PART 1

The explosion of data logging equipment has made the technology accessible to the club competitor, but where on earth do we start when we are looking to buy? Graham Templeman guides us through the minefield in the first article of a new series

IN RECENT years, data logging has moved into the range of things that any club competitor could reasonably aspire to. Using the traditional unit of currency of a set of new tyres, entry level systems now start at less than two sets and climb from there.

Six sets buy you all the information you could ever want. On the other hand, if you look round any collection of competition cars you will quickly find ones fitted with loggers that are not being used.

Over the course of the next few issues, we will be looking at whether or not it is worth making the investment and if you do, how to maximise the payoff and avoid being stuck with another expensive shiny bit that does not earn its keep.

Do not forget that the costs are not only financial but also in terms of effort. Time spent managing the system and the data is time not available for other activities. There is no sense in knowing exactly how fast you were in the few minutes leading up to the accident that happened because you did not have time to do a spanner check.

However, the benefits can be enormous. You quickly build up a bank of data that is invaluable for both preparation and competing on the day. Before the event, and wearing the engineer's hat, you can look at the data for guidance on gearing (speed and revs), suspension settings (damper travel, g force and corner radius) brake cooling (time under braking and time for cooling) and a whole lot of other things.

As driver, you will be able to look at braking points, early or late apexes, cornering speeds, sector times, gears and target revs for the straights and which corners caused you the most trouble the last time you were at that circuit. This and much more will be available before the event and it will all be based on hard evidence, not hazy recollection.

At the event, you can compare performance with the last time at this track and with what you now expect of the car and the driver. You can even check the old driver's excuse that the engine lost power through the day.

To help you select the right system for your application, we have put together a chart that shows facts about what each manufacturer calls their entry level system. The unit in currency is deliberately vague because, as ever, there are all sorts of side issues involved in making such a major purchase. Do your own negotiation and do not overlook important issues such as whether trackside beacons are necessary and ▶

> "Benefits can be enormous. You quickly build up a bank of data that is invaluable for both preparation and competing"

