

A simplified graphical method for describing Limited-Slip Differentials in equilibrium conditions

Andrea Patrucco, Vehicle Dynamics Engineer, Applus+ IDIADA Spain, 29/8/2025

The views, thoughts and opinions expressed in this article solely those of the author, and they do not necessarily reflect the official policy or position of Applus+ IDIADA.

Introduction

Recently I stumbled upon multiple reviews of the new trio of Stellantis offerings: the Peugeot E-208GTi and its two derivatives, the Alfa Junior Veloce and the Lancia Ypsilon HF. These can be classified as sporty, electric B-segment vehicles, or, to use an anachronism, “hot hatches”. They come with a rather powerful powertrain, with 345Nm and 280HP delivered to the front wheels through, how interesting, a Torsen differential.

A Torsen differential coupled to an electric powertrain sounds rather unusual, in fact if you look it up in Wikipedia there’s an extensive list of vehicles which allegedly come equipped with one, and none of them is electric; Stellantis themselves are boasting it as a “segment first”.

Trying to get an idea of the effects of this solution in terms of vehicle dynamics, I found it rather difficult. There are plenty of videos and articles detailing the marvels of its kinematics (which I would not be able to explain), but none offer a simple answer to the question “Your powertrain is delivering X torque, tires are spinning at Y speed, the friction limit is Z: what happens?”

Hence, I took pen and paper to try it myself: and since I think the solution is rather neat, I’ll share it here for the joy of my 5 LinkedIn followers.

Differentials, Locking Differentials, Torsen

No article about locking differentials can exist without its “why do we need it” and “what happens if we lift a wheel” part: for those of you who know, just skip further ahead.

Whenever you need to deliver torque to two output shafts which must be able to turn at different speeds, you need a differential. Without going too far (with center differentials), a vehicle axle is a perfect example: during a corner, the inner wheel needs to turn slower and outer one faster: here comes the open differential, ubiquitous in road vehicles, which allows just that: at any speed, torque coming from the differential is evenly split to its two sides.

The most common way to explain the limitations of this setup is considering the case in which one of the two wheels is lifted, or experiences a sudden loss of grip: the resistant torque coming from its shaft reduces drastically, and, since torque must be evenly split, the other shaft is unable to transfer any either, leading to the inability, by the vehicle, to transfer its power to the ground.

In this situation, the use of a locking differential becomes clear: if we introduce some friction between the two shafts, once the grip-less wheel starts to spin, it will start transferring torque to the other, slower one through friction. Which is not only useful in this case, but has some positive effects in vehicle dynamics, too. For example, if we are trying to exit a corner fast, load reduction at the inner wheel (due to load transfer towards the outside) will limit the amount of torque we can deliver through

the tire, causing it to spin: with an open differential, this will also limit the amount of torque the outer wheel can transfer, reducing longitudinal acceleration. If we introduce friction between the two wheels, part of the tendency of the inner wheel to accelerate will be absorbed by the outer one, increasing our exit speed.

Some concepts to bear in mind when talking about all locking differentials are that:

- They always transfer torque from the faster wheel to the slower one
- The torque transfer is made through friction, i.e. with an energy loss. A nice way to visualize this is to think that the slower wheel is “borrowing torque” from the faster one. Since the output power is the sum of the multiplication of torque by speed at the two ends, torque moving from a higher-w to a lower-w side will imply a lower output power.
- Their ability to keep wheels locked is generally limited: that is, once the difference between resisting torque at the faster end and at the slower one exceeds a given amount, the faster wheel will start spinning, again limiting (to a higher value than an open differential) the slower wheel torque.

There exist many types of locking differentials, depending on cost and its intended usage. For off-road applications (wheel-lift kind of situation) there are manually locking differentials; for more performance-oriented applications there are limited-slip differentials (LSD) with clutch plates, and somewhere in between lies Torsen which, with its relative low cost, maintenance and noise, offers a more educated alternative to clutch plates, fit for powerful road vehicles who aim to deliver part of the advantages of an LSD without most of its shortcomings.

The Torsen name stands for Torque Sensitive, which to my understanding refers to its ability to bear different locking characteristics when the torque is “flowing” from the input shaft to the wheels (when accelerating) or the other way around (when coasting or braking). This is obtained through the differential geometry and its materials; normally, it is used to increase friction when accelerating and decrease it when braking.

