Leading-Edge Motorsport Technology Since 1990

Racecar engineering™

Reprint • www.racecar-engineering.com

Seven-post rigs

Shaking down the best in suspension testing tech

All **shook** up

If on-track running is severely restricted then there are always test rigs to help evaluate and develop your suspension. We spoke to the movers and shakers in this sector for an insight into the technology that's rocking their world

By GEMMA HATTON

If you're not moving forward in motorsport, then you are standing still – the cliche that underpins every racing team across the globe. That burning hunger to claim the top step on the podium is what drives teams to continuously develop, improve and optimise the performance of their racecars and drivers. Unfortunately for them, the governing bodies have a different role to play and roadblock the teams at every turn through regulations. Arguably, the most challenging of these roadblocks is the heavy restrictions on track testing, which is often implemented to try and establish some control of costs.

With no better platform to test a racecar than a track, the teams react by spending even more money on developing driver-in-the-loop simulators and other simulation tools, but also test rigs where the real car can be evaluated. 'The use of seven-post rigs depends on the culture of the team. For us, we believe it is extremely important to test the real vehicle as

Renault is a great believer in making use of shaker rigs and it says that in the past it has developed technology like its F1 mass damper with them

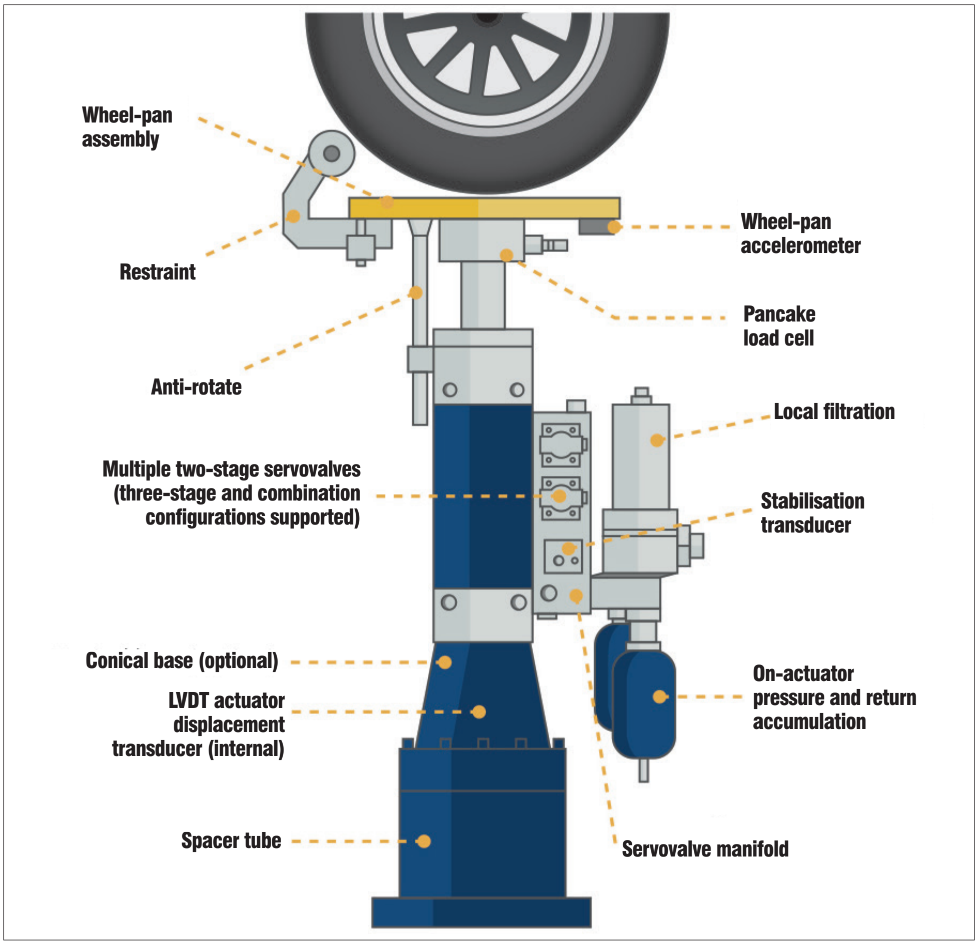

Displacements are driven into the car's wheels through the wheel pan actuators. These can switch between two-stage and three-stage servovalves depending on the required power of the test, or on the application of the vehicle that is being tested

'We believe it is extremely important to test the real vehicle as much as we can, to help optimise the car when at the race track'

much as we can to help optimise the car when at the track,' explains Rene Torcato, lead R&D engineer at Renault Sport F1. 'This has been proved by past championships; mass dampers were developed here using these rigs, so we don't see our future without them.'

Shake down

The primary purpose of a seven-post rig is to not only understand the suspension design, but also optimise it for each track. This is achieved in two ways. Firstly, frequency sine sweep tests are conducted to characterise the vehicle. This is where the input signal starts at high amplitude and very low frequency and, with time, the frequency increases, while the amplitude decreases. From this engineers can understand how the car responds to that frequency range.

'We also use our seven-post rig to replay the profile of different tracks to see how to set up the racecar and what compromises to make between grip and aero levels, depending on the demands of each track. For example, Monza is more aero dependent, while Monaco is more grip dependent,' Torcato says. 'The advantage of this is we can make a set-up change and immediately evaluate this change on the first five tracks of the season, for example. This allows us to see how that change affects the car's response to specific corners under specific conditions at each track independently.'

Good vibrations

As well as optimising overall suspension performance, seven-post rigs enable teams to test specific components in isolation or investigate particular issues in a controlled environment. 'Some teams want to analyse specific scenarios such as hitting kerbs, which can cause the car to unsettle, so we have a kerb file that teams can use,' says Christer Loow, engineering manager at Ohlins. 'When you actually stand next to the rig, the movement of the suspension can be extremely violent, but that's what happens in reality. Sometimes it's quite useful for teams to see close-up how the suspension components move because they may realise that they need to make specific parts stronger, which is something they would never see at the track, until it broke.'

Of course, there are some limitations with rig testing as not every condition can be modelled; variables such as wind direction, tyre wear and driver behaviour can't be incorporated into these rig tests. However, interestingly, this allows the suspension behaviour to be isolated and tested within a very controlled environment, so only the desired variable is changed, while everything else remains constant, which is almost impossible to achieve at a race track.