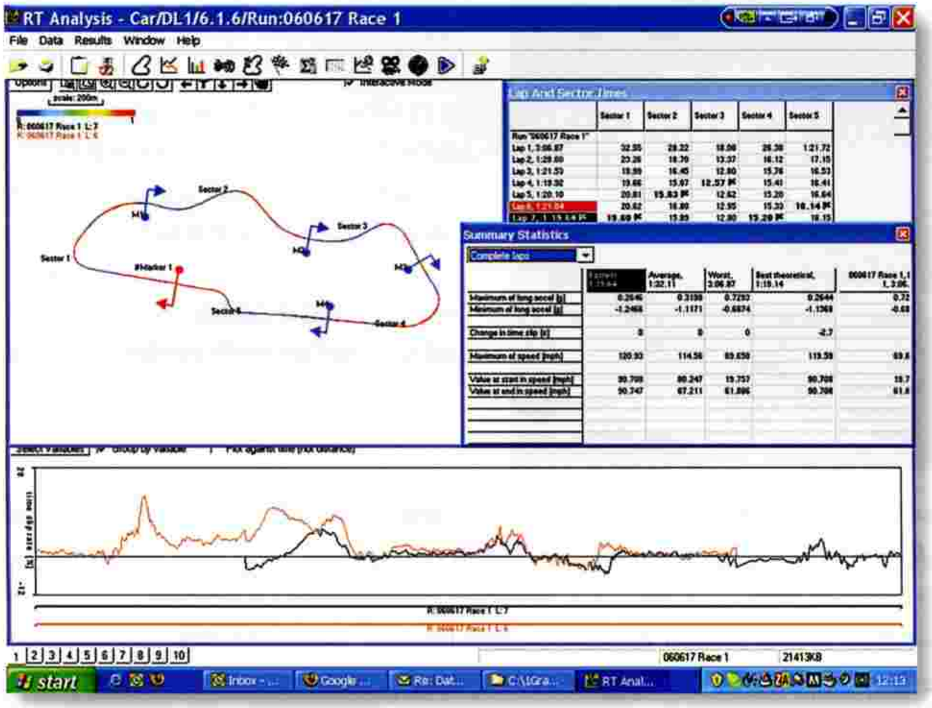

	Sector 1	Sector 2	Sector 3	Sector 4	Sector 5
Run "060617 Race 1"					
Lap 1, 3:06.87	32.55	28.22	18.06	26.38	1:21.72
Lap 2, 1:29.80	20.26	18.70	13.37	16.12	17.15
Lap 3, 1:21.53	19.99	16.45	12.80	15.76	16.53
Lap 4, 1:19.92	19.66	15.67	12.57 ⚑	15.41	16.41
Lap 5, 1:20.10	20.81	15.83 ⚑	12.62	15.20	16.64
Lap 6, 1:21.04	20.62	16.80	12.95	15.33	16.14 ⚑
Lap 7, 1:19.64 ⚑	19.60 ⚑	15.89	12.80	15.20 ⚑	16.15

Complete laps

	Fastest, 1:19.64	Average, 1:22.11	Worst, 3:06.87	Best theoretical, 1:19.14	060617 Race 1, L 3:06.
Maximum of long accel [g]	0.2646	0.3199	0.7293	0.2644	0.72
Minimum of long accel [g]	-1.2468	-1.1171	-0.6874	-1.1369	-0.68
Change in time slip [s]	0	0	0	-2.7	
Maximum of speed [mph]	120.93	114.56	89.650	119.59	89.6
Value at start in speed [mph]	90.708	80.247	15.757	90.708	15.7
Value at end in speed [mph]	90.747	67.211	61.886	90.708	61.8

LEFT The analysis software can offer huge benefits to club competitors, but it makes more sense if we understand how the data got there in the first place

OW TO AVOID BEING DAZZLED BY DATA

QUICK GUIDE TO ENTRY-LEVEL UNITS

Model	Price (Sets of Tyres)	Dash?	Standard (built in) Channels	Number of External Inputs	Resolution (Bits)	Sampling Rate (Hz)	Memory
Aim XGLog	2.25	Yes including lap times	lat + long g/ rpm/ wheelspeed	5 Analogue 2 Wheelspeed	12	200	8Mb
Astra Microlog	2.5	Slimline readout of 2 channels	speed/rpm	2 Analogue	8	10	21 minutes
Motec ADL 2	7	Yes including lap times	lat + long g/ rpm/speed/ lap timing/ pressures etc.	8 Analogue 2 Wheelspeed 4 Thermocouple	12	1000	8Mb
Pi X-Sport Dash and Compact Logger	5	Yes including lap times	lat + long g/ speed/rpm/ pressure/ temp	None (other models provide for external inputs)	8 and 10	1 (temp + pressure) and 10	Compact Flash
Race - TechnologyDL1	1.75	Extra	lat + long g/ speed/rpm	4 Frequency 8 Analogue	16	200	Compact Flash
Stack ST8802S	7	Yes with predictive lap timing	Lat g/rpm/ wheelspeed/ lap timing/ pressures & temps	9 Analogue	8	1 to 20	Up to 240 minutes

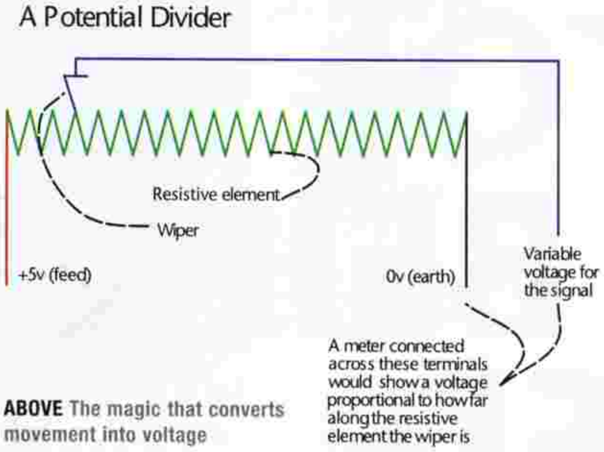

ABOVE The magic that converts movement into voltage

LEFT & BELOW A variable resistor being used to measure throttle position (left) and steering position (below)

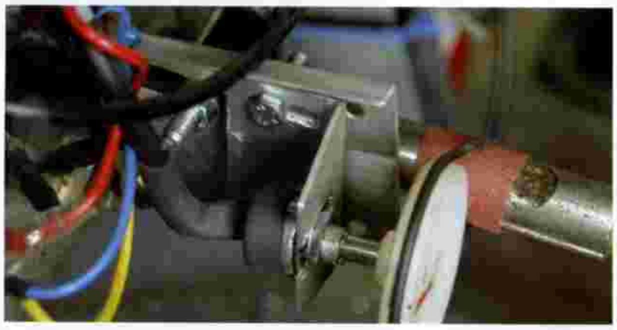

what sensors are included in the package.