The limits of its action are usually expressed through the *Torque Bias Ratio*, which is defined as the maximum ratio between the torque applied at the slower wheel, and the one at the faster one. A TBR of 5 means that the slower wheel will be delivered up to 5 times the torque of the fast one: ironically this implies that in the famous “lifted wheel” situation, the Torsen will be of little help, as multiplying a zero available torque at the fast side will result in, again, zero.

Differential basic equations

Suppose that we know the torque applied at the input shaft of the differential and the current angular speeds of the wheels, and that we want to determine the output torque at each output shaft.

In this condition, we can use the following system of equations:


$$\begin{cases} M_L + M_R = M_C \\ M_L - M_R = M_B(\Delta\omega_{L-R}, M_C) \end{cases}$$

That is:


- The sum of the two torques equals the input torque at the differential M_C

- The difference between left M_L and right M_R torques equals the locking torque M_B , which is a function of the speed difference between the two output shafts $\Delta\omega_{L-R}$, and the input torque.


By writing the M_B function in a sufficiently flexible way, one can describe the behavior of a large variety of differentials. A typical method employed, for example, in my suite of vehicle models, is to use the dependency on the sign of $\Delta\omega_{L-R}$ in torque differentials through a $\tanh$ function: this eliminates potential discontinuities (hampering time-integration) and the need to implement a rigorous locked/unlocked check on differential braking torque. It should be noted that this comes at the cost of a loss of fidelity in close-to-lock situations (since some however small speed difference will remain between the two wheels).



In the above plot, we see the behavior of a “Ferguson” type differential (locking torque depends on speed difference): M_B (here called C_f) depends on $\Delta\omega_{L-R}$ and there is no input torque-dependent component.



In this second example, we see the amplitude of M_B depending on input torque, for a ZF or Torsen type differential: the amplitude of the maximum locking torque is proportional to the input torque, with two different slopes depending on its sign. Note that different shapes are indeed possible (with different kinds of differential); by adding preloads to the differential components it is possible to obtain an M_B curve which is vertically offset with respect to zero as in the following example:



In building functions to describe different kinds of differentials, I took inspiration from the notes to [1].

More in general, we can express the *maximum locking torque* of the differential as:

$$M_B = a * M_C + b$$

Which is of particular interest in the Torsen differential case, where:

$$a = \frac{TBR - 1}{TBR + 1}; b = 0$$

Most sources report a TBR range, for Torsen differentials, between 2 and 6, corresponding to the following sensitivity values:

TBR	a
2	0.33
3	0.50
4	0.60
5	0.67
6	0.71

You can find a more rigorous formulation of this, as well as relation between TBR and the differential's efficiency, in section 3.14 of [2].

Even though useful for time-integration, the question we normally ask ourselves when thinking about a differential are “reversed” with respect to the form of these equations. That is, it is more interesting to think in terms of available resistant torque (or tire grip) and of potential lock, than of known input torque. The questions we want to answer are:

- For a given condition, will the differential lock?
- How much of the available power will we be able to transmit to the tires?

The proposed procedure will answer them, though at the cost of some painful simplifications.

Simplified Differential Solution - Assumptions

Let us focus on the *driving* part of the differential's range, that is, the drivetrain is transferring a positive amount of torque to the input shaft of the differential. Let us also focus on the axle case only: we will not consider uneven static torque distributions (which are possible in case of center differentials). We are studying an *equilibrium* condition, i.e., wheels have no ability to absorb part of

the transfer torque by increasing their speed. At the same time, an infinitesimal speed difference between the wheels is sufficient to allow for the full torque transfer between the two output shafts.

We will consider the following to be true:

- The longitudinal force generated by each tire is a linear function of its slip rate until its maximum value; then, it remains constant. As a consequence, the amount of resistant torque available at the i-th shaft is linear with speed, and plateaus when it reaches the maximum value:

$$\omega_i = \begin{cases} k_i(\omega_i - \omega_{i0}) \text{ for } \omega_i \leq \frac{M_{Ti}}{K_i} \\ M_{Ti} \text{ for } \omega_i > \frac{M_{Ti}}{K_i} \end{cases}$$

Where we called M_{Ti} the limit resisting torque for the i-th wheel (depending on tire friction), and ω_{i0} the angular speed of the i-th wheel in absence of slippage with respect to the ground.

- The torque available at the differential input shaft does not depend on the speed of the input shaft. This assumption would have been absurd a couple of years ago, but it is surprisingly reasonable in the era of EVs and of modern, turbocharged ICEs, which often manage to keep a constant maximum torque value along a wide range of engine speeds.

