Train Separation Distance Standards, basic and extended, has been reworked to take into account the movable parts of switches, and also the resetting of the switch behind a diverging train. As a consequence, this item now has a rational foundation, no longer just a rule-of-thumb. Full details are given in the section ‘Consequences of the Results’ on p.12. This increased the buffer size, and this in turn required the reworking of all the numerical results. One unfortunate consequence of this has been the loss of a very popular capacity, 64tph, which is simply no longer available. All this explains why it has taken so long to get this latest version ready for publication. Having added this new section, I decided that it unbalanced the entire article. Accordingly, at v4.1 it was removed and published instead in a new article, Volume 2 of the present one, intended for advances topics, which Change of Line Speeds undoubtedly is. It has been replaced, in the current article by a stub, explaining what has happened.

I assert that, in determining the scheduling of Same Speed railways, and the preservation of capacity across different line speeds, that I have justified and validated the capacity values presented. I can now define the behaviour of a train at every stage of its journey, to a split second: the speeds, locations where speed changes and how it changes, the wait time at a station and so on. Same Speed railways, when operating to maximum capacity, are so tightly scheduled that there is no spare time anywhere in the schedules. As a consequence, many of the features of conventional railways that one simply takes for granted and hardly thinks about, no longer apply. Thus, there is no free choice of line speed, only certain 'Sweet-Speeds' are available. The wait times of trains at stations are prescribed to a split second. A train may not simply decelerate at will, because the following train is already as close as it is permitted to be; the main line has to split in two, or a loop line diverge, to allow the train to get out of the path of the following train. For High Speed railways, a special extended train separation standard is necessary to allow a train to decelerate down to the turnout limit speed on the main line, before it can start to diverge.

I believe I have fulfilled my promise, in the Summary section, that I would justify the capacities claimed for Same Speed railways.

What I have produced is, in effect, a first version of the Functional Specification of the automatic train control system; these are the things which it needs to be able to do. That challenge has to be dealt with by the people who design and build ATCs, it's not something I can do, but I can and will tell them what they need to provide.

Frankly, the line capacities on offer for Same Speed railways are, except for a few very special cases (such as the central core sections of metros, Crossrails especially), more than we can readily find real trains to pack into. This is a delightful problem to have!

The approach I recommend is to **schedule** everything to maximum capacity, but leave the majority of trains phantoms. There is absolutely no penalty involved in having the majority of capacity slots running around empty, and it offers enormous convenience in operational flexibility.

I believe now is a good time to publicise this article, having addressed reasonable concerns about the promised capacities. But this is far from the end of the story. I have demonstrated the techniques required for scheduling Same Speed railways, but I haven't actually yet scheduled any – not to these latest, rigorous standards, that is.

The section on Terminal Stations has been part of the article since the beginning. I have opposed HS2 Ltd.'s plans for a terminal at Euston from the very beginning, arguing my case with everyone I could think of, and met with a consistent blank wall of indifference, even from the supposedly specialist railway magazines. Every time the matter (indeed, HS2 in general) has come up in parliament, it has simply gone through on the nod, without any effective scrutiny. Final implementation of those plans will result in not so much an immediate and palpable catastrophe, but the frivolous frittering away of a unique opportunity for a step change in quality of cross-London travel. Gradually, the realisation will dawn that Euston is a huge disappointment, grossly oversold in what it could deliver, and that perhaps we should have considered the alternatives more carefully. It is, of course, no satisfaction at all to be proved right, after the damage has been done.

The current delay to the Euston implementation, with initial termination at Old Oak Common, gives us a last chance to get it right. In this new version, the Terminal Stations section has a new argument, not available to me before: with validation of the capacities available, I can now credibly offer HS2 Ltd. 32tph for HS2, (far more than their terminal could possibly accommodate,) **and,** with constant capacity over sections of different speed, demonstrate how to get if between Old Oak Common and Stratford.

Same Speed railways are now as much about (development opportunities for) the classic network, as for High Speed. I have several ideas on these opportunities to work on. I shan't be out of work any time soon.

The current version of this article is thus by no means the end of the story.