Before making the choice, we need to look at some basic concepts. The logger itself is simply a device for reading analogue data – voltages and frequencies – and storing them as numbers that can be read in a software package. The pretty pictures of the analysis software will make much more sense if we understand the way that the data got there in the first place. So let's look at how it happens.

RESOLUTION

The analogue data is converted into number form by a process known as analogue to digital conversion or ADC. A logging system uses binary digits to optimise the use of the sometimes limited memory on board the logger unit. This keeps the file sizes low and reduces download times.

Computers love binary numbers because the ones and zeros make the best use of the on/off states of computer memory but we should think about the form of the number. An eight bit number consists of a mixture of eight ones and zeros and is capable of registering up to 256 separate values. Ten bits will count to 1,024 and 12 bits to 4,096. The reason that we should care about this is that it is one of the things that determines the level of precision that the system can achieve. Eight bit accuracy is common and adequate for many uses although 12 bits are beginning to appear in entry-level loggers.

In the case of a throttle position sensor that measures the 90-degree rotation of the throttle spindle, 8 bit accuracy will split this 90 degrees into 256 separate values – or less than half a degree – which is more than adequate. In fact we could probably understand what the driver was doing using a simple five point scale of no throttle/a bit of gas/half throttle/almost full/flat to the floor.

On the other hand, if we are trying to understand the suspension behaviour of an off-roader with 250 mm of total suspension travel, 256 divisions will only tell us to the nearest millimetre. Given that we are unlikely to be able to create a mechanical linkage that takes advantage of the full range of sensor movement, we might only use 75% of the available measurement so we would be down to a maximum resolution of 1.3 mm. Add in the fact that because we are converting from smooth analogue values to discrete digital steps, the voltage that the ADC translates into 26 mm could represent anything between 25.35 mm and 26.65 mm.

So three consecutive readings that the system interpreted as 26 mm could have been 25.4 mm, 26 mm and 26.6 mm. If we then want to go on and examine damper shaft speed and acceleration we are building-in errors upon errors because the position data will be used to calculate the speed and acceleration numbers. We would end up with an impressive set of numbers that we ought not to trust at all.

> "We could end up with an impressive set of numbers that we ought not to trust at all"

THE POTENTIAL DIVIDER

The magic that converts movement into voltage is nothing more complicated than a variable resistor. This is configured as a potential divider so that the voltage input at one end, usually 5v but sometimes 12v, is reduced to a signal voltage that is proportional to the position of the linkage. The diagram shows what happens.

The supply voltage is fed into one end of the variable resistor and the other end is grounded. The third wire – attached to the resistor's wiper – sees a voltage that varies according to how far along the resistive element it is. At one end there is the full supply voltage – because the resistance does not come into play – and at the other there is no voltage. Guess what happens in between. Attach a mechanical linkage to turn the resistor and we have a way of measuring movement. The photographs show variable resistors being used to measure throttle position and steering. ▶

ABOVE Sampling rates are important. Ten samples a second are fine for driver inputs but even 100 samples a second will give sketchy suspension data if the car is covering 50 metres per second

SAMPLING RATE

The sampling rate can also be important. There is a difference between the rate at which a parameter is sampled and at which it is displayed in the software. To do the job properly we should be sampling at least twice the rate at which things happen in order to be sure of capturing all the events, but again it depends on what you are trying to measure.

Many loggers will tell you the battery voltage but you really do not need that figure updating more than once per second. Driver inputs – steering and throttle – are OK at 10 samples per second and acceleration traces – lateral and longitudinal g – look very jagged when displayed at more than tenth of a second intervals.

