Line Capacity vs. Speed for High Speed Railways

Introduction and Background

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.

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 dependant 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. That is **not** what the present article is about.

Same Speed Railways, of which High Speed Railways are an important category, (the only other important category at present being metro systems – 'Low Speed Railways',) have the fundamental property 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 Initially they will run at the same performance as the existing trains, until these have been progressively withdrawn, when the improved characteristics of the new stock can be taken advantage of.)

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. (It's rather more complicated than that, naturally, but that is the essence.) The important consequence of this is that same speed lines **do** have a precise value of line capacity, for a specific value of line speed; this varies as the line speed itself varies, and this can all be analysed mathematically, and its characteristics deduced. **That** is what the present article is about. It is concerned specifically with High Speed lines, because of their intrinsic importance, and also because they have certain characteristics peculiar to themselves, not present with other categories of same speed line.

The main reason for analysing line capacity is to find out how to optimise it, for some prescribed objective. Line capacities of high speed lines, likewise of metro systems, 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 high 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 all 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.)

TSD(b) has a dynamic and a static component. The dynamic component is readily calculated, knowing the average deceleration rate. The static component consists of the length of the train (assumed to be all the same, otherwise take a maximum length) and a rule-of-thumb-type buffer zone (since the train must stop some distance – the buffer zone – behind the back end of – hence the train length – the train in front). The capacity is readily calculated (I use Microsoft Excel©) for a series of line speeds, and the results plotted automatically. The graph of capacity against line speed follows.