'For new tracks, we can actually build up a track using real data,' Torcato says. 'So we start with a bit of white noise and then from the response of the vehicle to the known inputs we can do some mathematics so that the rig understands what input is required to

A rig allows the suspension behaviour to be isolated and tested within a very controlled environment

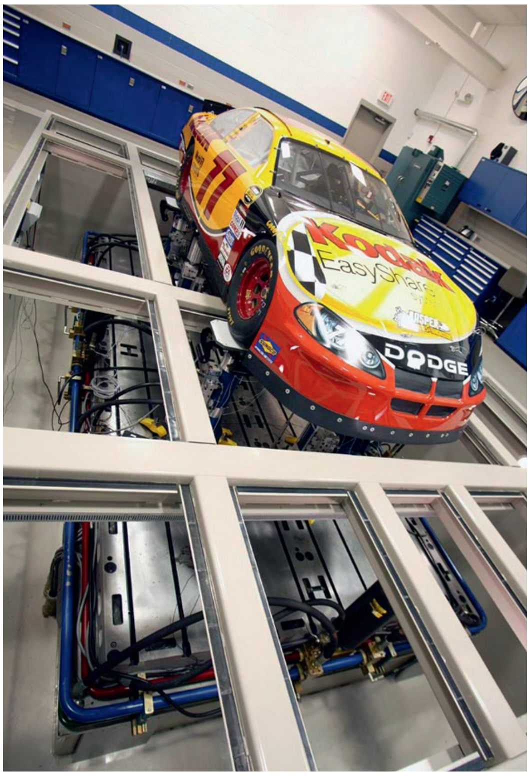

For oval track racers such as NASCAR stock cars the triangular configuration of the aeroloader actuators are often placed with two on one side and then one on the other, rather than one at the front and two at the back. This is to help replicate the lateral forces

have the output that matches the real data. It requires several iterations and the software will automatically start modifying the input until we have the output from the sensors on the rig reading similar to the car on track. We try to automate our processes as much as possible, but it's still important to allow some flexibility if we need to change anything.'

There are a range of static and dynamic test rig types teams use to optimise their set-up. Four-post rigs are arguably the most common, and this is where four servo-hydraulic actuators are used to drive precise displacements into each wheel of a racecar via an instrumented wheel-pan. These displacements are determined by a test track specific drive file which has been derived using iterative techniques from track data, including hub accelerations, pushrod loads and damper displacements.

Shaken not stirred

As useful as four-post rigs are, they have their limitations, chiefly because at over 100mph the suspension has to cope with an entirely new input: aerodynamic downforce. To simulate this downforce, additional aeroloader actuators are used, most commonly three but sometimes four, converting the original four-post rig into a seven- or eight-post rig. These extra aeroloaders are programmed by an aeroloader map within the control system and need to attach to a rigid area of the vehicle's chassis, which is why they are usually mounted in a triangle configuration, with one at the front (under the nosecone) and two at the rear for single seaters. However, for oval track cars like NASCAR, two of the actuators will be on one side, with the third aeroloader on the other side. This triangular configuration allows the effect of weight transfer under acceleration and braking as well as pitch and roll to be more accurately simulated.

The aim of the aeroloaders is to pull down on the entire vehicle chassis, to simulate the overall effect of downforce. The challenge is, that these aeroloader actuators essentially act as viscous dampers, which, unless mitigated, can disrupt the natural body motion of the racecar as it responds to the controlled inputs at the wheel-pans. The loop-gain of the aeroloader load control cannot generally be increased sufficiently enough to minimise this damping to an acceptable level without the risk of instability. Furthermore, the mechanical impedance of the car at the attachment points is also included in the load control loop and so is heavily affected by any set-up changes.

To mitigate these effects, one solution is to include some form of mechanical compliance, such as a spring arrangement, within the aeroloader actuator load path. Here the actuator is coupled to the vehicle through a 'compliant link', which is essentially a clever set of springs that isolates the motion. Therefore, load can be applied to the vehicle, while retaining sufficient compliance so that the racecar is still able to move, without having to move the actuator. Unfortunately, when used in isolation this compliance means that the large downforce loads experienced during braking cannot be simulated and, depending on the configuration, the mass of the moving parts may increase the overall mass of the racecar as well.

Linked-up thinking

Another strategy is to try to predict the velocity of the test vehicle chassis and to control the actuators so that they follow the car body's motion. However, these velocities would change with any set-up modification and the measured velocity would tend to de-stabilise the aeroloader load control loops.

'If you only use a compliant link that can give you one set of issues, while using velocity feed-forward alone can lead to stability challenges,' says Frank Blows, chief engineer at Servotest. 'However, by combining the two and using a compliant link together with a velocity feed-forward algorithm, then you achieve a workable solution to simulate the effect of downforce on a seven-post rig.'

This design not only removes the de-stabilising effect of the feed-forward signal, but the impedance is no longer included in the aeroloader control loop, making it independent of set-up changes and the mass of the moving parts is driven by the aeroloader itself.

To translate the electrical demand signal from the control system into the movement of hydraulic fluid in the actuator, two stage servovalves are used. These are low flow devices at around 40l/min and relatively low power, up to 20kW. 'The beauty of this type of servovalve is its fast response to command signals, they are even capable of operation at over 200Hz,' says Martin Jones, motorsport market manager at Moog. 'They are similar in principle to the much

The main tests conducted on seven-post rigs are sine sweeps which help to characterise the suspension, and track replays where constructors or teams can tune the suspension set-up to optimise performance before the car actually hits the track