If there are ways around these two assumptions which do not dramatically increase the complexity of the solution, I didn't find them.

Simplified Differential Solution – Approach

The proposed method is based on a graphical representation of the main torques involved in a differential's status. We will first build a plot using the right side shaft torque as the X axis, and the left side shaft torque as the Y axis. This represents the space of all the possible solutions to the “differential problem”, identified by (M_L, M_R) tuples.

This is, sadly but expectedly, not my invention, as it is described as “common plot” in chapter 20 of [3] and used to discuss various limited slip differential characteristics, for example, in [4].

There are a few additional geometrical entities we can add to further restrict the search of possible solutions.

- The plot will be limited to the right side by the available friction of the right tire (yellow line): at no point the wheel, in equilibrium, will be able to exceed this amount of transmitted torque.
- The plot will be limited to the upper side by the available friction of the left tire (magenta), for the same reason.
- The plot will be limited, on both sides, by lines representing the limit M_B obtainable for a given combination of M_L, M_R . They can be found, knowing sensitivity and preload of the differential, through the equations:


$$M_L = \frac{1+a}{1-a} * M_R + \frac{b}{1-a}$$

$$M_R = \frac{1+a}{1-a} * M_L + \frac{b}{1-a}$$


- A further limitation is given by the maximum input torque available at the input shaft. This condition is limiting the sum of M_L and M_R , hence it will be described by:

$$M_R = M_{C,Max} - M_L$$

We will hence obtain the following plot, in which the area identifying impossible solutions has been grayed out.



Now that we know that torques will necessarily fall into the white area, we can go ahead with another geometrical construction.



In this plot, we have added a bisector in green. This line simply represents the equation:

$$M_L = M_R$$

That is, an equal torque distribution between the left and the right wheel, corresponding to the “open differential solution”. In a perfect, lossless open differential, possible solutions to the torque split are always found on this line. The line will cross:

- The left wheel available torque limit (magenta)
- The right wheel available torque limit (yellow)
- The maximum torque at the differential line (red)

Whichever intersection comes *first* (that is, at the lowest left or right torque coordinate) determines the solution. In our case, the left wheel has a lower friction limit, hence it reduces the applicable torque to the right one as well to 400Nm.

Now, considering the locking differential case and the assumptions we made (that is, wheels keep delivering their maximum torque value for whatever higher speed they reach, and the maximum input shaft torque remains constant with speed), we can further elaborate on the concept.

Tire friction limited case

If the first crossing involves a wheel friction limit, we can imagine that the torque-limited wheel will accelerate up to a point where it can run faster than the other one, leading to a torque transfer which is only limited by the maximum friction between the wheels. Graphically, this means moving along

the vertical or horizontal (in our case) line limiting torque, right or up, to increase torque on the other side (right, in our case).

At some point, we will reach one of the following:

- Maximum input torque limit: we are reaching the available torque (which is great news for longitudinal acceleration of our vehicle). The speed of the *slower* wheel, which is always the non-friction limited one, will be determined by dividing the identified torque on its side by its “stiffness” k_s and adding its zero-slip speed:

$$\omega_s = \frac{M_s}{k_s} + \omega_{s0}$$

The speed of the faster wheel can be determined as follows:

$$\omega_f = \max\left(\frac{M_f}{k_f} + \omega_{f0}, \omega_s\right)$$

That is:

- If the angular speed the wheel reaches at its maximum friction is already larger than the speed of the slower shaft, this is sufficient to ensure torque transfer to the slower side. This implies that the differential is not locked, and that the two shafts are generating torque transfer through friction between them, rotating at different speeds.
 - Otherwise, we have to increase the speed of the faster side to ensure it *remains the faster one*. Since the amount of torque delivered is constant for every slip (hence, speed) after it peaks, we can accelerate the wheel up until it reaches the speed of the slower shaft. The differential *locks*.
- The other wheel’s torque limit. This means that we are not able to transfer the whole torque generated at the input shaft of the differential, but we are fully exploiting the available friction. Torque at both sides is determined by the tire limit; speeds are not really determined (since, due to the shape of the speed-torque curves, any slower-side speed greater or equal to $M_{Ts}/k_s + \omega_{s0}$ technically withstands the constraints, and the faster one is always able to keep up), so we can assume a “best case” and just impose

$$\omega_s = \frac{M_{Ts}}{k_s} + \omega_{s0}$$


Following the same procedure as in the previous case to determine the speed of the faster side.