> "Keep an eye open for whether or not your shortlisted suppliers provide free updates of the software"

However, if you are looking at suspension movement, 100 samples per second is only at the lower end of what is reasonable. This makes sense if you think about it. The car might be travelling at 50 metres per second, so even if you take 100 readings per second, we still only know what is happening every half-metre. This gives us some insights about roll and pitch but it falls far short of being a substitute for a seven-post shaker rig.

FREQUENCY

Data loggers also sense frequency to enable them to count engine revs and wheel speeds. The methods used are more complex than for sensing movement and some systems are capable of cleaning up a noisy signal where others will only deal with a nice sharp signal. It is a case of following the manufacturer's instruction.

From the point of view of choosing a system, the maximum frequency, as opposed to the sampling rate, that the system can read is not usually significant. The manufacturers measure this in Hertz (cycles per second) and it will generally be a big number. Wheels rotate relatively slowly – not more than 50 revolutions per second – and even 16,000 rpm motor cycle engines are only running at 267 Hz. Do not forget to factor in the number of pulses per revolution, so a wheel speed sensor looking at 4 bolts would be in the 200 Hz range but trying to measure engine speed by pointing a sensor at a 20 tooth crank pulley lifts us into the 5,000 (5 kHz) range.

LEFT & BELOW The opposite ends of the spectrum for entry-level systems: Race Technology's DL1 (left) and Stack's ST8802S. The DL1 will be used to collect illustrative data for future articles in this series

MEMORY

It is a bit difficult to be prescriptive about the amount of memory you need. More is obviously better than less, but much depends on how big are the run files that the system generates. Some loggers have less memory because they need less and the important question becomes not how many megabytes, but how many minutes of logging can be achieved between downloads. This is one enormous advantage of loggers with removable memory cards – the other being that the data geek simply removes the card and leaves the car to the engineers.

SOFTWARE

Most manufacturers are now prepared to put their software and some sample data on the net so that it is possible to test drive the package before buying. This gives you a chance to play with it but it is still not the same as getting to know the software in detail. However, at least you will get some feel for what you can and cannot do. Keep an eye open for whether or not your shortlisted suppliers provide free updates of the software or whether you have to pay for updates

SHOPPING TIME

With the basic understanding in place, it is time to look at the table and the manufacturers' websites and brochures. It is the usual price, performance and features trade-off to find the systems that best meet your needs.

Once you have narrowed the choice, phone the manufacturer and ask to talk to a sales engineer. You are bound to learn things and will get a feel for how good the after sales service is going to be – and these systems are complex enough to guarantee that you will need some external help at some time or other.

Next month we will look at what data is available and how to make sense of it. ■

ACCESS ALL AREAS

PART 2

Data logger now chosen, the world – or at least the racetrack – lies at our feet. But what do we do with those pretty pictures once we have them on screen? Graham Templeman gives us a few clues

IF YOU are not careful, data will tell you the blindingly obvious.

The throttle histogram almost invariably shows that lap times are minimised when the driver spends more time on the throttle. You can also learn that the fastest laps come by carrying speed through the corners and on to the following straight. Did we really need to spend a week and the price of three sets of tyres to be told that?

We have to look for more subtle signals and it is best to have some sort of organised method for finding them. And a step before that, we will look at some sort of organised approach to collecting, analysing and managing the data.

Even the smallest teams need a strategy for managing the data. It sounds obvious, but the logger needs switching on every time it goes on track. It also needs downloading whenever it comes back again and the data needs looking at. The driver is a good person to be given the responsibility of switching it on (it makes him or her feel important and avoids the 'I thought you switched it on and now we will never know how I broke the lap record!' sort of a tantrum).

Downloading is best done by someone else in the team because drivers are always adrenaline-fuelled when they get back and would probably break something anyway. Interpretation is a multi-stage process involving everyone at some point.