Two-stage electro-hydraulic servovalve

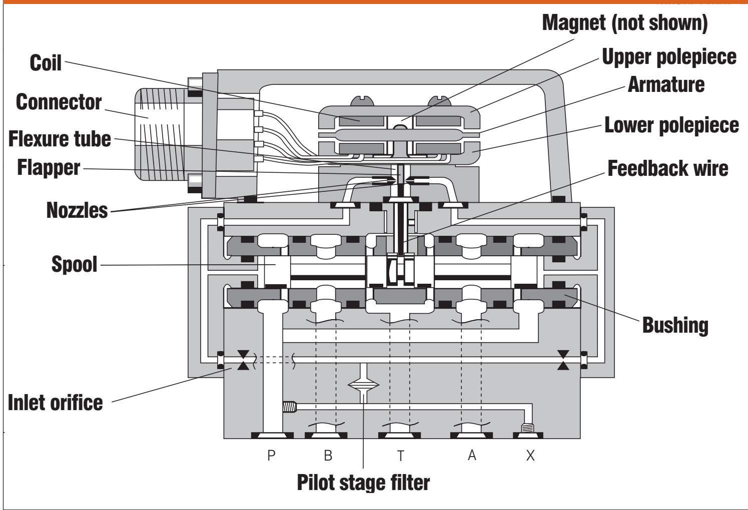

Three-stage servovalve D792 with pilot valve D765 series

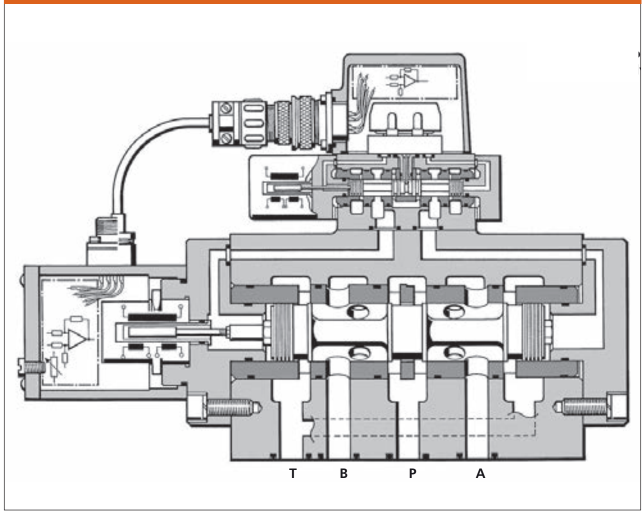

Cutaways showing a Moog two-stage servovalve (left) compared to a three-stage servovalve. These components play a crucial role in operating the actuators on seven-post rigs

smaller E024 valves used in F1 brake by wire systems [see p64]. Essentially the servovalve is an electronically commanded spool valve that can be applied to provide fast and accurate control of force or position through a double-acting cylinder. The power limit inherent in this type of valve can be overcome by using two or three units in parallel to supply enough hydraulic power for fast vehicle movements while retaining the high resolution required to follow complex waveforms.'

Rough and ready

For other types of race vehicles such as rally cars and other off-road cars, higher-performance rigs may be required. These typically offer longer-strokes and higher wheel-pan velocities. While two-stage servovalves are still required for some of the testing on these rigs, the high velocity tests may require an alternative approach. For such rigs high-flow three-stage servovalves can be added, typically with hydraulic isolation slices when being operated in low flow mode.

A three-stage servovalve is essentially a two-stage servovalve sitting on top of a larger high-flow valve. The two-stage valve acts as a 'pilot' to a much larger and higher-flow third-stage spool, which has the capability to shift substantially more oil into the wheel-pan actuators.

For aeroloaders, neither conventional two-stage or three-stage valves provide sufficient performance. Therefore, these actuators use dual electronic feed-back valves (EFBs) which deliver the fast response required by the velocity feed-forward control to accurately track the movement of the vehicle's chassis.

'In normal mode, teams would operate the rig with a couple of two-stage valves. But if a customer was testing a rally application, then not only do they need actuators with a bigger stroke, but they would also need to enable the three-stage servovalve to achieve the higher wheel velocities and drive the rig,' explains Peter Rogers, marketing manager at Servotest. 'The servovalve configuration can be switched automatically to suit the application, with two-stage valves delivering the control required for conventional seven-post tests while retaining the performance advantage of a three-stage valve for those required cases. It would be like trying to go shopping in a Formula 1 car; you would get there quicker, but you wouldn't be able to control it enough to park it. Similarly, teams mostly utilise two-stage servovalves because they are controllable. However, when they really want to go racing, they use the three-stage servovalve.'

Rigged for efficiency

As with any simulation tool, efficiency is key to achieving realistic results with rigs. With seven-post rigs the fidelity of reproducing time domain signals is critical to providing a realistic physical simulation. One key to this is minimising the effects of friction within both the wheel-pan and aeroloader actuators. This can be done with hydrostatic bearings at the rod end as well as the base end of the actuators. Hydrostatic bearings use two pairs of opposing rectangular pockets of high-pressure oil within each bearing head to both centre the piston rod and lubricate the bearing surfaces, eliminating the friction and wear that can occur with less sophisticated polymer-bearings.

Servotest also utilises a seal-less piston design on its actuators, which relies on tight tolerances between the piston and the cylinder bore to minimise leakage and eliminate the motion distortion due to the static friction (or stiction) that is characteristic of sealed-piston actuators at each end of their stroke. Also, the surface of each wheel pan has been made to be extremely low friction to deliver the highest possible quality of motion simulation.

Another, often overlooked, practical aspect of seven-post rig design is that test rig operating pressure and associated actuator force rating (and therefore piston area) can be matched to the test vehicle characteristics and the required rig performance. For Servotest, this allows the size of the hydraulic power supply to be optimised with an associated reduction in long-term running costs, something that just isn't possible for suppliers with a discrete range of actuator sizes and a fixed operating pressure.

Quake up call

It is not just the individual components that require careful design consideration, the overall rig does too. The pit layout can vary depending on the size available and the rigs are often mounted on a heavy, often concrete, seismic base. This ensures that vibrations from the rig are not transmitted to the rest of the building, which is essential if there are manufacturing or metrology offices nearby, but it also avoids any vibration from adjacent facilities affecting rig performance. This seismic block can either be 'in-ground' where it is connected solidly to the ground, resting on rock or sand, or 'floating' where it is suspended on metal spring-boxes or air springs, although the latter is more expensive with high maintenance costs.

Typically, seven-post rigs are built into pits with the active components housed below ground level. This means that the racecars can be wheeled directly on to the

With seven-post rigs the fidelity of reproducing time domain signals is critical to providing a realistic physical simulation

wheel-pans without the need for lifting equipment. Above-ground installations are also possible, but they are less common.

In-ground pit designs vary hugely depending on the company and location. 'At one extreme there are relatively small and shallow pits that are accessed via a trap-door,' says Rogers. 'Typically, these pits can only be accessed safely when the seven-post rig is de-energised. At the other extreme there are large deep pits that are complete with an inspection walkway channel around the bed-plate, typically accessed by a separate staircase. These pits allow commercial clients, sponsors and race engineers relatively safe viewing access to the underside of the vehicle while testing is in-progress. In these cases the hydraulic hard-line – the pressurised oil distribution system – is typically mounted either on to the side of the seismic block on the inside edge of the walkway, or on top of the seismic block itself, allowing unobstructed access around the rig.'