- The locking torque limit condition (blue line). In this case, we are not able to deliver the full amount of torque available at the input to the output shafts due to the differential’s limited ability to transfer torque from the faster to the slower wheel. The speeds of the wheels are calculated as in the first case (input-limited).

One may note that, for the limit the case of an open differential, the third condition would always be reached: this may bring to situations where we would flag the differential as “locked”, even though the differential can’t lock. Let’s consider, for example, the case of an axle with identical zero-slip values for wheel speeds (for example, in a still vehicle). Let’s assume the rather odd case that the left wheel has a very low friction limit, but a very high *stiffness*, while the opposite is true for the right wheel: high friction limit, low stiffness.

Limit torque is reached for $M_L = M_R = M_{TL}$, however the right wheel speed must turn faster than the left one even to reach the lower, left tire grip limit. The maximum torque is reached for a determined right angular speed, and for whatever left speed great enough to ensure limit friction is reached.

Following the proposed algorithm, we will find that $\omega_L = \omega_R$, which apparently contradicts the open differential philosophy. However, if we introduce the slightest amount of resistant torque (for example, due to the differential friction) and wait long enough, the left angular speed should indeed reach the right's, since it will be delivered *slightly more friction* than it can withstand, until it reaches the faster wheel speed.



In the following four graphs, we can see:


- Upper left: torque limited by the left wheel torque, then by the input shaft.
- Upper right: torque limited by left and then right wheel torque.
- Lower left: torque limited by left wheel, then by the locking torque
- Lower right: torque limited by the right wheel, then by the input shaft torque.

The black x highlights the solution found by the algorithm.

Input shaft torque limited cases

Up until now, we have assumed that the first limit we would encounter while going upwards and towards the right with a free differential would be resistant torque at one of the output shafts. But what happens when, instead, it is the torque at the input shaft limiting the output values? This is

indeed the case in the majority of boring, normal driving conditions, where we input a torque well below the one that would make the tires spin.



Let's consider the case here above. Tire friction is well above the input shaft torque. Thanks to the known limits of the differential (blue lines), we know that the operating conditions of the axle will be found in the triangle identified by the red-dotted intersection points between the input shaft torque line and the blue ones, and the origin. Since we are able to fully transmit engine torque to the ground, we can assume our solution will be found on the red line, but at which height?

To determine this, we'll first try and find the hypothetical solution under the assumption that the differential is, indeed, locking. We can hence simplify the expression for input shaft torque, that is:

$$M_C = k_L * (\omega_L - \omega_{L0}) + k_R * (\omega_R - \omega_{R0})$$

Into the following:

$$M_C = \omega * (k_L + k_R) - \omega_{L0}k_L - \omega_{R0}k_R$$

Where we replaced the two separate angular speed expressions with a single one, reflecting the fact that both wheels will move at the same speed when the differential is locked. Solving for ω we obtain:


$$\omega = \frac{M_C + \omega_{L0}k_L + \omega_{R0}k_R}{k_L + k_R}$$

We can further use the angular speed we found to calculate the single torques acting at the two differential ends.

$$M_L = k_L * (\omega - \omega_{L0})$$

$$M_R = k_R * (\omega - \omega_{R0})$$

Hence, the torques at the two sides will depend on the tires' stiffness and on their initial velocity. The graph above shows the example of a vehicle, with a track of 1.5m and rolling radii of 0.33m, performing a $\sim 5\text{m/s}^2$ corner at 100km/h. The two zero-slip wheel speeds are 83.77 and 84.58rad/s, respectively, a rather low difference considering k values (identical for the two sides) of 300Nm*s/rad. But what happens when zero-slip wheel speeds are more spread apart? Let's reduce speed to 50km/h leaving all other parameters constant (including lateral acceleration): the vehicle is now driving on a much lower (1/4) radius corner, with wheel speeds of 41.27 and 42.91rad/s, respectively. Following the same approach, we find the situation represented in the below graph:



The red X identifies the locked-differential solution. Since ω_{L0} and ω_{R0} are more spread apart, the corresponding torque difference is larger, and the inner wheel now carries an amount of torque which is well past the operating limits of the differential. In this case, we must “go back” along the red line to the first valid operating point, that is, at the crossing between the input torque limit and the locking torque limit.


We can then recalculate speeds with the usual formula:

$$\omega_i = \frac{M_i}{k_i} + \omega_{i0}$$

In this condition, the differential is still able to transfer torque from the faster to the slower wheel; however, it is doing it in a non-locked fashion, that is, the two wheels are rotating at different speeds, with a friction-driven torque transfer between from the faster to the slower one.