Line Capacity vs. Speed: the Results

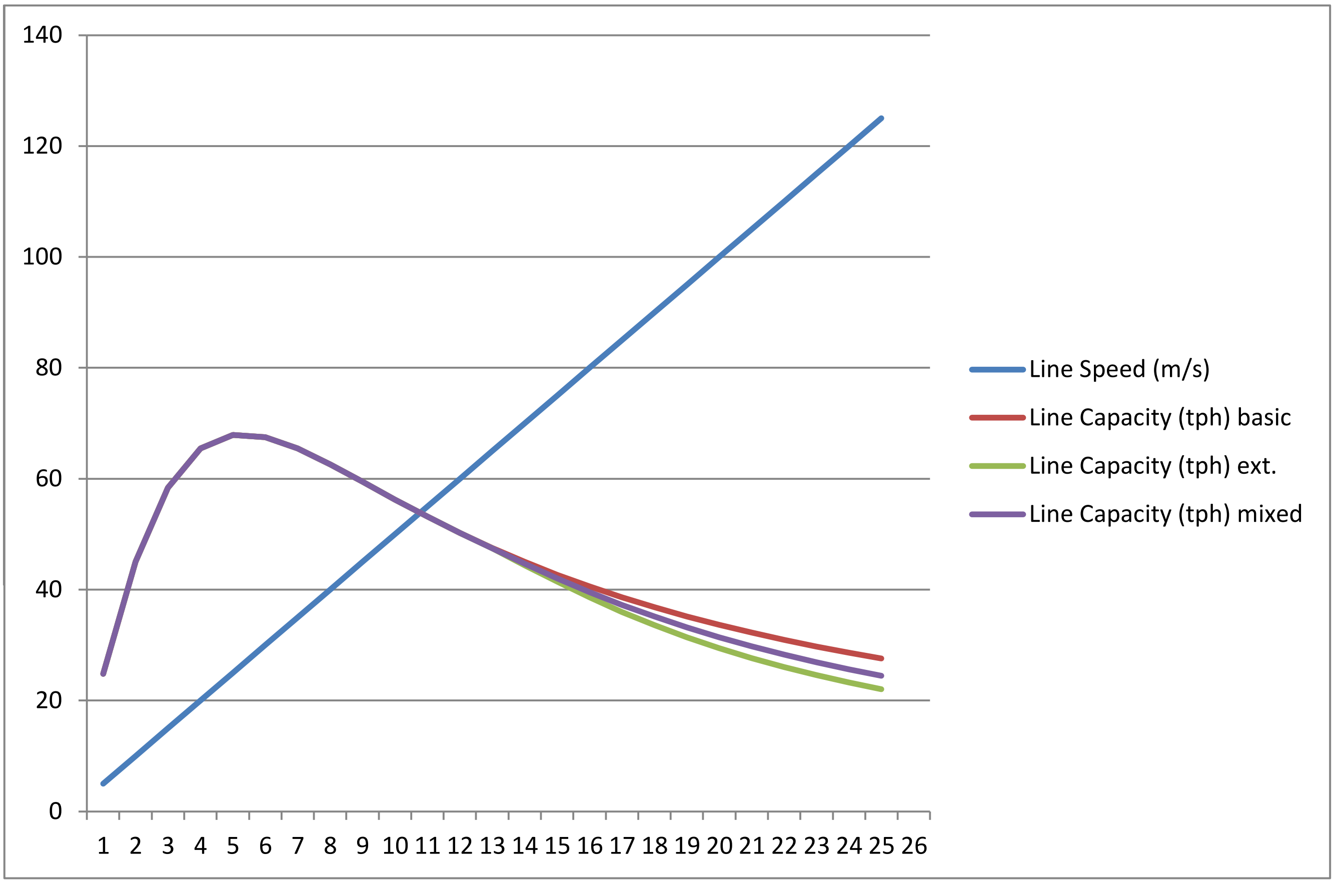

[A word of explanation for readers who may not be familiar with Excel© charts: The results are calculated for a series of line speeds, with a constant increment, of 5m/s in the above case. The values along the x-axis are the ordinal numbers in the sequence of these values, thus 1 is 5m/s, 2 is 10m/s 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. Thus the seventh result set has speed 35m/s and capacity c.65tph (actually 65.45tph). To get from m/s to kph multiply by 3.6, and to mph by 2.237.]

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 separation (the several variants of which, basic, extended and mixed will be explained shortly). It is, in fact

capacity = speed / separation.

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.

A further graph will elucidate:

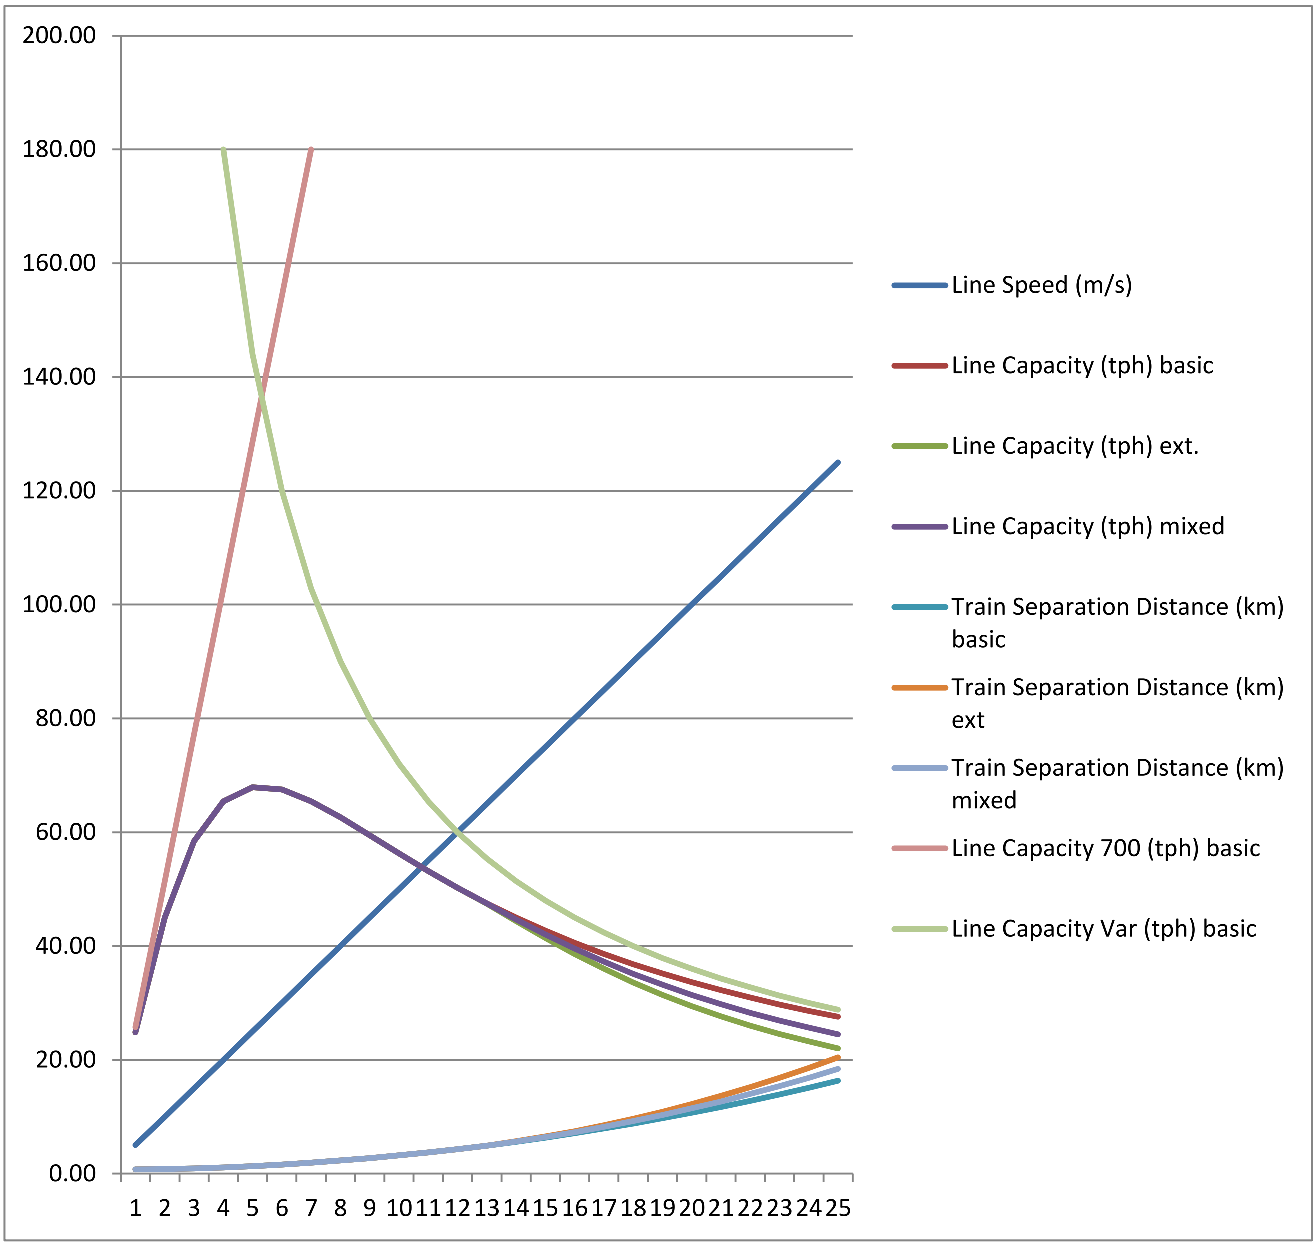

This second graph, above, adds the actual separation distance, expressed in km (thus indicating, for example, basic separations of 1.5km at 100mph, 45m/s, and 10.7km at 225mph, 100m/s).

It also adds two further quantities, the capacity value determined only by the fixed component ('line capacity 700' – the odd name alludes to the fixed component of the separation, in the current treatment, having the numerical value 700m), and the capacity value determined only by the variable component ('line capacity var').

This last 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.

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

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.

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 above graph has a very familiar shape with, at low speeds, capacity increasing rapidly with speed, until it reaches its maximum, 68tph, at the astonishingly low (to me!) speed of only c.26.5m/s, 59mph. (This may not be so surprising to metro operators, whose trains perform in precisely this speed range.) 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 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, 60tph, of a line speed of 225mph, 100m/s, which is 29tph. 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 59mph, and will certainly not be consoled by the thought that by so doing, another 67 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). I don't, personally, have the knowledge or skills to form a defensible opinion of what the maximum speed should be, but have the uneasy feeling that 225mph is probably going beyond what is justifiable, and 250mph almost certainly is. It may be that the existing operators of high speed railways have got it just about right, at 300kph, 186.5mph. As a trade-off between speed and capacity, I'm coming to the conclusion that 140mph is a good compromise – theoretical capacity c.50tph. (See also below, on maximum turnout speed.)

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 pointwork. There are no points available which allow a diverging train to do so at line speed. The fastest points currently available (at least, as at 2014,) allow for a **maximum turnout 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 points. (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 the back-end of the train has crossed the junction.)

The graph of capacity splits into three strands at the higher speeds. 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, 29tph, 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. (In practice, of course, we would always choose TSD(e) throughout, since 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 230kph, and are certainly significant, but still small. All this stuff is expounded at length in the web 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 maximum turnout 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 section on Through Stations for an overview of the Capacity Slot model. In effect this fits a phantom straight-ahead train between them.)

A strong reason why, as mentioned above, I favour 140mph as a good compromise between speed and capacity is that it is just within the current maximum turnout speed. Trains may therefore diverge from / converge with the main line at full line speed. This removes the need for the TSD(e) standard; TSD(b) suffices throughout, with major simplification. What this means, precisely, is that a following train never perceives a preceding, diverging, train decelerating **on the main line** – it has already diverged, at full line speed, before beginning its deceleration on the station loop. (This means that the entire deceleration / acceleration is now performed on the station loop, but the length of the loop is essentially unchanged (slightly reduced, in fact) from those higher-speed cases, since there only the deceleration from / acceleration to the turnout limit speed – 143mph – is performed on the loop.) It must be admitted that the journey times under this standard compare very unfavourably with the true HS line – c.50% higher – albeit still impressive compared with any mixed-traffic line.