Heavy, often concrete, seismic bases ensure that vibrations from the rig are not transmitted to the rest of the building

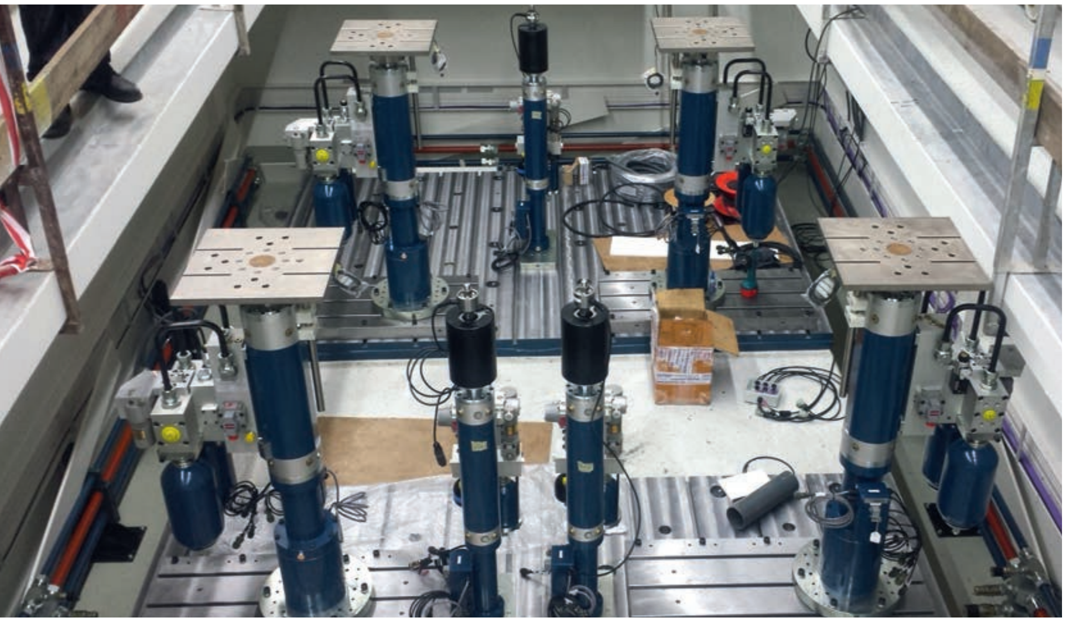

The seven-post rig at TRE, which is equipped with Servotest running gear. You can see in this picture the four wheel-pan actuators as well as the three aeroloader actuators, complete with the compliant link at the top (the black cylinder section)

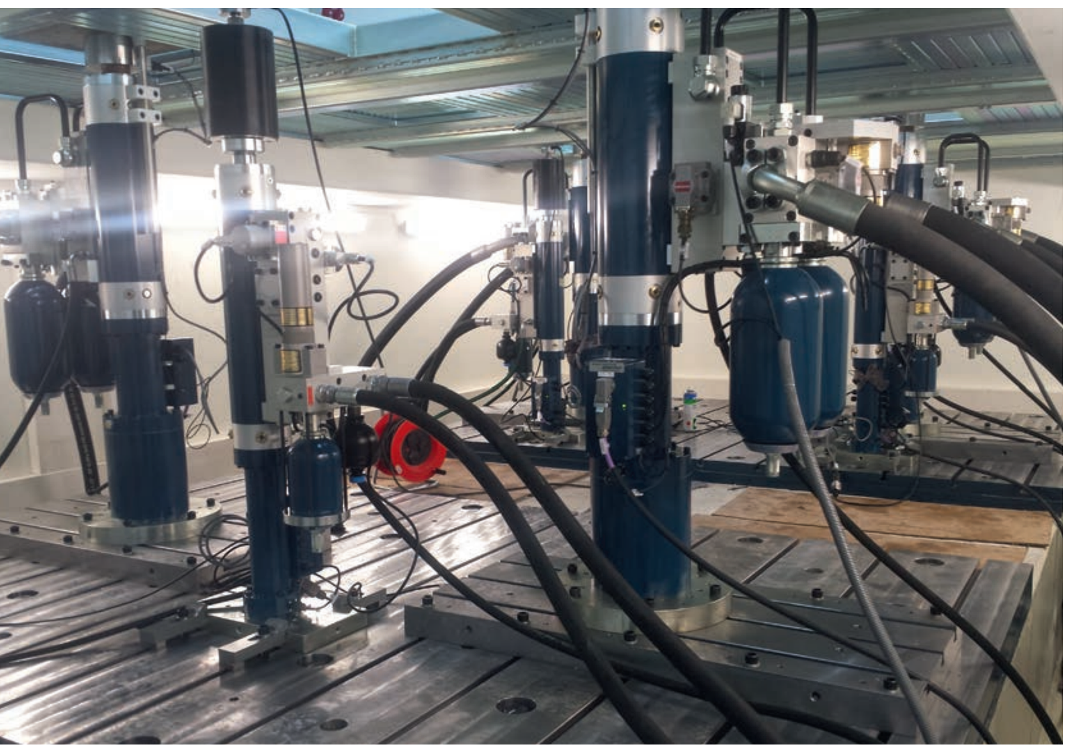

Wheel-pan actuators can be moved laterally and longitudinally on base plates while they also rotate on their own base discs

Cooler shaker

In hotter countries, often the pit area is also fully air-conditioned and climate-controlled. This not only makes the pit environment more comfortable for the rig engineers, but also provides a uniform temperature environment for the test vehicle, suspension components and measurement instrumentation. In some cases this has been reported to have a noticeable impact on test result consistency.

The location of the actuators themselves also need to be adjustable to cater for cars with different wheelbases and track widths. This can be done with four simple bought-in T-slot bed-plates which are bolted together to form a large uniform space. The wheel-pan and aeroloader actuators can be re-positioned in both the lateral and longitudinal directions and held in place by machine clamps. Alternatively, large industrial magnetic bases under each actuator can be locked into place and then unlocked electrically, in which case spring machine ball-lifters, mounted in the actuator base, facilitate the repositioning of the rig components without the need for external lifting equipment.

As ever with any type of simulation, correlation is the biggest challenge. 'Each of our simulation tools need to complement each other as some things are easier to build and run on the real car, while for other things it's easier to create a model or test on the simulator,' says Torcato. 'For example, when we have quite different car designs from one season to the next we use simulations to help us understand the compromises at each track. But each tool has its limitations and it's important to understand what they are and how to compensate for any error that you might have. The aeroloaders, for instance, can result in extra damping on the racecar because they are not perfect in generating the procedural loads, so you need to understand those limitations.'