Data should be examined immediately at the end of each run. It should also be looked at again without the inevitable time pressures. The first attempt gets a general impression of how

FIGURE 1 TIME SLIP

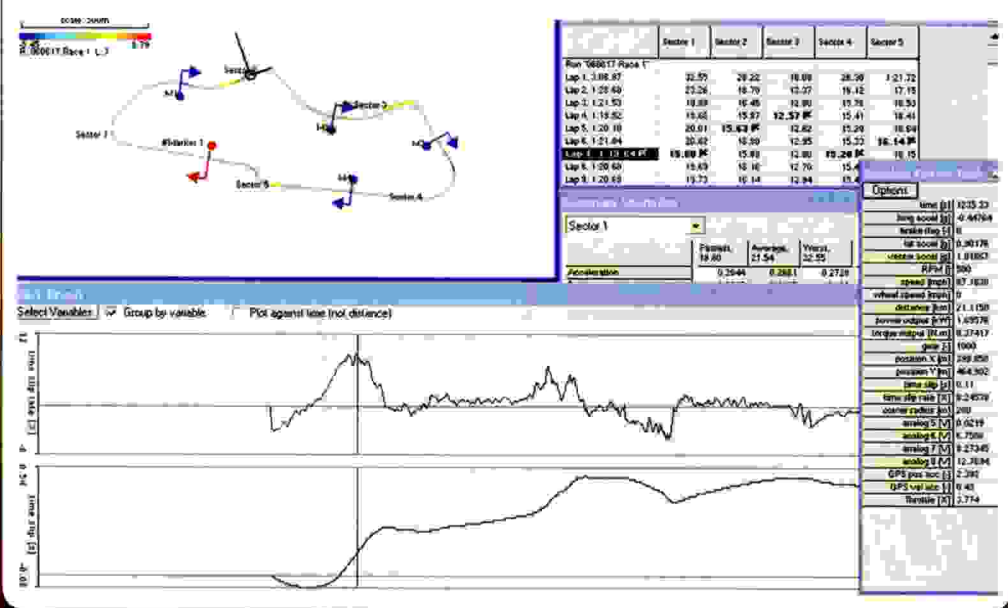

LEFT The upper black trace shows the rate of time slip, the lower one the gradual build up of time lost through the lap. The RT software shows nothing for the first part of this lap (Lap 7) because Sector 1 was the fastest sector time of all. It is difficult to interpret from a static trace like this, but on a computer the cursor could be moved to the point of interest and the location on the circuit could be identified from the map. In the trace shown here, time was saved (0.1 sec) under brakes for the fast left hander (the Craner Curve) but more than that (0.25 sec) was lost under braking for the fast right hander (the Old Hairpin). The map has been specified to show rate of time loss from blue (good) to red (bad) things were going, the second, third and fourth attempts can be more leisurely, often just browsing to see what you see, but also equally often with a specific goal in mind – such as evaluating the performance of the brakes or whatever.

There is always a big temptation not to look at the data after the last run of the day. Everyone just wants to pack up and start the journey home. This is a pity, because it is always best to look at things while everyone is still together and while events are still fresh in people's minds.

Sadly, the excitement when the first data comes out of the system soon fades when it simply looks like a bunch of uninterpretable squiggles. The traces hold a huge amount of information but it has to be decoded and this can be pretty daunting. But if it is approached methodically, you will extract an awful of useful stuff.

> "The traces hold a huge amount of information but it has to be decoded and this can be pretty daunting"

The logger gives you the *data* (individual facts, like the fact that the engine was pulling 7,200 rpm at the apex of a bend) and it puts lots of bits of data together in the form of charts and histograms from which it is our job to extract *information*, but what we are really after is *knowledge*. The logger won't tell you that there is understeer at turn-in which could be easily fixed by adjusting the dampers. Not in so many words, but with a bit of experience it becomes possible to identify the shape of the lateral g curve that gives you the clue.

So assuming that the driver is away doing the 'Did you see me drive round the outside of you at the Hairpin' routine, where can we start? We should look at the basics, just to reassure ourselves that there are no problems waiting to catch us out later.

If we are logging temperatures and pressures, a quick scan of the whole session can be made to ensure that they have been within limits. Revs come next. Has the engine been buzzed and if so how seriously? This leads on to gear ratios – is the engine reaching peak revs at the end of the fastest straight? Is every corner being taken at engine revs that are within the power band? If not, what changes need to be made to the ratios?

Then look at performance level in general. You pretty soon get a feel for the sort of performance to expect from the car on a good day. Much of the data on show here and in future issues is taken from Sports 2000 where cornering of 1.6g is about right and braking of 1.3g. A small amount of aero downforce is hoped for and not always achieved. Acceleration in the lower gears can be .4 to .5 g, tailing off to .2 at higher speeds. But these figures represent the ballpark that we need to be in and if not, questions need to be asked and solutions found.