In this example, the inner wheel receives torque from the outer one, which is turning faster, and ends up slipping more than it would in the open differential case, leading to an uneven force delivery generating a moment opposing the vehicle's rotation. Which, by the way, is not enough to justify understeer effects, as the lower longitudinal slippage requested to the outer wheel might lead to an increase in lateral force, which, in a front wheel drive vehicle, may more than compensate the aforementioned torque.

With stiffer tires, the amount of locking torque required to keep speeds equal increases, leading to an increased likelihood of exceeding the locking torque limitations. And in case of different stiffnesses at the two vehicle sides, we could find ourselves in the situation below:




Where the inner wheel *would turn faster* than the outer one due to its low stiffness, and hence, torque transfer happens in the opposite direction, inner to outer, creating a yaw torque in the direction of the corner. This, and the case shown a couple of pages ago when dealing with the “left wheel torque limited” case, explain why cars with a locking differential tend to show less limit understeer (FWD) or more oversteer (RWD) in presence of large lateral load transfers.

If, while moving along the red line, we cross one of the tire's limit friction line *before* the differential locking one, the solution is similar.




In this case, our ability to lock the differential at a given speed is limited by the left tire's ability to absorb torque, and its further acceleration does not lead to any torque benefit to the right side. More than a *friction limitation to deliverable power*, this constitutes a *friction limitation to torque imbalance*: we are still able to completely use the input shaft torque; however, the amount of moment opposing rotation is limited by friction.

Which is also interesting when considering the case of a progressive acceleration out of a corner: we are slowly moving the red line further away from zero by applying an increasing amount of engine torque; with an open differential, at some point the left wheel will start to spin, limiting torque at both ends; with a locked differential, we will first encounter the “torque imbalance” limit: the inner wheel will start spinning *before* with respect to the open differential; then, when we reach the $M_L = M_R$ point, the differential will lock to allow the outer tire to keep generating force, up until either the right tire limit friction will be reached, or the right locking limit. The “trajectory” of M_R , M_L in that condition is shown in the below graph.



Note that this whole methodology is not necessarily limited to Torsen-type differentials: it can be used to determine the status for every combination of preload and sensitivity, such as in the following case:



For those that are less scared than me by non-linearities, I should mention that section 3.14.13.2 of [2] shows that it is indeed possible to find a still reasonably simple solution to determining a differential's state integrating it with a proper tire model. It starts, however, to get a bit further from the "back of the envelope calculation" nature of the method presented here.

Conclusions

All things considered, do I think that fitting a Torsen differential onto a B-segment EV is a good idea? From a VD standpoint, yes. For a rather powerful (280HP, 345Nm) powertrain coupled to front wheel drive, advantages in terms of traction and of understeer reduction should be significant. It probably manages to give the vehicles (which are based on a general-purpose and not necessarily fun-oriented platform) a very distinctive character at the cost of some added complexity and cost, which I think the “usual” modifications (static geometry, spring/ARB stiffness, ride height) could not markedly achieve. Furthermore, it brings back to this class of vehicles some “interestingness”: it’s a product marketed to the kind of enthusiast who is likely to appreciate having some unique mechanical feature producing peculiarities he can adapt to, and talk about.

My only doubt would be about power efficiency: for a vehicle whose range characteristics are possibly as important as its road manners, a locking differential introduces a further source of power loss; however, it is not completely clear (to me) whether the weight of these additional losses is significant under normal driving conditions – which may be an interesting direction for further research.

Some other ones would be, understanding what happens with center locking differentials with an uneven torque split between front and rear axle, and how to solve a system with three differentials, with limit available torque at front and rear depending on the separate front and rear differential solutions: this is something I’m currently working on (not happy about the *additional* simplifications I have to apply in this case).

In the unlikely case it’s not immediately proven totally wrong, I hope my approach will constitute for somebody else the source of pleasant, mathematical time waste it is for me.

Bibliography

- [1] L. Galante, "Lezione sui differenziali," *Experis MotorSport Academy*, 2016.
- [2] M. Guiggiani, *The Science of Vehicle Dynamics - Second Edition*, Springer, 2018.
- [3] D. L. M. William F. Milliken, *Race Car Vehicle Dynamics*, SAE International, 1995.
- [4] M. Gadola, D. Chindamo and B. Lenzo, "On the passive limited slip differential for high performance vehicle applications," *International Symposium on Advanced Vehicle Control, AVEC 2018, Beijing, China*, 2018.