Testing data

As well as inputting track data into the rig for correlation studies, it's essential to use data from the rig in the set-up at the track. 'It's important for teams to take the data they learn from the rig and go back to the same track and physically test,' Loow says. 'One team understood the importance of this verification and physically tested changes at the track that not only improved the car's performance, but also those changes that worsened the performance as well, to really understand the correlation between simulation and reality.'

Simulation tools like rigs will only grow in importance in racing. Because these days race teams often need to have a racecar that is fast before they even roll it out of the garage.

Rigged-out

A seven-post rig test on a GT car such as this Porsche can log up to 70 signals at 1000Hz. That amount of data requires clever control strategies and accurate sensor measurement

Our two-part examination into shaker rigs continues with a look at the wide variety of hi-tech virtual and physical paraphernalia that's needed to make these amazing machines perform at their very best

By GEMMA HATTON

In last month's issue we delved into the fascinating world of shaker rigs, exploring the concepts behind them and looking at new innovations in the sector. This month we will be taking this a step further, venturing into the virtual dimension of rig testing, while also looking at the various measurement tools.

There are many questions to ask. With a F1 seven-post rig test logging approximately 100 signals at 1000HZ, just how is this amount of data measured, stored and managed? How can you synchronise seven actuators to move at mm-perfect displacements? How can you simulate tyre squash when the wheels are stationary? As ever, there is only one way to find out, and that is by talking to the experts.

At first glance seven-post rigs can be a bit of an anti-climax, as what you see is a car simply sitting on four metal plates that move up and down. But strip back the floor and underneath you will find a web of data acquisition points, along with complex hydraulic actuators to achieve the precise synchronisation and displacement of the wheelpans and aeroloaders.

Above 100mph, aerodynamic downforce has a major effect on the suspension. This is why in most motorsport applications three aeroloader actuators are added, with usually one at the front and two at the rear, along with the four wheelpan actuators at each corner. The wheelpan actuators drive displacement inputs into the wheels, emulating the bumps on a track, while the aeroloaders are attached to the chassis and simulate the overall effect of downforce by pulling down on the car. All these combine to create the seven 'posts' of the rig.

Double data

Rigs require two sets of sensors to capture the data. Firstly, there is a vast array of sensors on the car to accurately measure the reaction of the chassis, suspension and tyres, while another set

At each corner wheelpan actuators drive displacements into the wheels to simulate track inputs. There are also the three aeroloader actuators which emulate the effect of downforce; usually with one at the front and two at the rear, as seen here

A Formula 1 seven-post rig test logs approximately 100 signals at 1000HZ, but just how is this amount of data measured, stored and managed?

of sensors is required on the rig to monitor the inputs and the rig's performance.

'On the car usually there will be pushrod loads, damper displacements, strain gauges measuring the drop-links for the rollbars and accelerometers on the wheel hubs and the body to measure pitch, heave and roll,' says Henri Kowalczyk, COO at Auto Research Centre. 'On the rig side, each wheelpan has a displacement sensor and accelerometer to know what the 'road' is doing as well as a load cell to measure the normal load that the car is seeing. For a seven-post rig, you then have the aeroloaders to worry about, which require displacement and velocity sensors in addition to load cells. You will also be measuring temperatures of the actuators and other rig performance parameters. In this way, you can tell from the input side what the rig is doing and from the output side what the car is doing.'

Some of the most crucial sensors are the load cells, used to measure the axial forces exerted by the actuators on both the corner posts and the aeroloaders. The wheelpan load cells mounted on the corner posts also feature inbuilt accelerometers, so that the dynamics of the displacement can be monitored.

Load cells

'The principles of a load cell design are quite generic although there are different types depending on application specific factors,' says Ian Papworth, applications engineer at Novatech Measurements Ltd. 'We use foil strain gauges that are either etched or deposited depending on the production process of the supplier. Effectively, a strain gauge consists of loops of very fine wire, only microns thick. These are attached to a metal structure, and when this structure deforms under load, it either stretches or compresses the wire grid, therefore changing its length and cross-sectional area. All the time this deflection is occurring, the resistance of the wire is changing. The foil strain gauges are wired up in a Wheatstone bridge and any change in resistance excites this bridge, forming a potential divide. The output signal is in the order of millivolts and is linearly proportional to the

Sensor table

Car	Rig	
	4 post Rig	7 Post Rig (additional channels)
4 x Hub accelerometers (FL, FR, RL, RR)	4 x Wheel actuator displacement (FL, FR, RL, RR)	3 x Chassis velocity transducer (near each aeroloader)
4 x Chassis accelerometers (FL, FR, RL, RR)	4 x Wheel contact load (FL, FR, RL, RR)	3 x Aeroloader velocity
4 x Damper displacements (FL, FR, RL, RR)	4 x Wheel actuator acceleration (FL, FR, RL, RR)	3 x Aeroloader displacement
2 x Damper temperature (FR, RR)		3 x Aeroloader load
Vertical body acceleration (positioned at cog)		Front ride height
Lateral body acceleration (positioned at cog)		Rear ride height
Longitudinal body acceleration (positioned at cog)		Air speed (from track data)
		Lateral acceleration (from track data)
		Longitudinal acceleration (from track data)

There are two sets of sensors required to run a rig test; those on the racecar and those on the rig itself. Additional channels are then required for seven-post tests and track replays

amount of load applied. Instrumentation is set up so that the calibrated output is displayed as an accurate force measurement.'

Wheatstone bridge

The Wheatstone bridge is a concept dating back as far as 1833. It converts small changes in resistance to a voltage signal. Typically, four resistive elements are connected in a diamond shape, with the resistance value of one element unknown and the other three known. From the output, this unknown variable can be determined through comparison with the other resistors. In the case of a load cell, the Wheatstone bridge is used in conjunction with the strain gauges described above, which replace all of the resistive elements and all have equal, known resistances. When a load is applied to the strain gauged structure and the wire deforms, either stretching or compressing, the resistance across the Wheatstone bridge is unbalanced, producing a voltage corresponding to the induced strain and therefore the applied load. The wire of these gauges are designed to work within their elastic limit, so the wire reverts back to its original shape once the load has been removed, rebalancing the Wheatstone bridge into the 'zero' condition.

'The trick when designing a load cell is making sure the way you design and fabricate the metal component is going to give you the correct amount of deflection and therefore output signal for the load,' Papworth says. 'If you get it wrong, you could end up with the wrong amount of deflection. Either the output signal will be too low, leading to a loss in resolution, or you can go the other way and make the metal component too soft so that it deflects too much. This is even worse as it potentially puts the material under huge amounts of stress leading to fatigue, cracking and even complete mechanical failure. Although simple in concept, load cells do require a certain amount of engineering know-how and expertise.'