As far as the current article is concerned, the only point of describing the above effects (some of which are decidedly esoteric) is to reassure the reader that they have been taken into account.

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

Through Stations

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.

But there is still a penalty, and it may be serious. It concerns line capacity, but doesn't directly affect the overall capacity value itself (for the precise meaning of this statement, see below). Any train in 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 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.

(Readers of a philosophical inclination may wonder precisely why there is this extra time penalty for stopping trains; after all, there is no change to the line capacity. I suggest that it is because the fundamental requirement for maximum capacity, that the traffic be homogeneous has actually been breached, in that some trains make station stops which other trains don't. I think that a more philosophically satisfying answer than an excursion through queueing theory would be.)

With this further information, the precise meaning of the earlier statement, to which attention was drawn, (second paragraph of the present section,) can now be given. Stopping trains at intermediate stations does not affect the line capacity, in that exactly the same number of slots are still available, but it may, and probably does, cause some slots to be wasted or, rather, only partially used.

In fact, (see the section 'Optimum Mixture of Non-stop and Stopping Trains' near the end of Appendix B of the 'Same Speed Railways' article for full details; the matter is seriously difficult, and I can only give a very brief summary here,) it is theoretically possible to schedule a mixture of non-stop and stopping trains such that there is no penalty in line capacity. The critical issue is to adjust the Station Stop Time (being the sum of the deceleration, station wait and acceleration times) so that it becomes equal to an integral number of capacity slot times, at that line speed. This is achieved, simply enough, by increasing the station wait time above the standard 3 minutes. Over the range of line speeds of interest (225kph – 400kph) this means that the station stop time is either 7 capacity slots, for line speeds of up to 341kph, or 6 slots above that.

The consequence of this is that the capacity slot stream on the main line is logically divided into either 7 or 6 streams, each such stream consisting of every 7th or 6th slot. Each of these streams is a (potential) station platform stop stream.at an individual station platform. (If not so used, it is simply a normal, non-stop stream.) A normal station, with 2 platforms in each direction, could be associated (in each direction) with two such streams, doubling the frequency of stopping services at that station. That is simple enough to imagine, but in fact the theory is entirely general; the station platform stop streams are completely free-standing and independent, enabling different stopping patterns and different frequencies; it can get very elaborate.

In practice, no route has services which are non-stop between origin and destination, but rather divides into sections over which some services are non-stop, and others stopping. **Every** train stops at York, for example, none pass through; some services terminate there, but most continue beyond York. (Some also originate there, of course.) The above treatment would apply over the individual sections.

We can thus schedule a regular interval stopping service, (or several, limited only by the 7 or 6 available streams,) interleaved with non-stop services, with absolutely no impact on line capacity. This sounds too good to be true, and, in a sense, it is. It is, literally, true, but the fundamental time unit is the capacity slot, not the elapsed hour. We thus get regular interval stopping service of one train per platform every 516, 684, 734 or 841 seconds (i.e. 8m36s, 11m24s, 12m14s or 14m1s) for line speeds of 225, 300, 360 or 400kph. If both platforms are in use there is a train every 222/296, 294/392, 366 or 420seconds (i.e. 3m42s/4m56s, 4m54s/6m32s, 6m6s or 7m exactly), the first two, being 7-slot streams, have alternately 3 and 4 slot intervals, the others, being 6-slot, having a regular 3-slot interval.

The only alternative approach is to decide what the station stop time, and thus the repeat interval, should be. For the line speed range of interest, only one such time is available – every 15 minutes (except for the very lowest speeds in the range, where it is every 10 minutes). The actual capacity slot value is inflated from TSD(e) such that so that 7 or 6 slots equals 15 (or 10) minutes. This of course applies throughout, to all services. Unsurprisingly, the capacity penalty is severe. The following chart summarises:

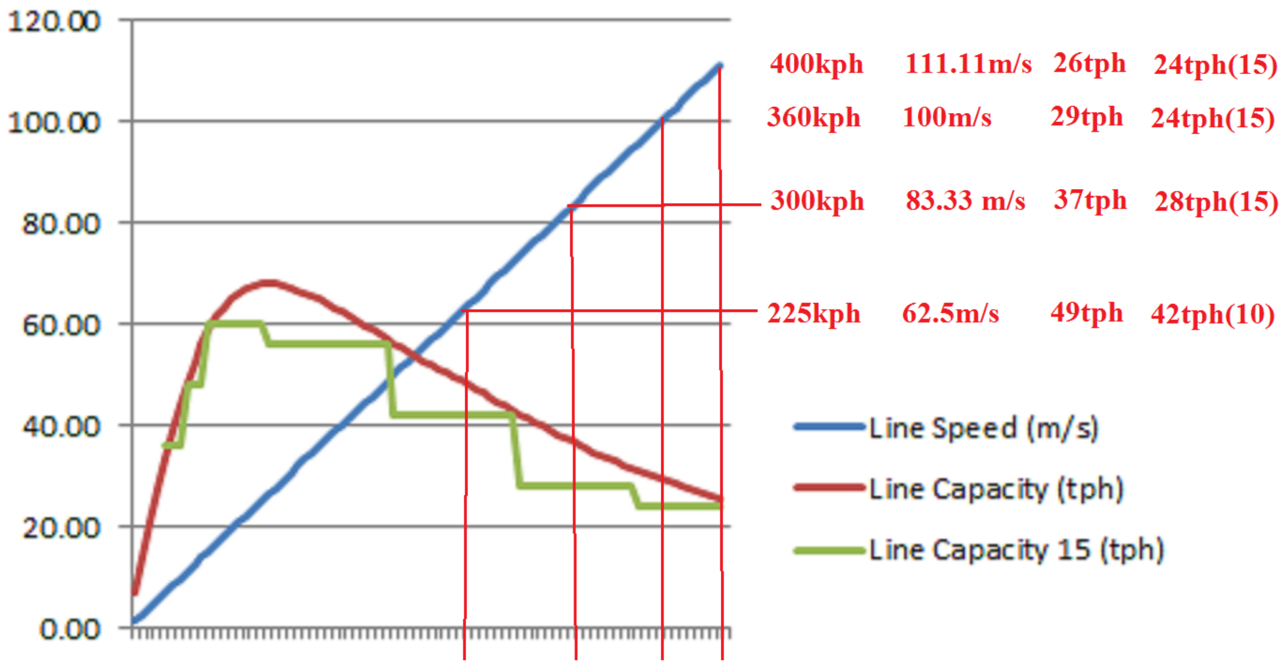

In the chart, the continuous line is the Extended Line Capacity, as discussed earlier. ‘Line Capacity 15’ is the available capacity when constrained to a 15 (or 10) minute frequency (over the line speed range of interest). As is clearly seen, the capacity penalty is severe. The four results of interest are quoted in full (so they don’t have to be estimated from the chart).

As I conclude the full treatment in ‘Same Speed Railways’:

‘All of this complexity is the consequence of the requirement to run a mixture of UHS non-stop and stopping services, with the requirement for the former to be able to overtake the latter. It has been a challenging matter to elucidate, and even more of a challenge to produce an intelligible explanation. I hope I have succeeded.’ I do indeed so hope.

Stations on the Main Line

Intermediate stations where **all** trains stop, (i.e. where all trains stop at all intermediate stations – I call this type of operation HS-Metro,) have a trivially simple effect on capacity: they impose a limit of 12tph per platform face, on the standard assumption of a 3 minute station stop + 2 minutes contingency. So the standard HS station of 2 island platforms, thus two platform faces in each direction, imposes an absolute limit of 24tph if all trains stop there. In most cases, this is hardly a restriction. If that isn’t enough, add one more platform face in each direction, getting an extra 12tph. 36tph is surely as much as anyone could reasonably need, (but you could just as easily have 48tph if you could make a case for it,) and is, almost certainly, in excess of the actual dynamic line capacity, as explained in the earlier pages. HS-Metro operation offers great simplification: no station loops are required (and therefore no high speed point work at stations) since there is no overtaking. It may be wondered if the capacity slot model is likewise redundant. In practice it is still needed, since it is necessary at route – as opposed to station loop – junctions, and all HS routes do in practice have route junctions, sometimes by branching within the same route, at others for connections between different routes – this is a **network**, after all, so of course they all connect up! The capacity slot model thus remains in use – if needed anywhere, it’s needed everywhere.

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 proposals 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 HS line of either type (overtaking or HS-Metro) 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 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 of the considered opinion that the 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. 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, Dover, Margate and Eastbourne, each of which, being served by only a fraction of the total, would need little if any new infrastructure.

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.