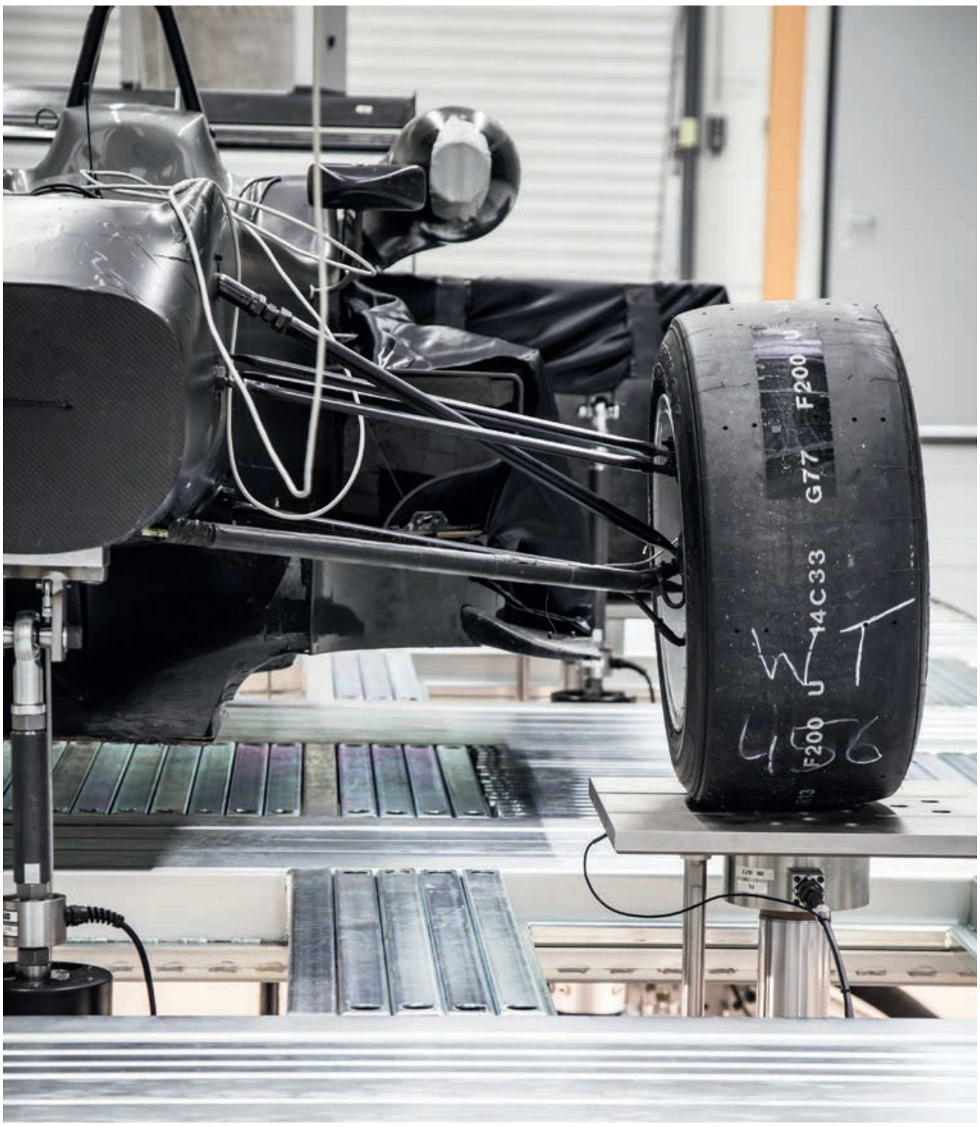

There are normally two sorts of test conducted on seven-post rigs. These are sine sweeps and track replays. The former characterises the suspension at different frequencies and the latter analyses the vehicle's response to realistic track inputs

Foil gauges

Novatech use foil strain gauges for four- and seven-post rig applications because they are the most stable and versatile, although the needs of many other industries to measure low forces at the highest resolution has opened up the potential for other load cell technologies to be utilised in the future. 'Piezoelectric devices are good for high frequency measurements and semiconductor gauges allow much more signal for the same mechanical deflection,' Papworth say. 'This could potentially increase the resolution beyond what foil strain gauges are capable of. But the technology that some companies are currently taking more seriously is fibre optics. They work almost like a radar gun where you measure the time taken for one pulse to travel down the fibre optic and bounce back. The strain element of the fibre optic wire can be attached to a material and, when deflected,

Four- and seven-post rigs tend to be adaptable and can be utilised to tune suspensions for pretty much any vehicle; whether it's a Formula Student (pictured) or a Formula 1 car

there will be a change in length of the section that the wave is propagating along so the return signal will be slightly out of phase. Fibre optics will come into their own in applications where large electromagnetic forces could severely effect the output of the load cell.'

Accurate inputs

The aim of a shaker rig is to simulate the effect of track inputs on the suspension, so the race teams can not only fully understand the characteristics of their particular suspension designs, but they can also experiment and tune the suspension to optimise their racecar's performance. To achieve accurate results, the inputs from the rig into the chassis have to be realistic, which is why the aeroloader actuators require the innovative 'compliant links' developed by Servotest, which we discussed last month. Furthermore, the drive file itself also has to be representative, and achieving this on a rig is quite a complex process.

Sine of the times

There are two main types of tests that are conducted on a shaker rig. The first is called a sine sweep. This is where the wheelpans are moved in a sine shaped movement. In this way the characteristics of the vehicle can be measured at each frequency, which is where the frequency of the wheelpan displacements gradually increases and the consequent performance of the racecar is measured.

Load cells are instrumental in measuring the displacement of the wheelpan and aeroloader actuators. They work by feeding back to the control system to achieve the precise control necessary to give realistic vehicle responses

The second type is a track replay. The first task before any form of track replay testing can be conducted is to generate a 'track file'. This contains the information on how to move the rig to match the car's suspension movements from the track. To do this the dynamics of the system need to be characterised by determining the frequency response function (FRF) or system matrix of the vehicle and rig combined as one system. This essentially measures the magnitude and phase of the outputs as a function of frequency compared to the inputs. These outputs are commonly shock potentiometers and pushrod load cells, while the inputs are the seven rig actuators.

'At the beginning of a track replay test, we have to develop the track file to replicate the

'The trick when designing a load cell is to make sure that the metal component is designed and then fabricated to give you the correct amount of deflection'

suspension movement from the track inputs,' says Christer Loow, engineering manager at Ohlins. 'We take the exact same sensors they use on the track and instead of recording them with the team's data acquisition system, we plug them into our data acquisition system. We then play random inputs for a few minutes, which moves the car randomly, and we measure the frequency response of the whole system.

'So, for example, we measure the effect on the left front suspension if you push or pull with the three downforce actuators and bounce with the four wheelpan actuators. We can then use this data to develop a computer model of the vehicle and the rig together as one system. We then invert that model mathematically. The inverted model is used together with the data

Socket set

Perhaps the most important performance factor of any racecar are those four rubber circles at each corner. Tyres are unquestionably a black art to design, model and analyse. Yet despite this complexity, advances in reliable data capture, correlation studies, and the available computing power these days, are all contributing to improving the accuracy of tyre models.

On a rig the main aim is to achieve an accurate representation of the suspension response from track inputs. However, these displacements are induced by the wheelpan actuators into stationary wheels. So not only are these wheels not rotating, but the consequent squash of the rubber under aero load is also not considered, leading to unrepresentative suspension measurements when compared to those monitored at the track.

But there's a way around this, as the physical tyre on the rig can be replaced by a virtual model, a load *vs* displacement DSP Simulink model to be exact, also known as a 'socket'. This model is inserted into the control loop in the DSP so it can influence the control in real time, and it allows the characteristics of the tyre to be modelled relative to car velocity and downforce, whilst the actuators drive the inputs into the wheel spindles. 'The Socket is a black box with as many inputs and outputs and as much complexity as the engineer judges necessary to model the dynamics under scrutiny,' explains Vincent Besson, R&D engineer at Servotest. 'For example, the effect of tyre growth in a spinning wheel or the influence of DRS on downforce can be simulated, as well as banking effects and a lot more which the teams do not tell us about.'

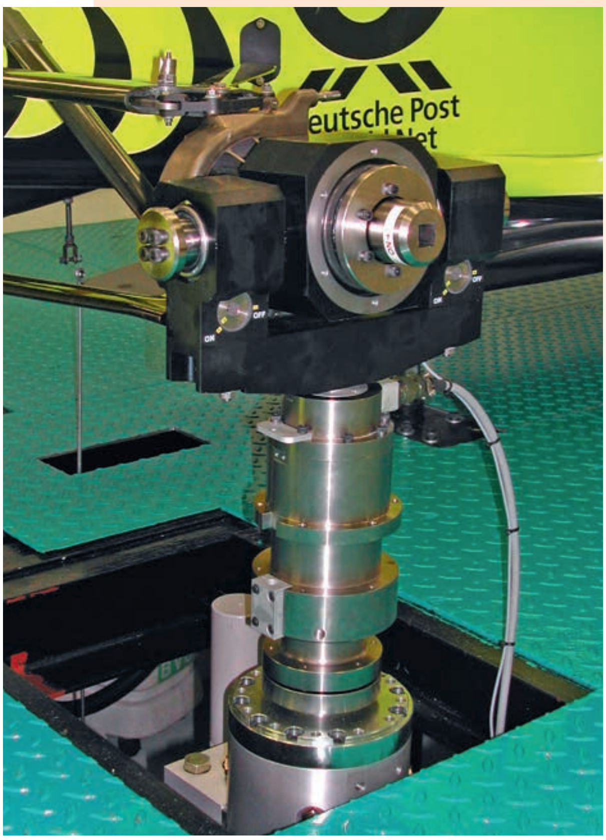

Sockets are virtual tyre models that replace physical tyres on the rig. They are integrated within the control loop to simulate the effect of tyre growth, and of downforce, during rotation. The displacements are driven into the wheel spindles (above)

The control system allows the operation of each actuator as well as the synchronisation between them. Here a single Optostar fibre optic purple cable on the node box allows noise-immune digital data communication

from the shock potentiometers and pushrod load cell data collected at the race track to calculate how we should move our seven actuators on the rig to replicate the suspension movements from the track.'

Essentially, the race teams use track data to iterate through different parameters to try and figure out what inputs are needed on the rig to recreate the outputs they saw at the track. Once this track file has been refined, this becomes a constant input into the test and therefore any measurements captured during the test are real responses which should match those experienced at the track.

Synchronisation

Another key factor in running realistic tests on a shaker rig is the synchronisation between all the wheelpan and aeroloader actuators. 'For a four-post rig, the key is to control all four actuators to exert their specific correct inputs and at the same time. Otherwise you'll be putting inputs into the car that it is not seeing at the track,' says Kowalczyk. 'For example, if you're inputting a heave input where all four actuators are supposed to be travelling at the same time and there is a delay between the front and rear, then you'll actually end up inputting some pitch, which is not what you want.

'For a seven-post rig, it's a similar problem,' Kowalczyk adds. 'But those additional aeroloader actuators that simulate downforce have to move with the car, because they are attached to the chassis. It's easy to apply a force to a single actuator, but it becomes harder when the car is moving and you are trying to make sure all of the actuators are tuned.'

To operate each individual actuator, as well as precisely synchronise them all together, requires an advanced control system and each company has developed their unique control theory strategy. 'This is where the magic happens,' says Kowalczyk. 'The control system for a standard hydraulic actuator is a PID [Proportional Integral Derivative] controller for displacement. So for a single actuator to displace one inch at a certain frequency in a specified time, it is a reasonably straightforward process, it is just a PID loop which you are able to tune. The key for multiple actuators is to synchronise them and each company has their own special control theory that allows them to do slightly different things.

'Our seven-post rig is a Servotest unit and we have developed our control strategy to allow us to apply downforce that varies with ride height maps,' Kowalczyk adds. 'This means you can take wind tunnel data and simulate aerodynamic changes because the ride height is changing. One of the stages of the control system is the PID control and that's not straightforward, which is why everyone has their own recipes.'

Pulsar control

One of the most popular control systems is the Pulsar digital servo-controller from Servotest. Pulsar is a second-generation real-time control system that utilises industry standard USB

A key factor in running realistic tests on a shaker rig is the synchronisation between all the wheelpan and aeroloader actuators

Control system data management method

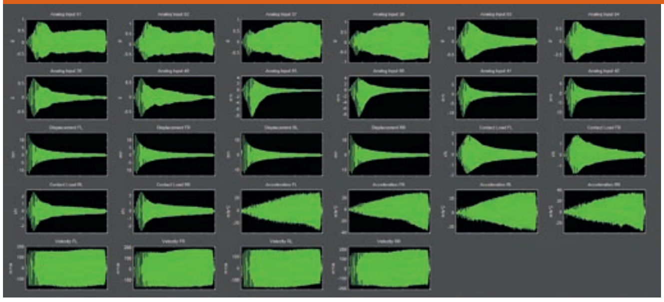

This is an excellent example of some of the inputs and sensor signals you can expect to see during a test on a top class seven-post rig, including live telemetry

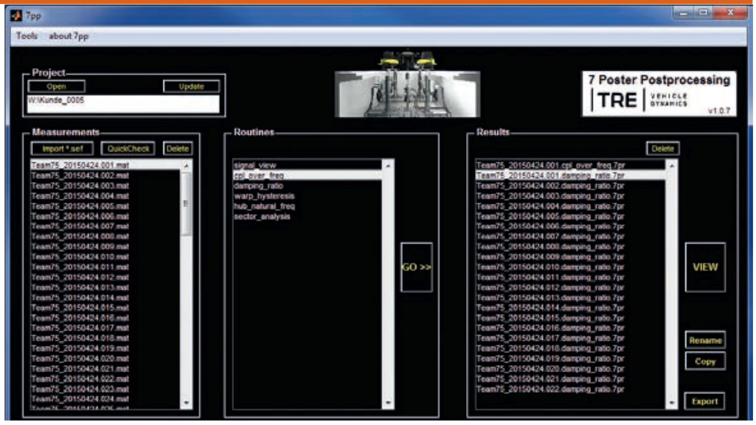

Recorded data can be exported into Matlab or Excel and analysed by TRE's custom toolbox

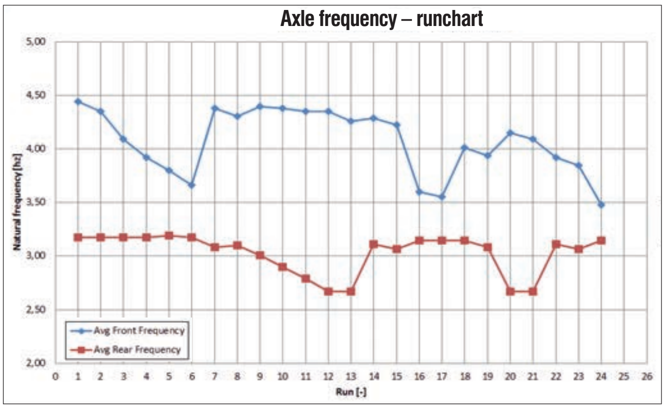

This shows how a plot of the basic analysis parameters recorded in Excel helps to illustrate the influence and progress of set-up changes throughout the rig test

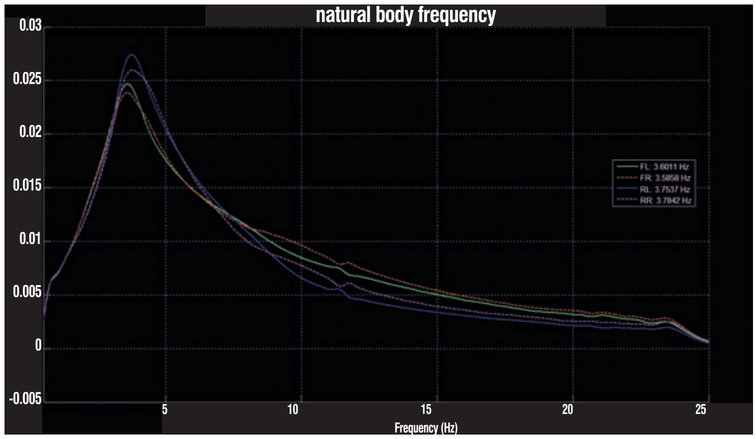

Graphs are automatically generated during post processing, with the displayed parameters depending on the customised code. Different companies tend to have their own approaches

'The ultimate set-up does not exist. It is always a compromise between a variety of performance related factors'

technology allowing for easy connection to any laptop or desktop PC. The new DSP within this controller is capable of 30 times the processing power compared to the older DCS2000 unit which means much more powerful control algorithms can be computed and implemented quicker. The Pulsar DSP also minimises gain and phase errors between the four wheelpans whilst ensuring robust control of the downforce actuators. Overall, all of these benefits mean that any set-up change on the racecar is fully reflected within the car's response, rather than the actuator response.

'We use the Servotest Pulsar system which is the main control system for the rig itself,' says Daniel Pfeiffer, who is the senior engineer at TRE GmbH. 'We also use our own TRE seven-post rig post processing toolbox which is very variable so the routines can be customised to fit the unique and specific needs of our customers. Essentially, all the rig and sensor signals are initially synchronised and logged via the Pulsar system and saved as a .sef file. This is then exported into Matlab or Excel format for further post processing using our own TRE toolbox and the raw data will then be handled and filtered differently, depending on the specified routines that we are using.'

'The benefit of Pulsar is it allows you to run all the complicated control mechanisms and loops such as the ride height maps and controlling the aeroloaders to be in sync with the car,' says Kowalczyk. 'There is also a lot of data coming in. We run some tests between 500-1000Hz so the control system has to be fast and do a lot of computation within that time, whilst continuing to log a lot of channels. You really need to eliminate the lag of all the inputs into the car, because the moment the actuators become out of sync then you are applying loads that the real car isn't seeing on the track, so the control system has to react fast.'

Translating data

Despite the monumental efforts and investments in technology and modelling to help deliver realistic results on shaker rigs, achieving representative suspension behaviour is by no means the last stage in the process. The final, and arguably most important task, is to translate the data and lessons learnt from the rig tests into useful results that the race engineer can implement at the race track.

'The ultimate set-up does not exist,' says Pfeiffer. 'It is always a compromise between a variety of performance related factors. Therefore, we provide our customers with a final set-up sheet which includes a list of various set-up options which are each specifically designed for a particular scenario. For example, we suggest optimum set-ups for changing weather conditions; fast tracks which require increased aero platform stability and slow tracks which focus on mechanical grip and traction. We also run set-up matrix tests where we sweep through a matrix of different set-up options such as damper settings to identify sweet spots for different performance parameters plus linear and non-linear effects on car performance. This provides the race engineer with a great guide for set-up work at the race track.'

Reality check

With track testing restrictions and the desire to win placing huge emphasis on arriving at a track with an already optimised racecar, seven-post rig testing is an essential tool for keeping one step ahead. But can future developments in shaker rig technology ever reach the accuracies required to replace track testing all together? Probably not, as no simulation ever fully matches reality. But one thing's for sure, rigs are continuing to edge closer to this.