This is where it is useful to have someone in the same class with whom you share data – it might just be a slow day and you might be chasing your tail. When you have built up a bank of data, you can look at your data for this circuit on previous outings.

By now, the driver is back and we can look at things in more detail. Drivers inevitably, and quite rightly, want to focus on the ▸

FIGURE 2 INCONSISTENCIES

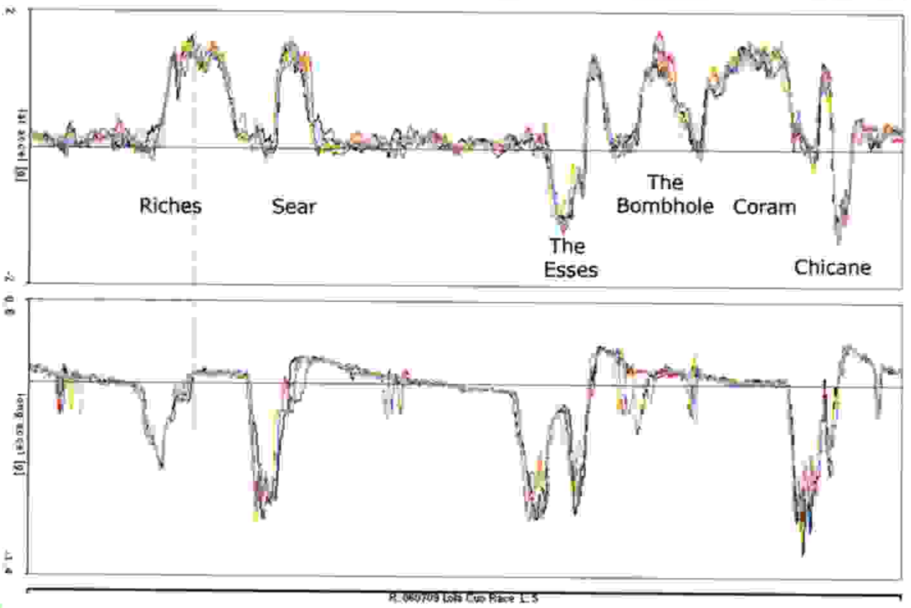

LEFT Overlaying a number of laps shows up where there are inconsistencies that would bear further investigation. The lateral g trace (top) shows that Sear, the second part of the Esses and the Chicane seem to be fairly well under control. These are the three slowest corners so we need to discuss with the driver where the problem lies with the faster corners. The long g trace shows inconsistencies in braking for the Esses and the Chicane, so again we need to establish whether it is a matter of brake set-up or driving technique fastest lap and whether any further time could be found. The tool for this is a graph showing time slip. This enables you to compare where one lap is faster than another. The name used and the form of the display varies from one software package to another but FIGURE 1 is taken from the R-T Analysis package which shows time slip (the cumulative time lost through a lap) and the rate of time slip (how much faster or slower this part of the lap is, measured as a percentage). These charts show precisely where time is made or lost and can then prompt further investigation as to why that has happened.

> "Overlay as many laps as the software allows. What you are looking for is inconsistency"

Useful as the time slip analysis is, there are other methods. One is to overlay as many laps as the software allows on top of each other. Almost any combination of speed, revs, steering and g forces will do, because what you are looking for is inconsistency.

Some corners will have nice consistent traces where one lap sits on top of another lap with hardly any difference. Some corners will have traces which are all over the place. The consistent corners are the ones where there is no immediate problem. The driver is in a routine that gets through them without any drama. Whether they should be taken faster or on a different line is a question for later in the analysis process. FIGURE 2 shows this approach. ▶

DL1 LOGGER

MOST of the data and analysis charts shown in this series have been compiled using Race-Technology's DL1 logger and the RT Analysis package. It was chosen because of the ease of set up, the large number of available channels and the fact that it is the lowest cost logger available. There can be no accusations that this was a high budget operation where any problems were solved with the application of copious amounts of money.

Unlike the others at this level, the DL1 uses GPS technology not only to draw very high quality maps but also to measure speed. The GPS measures its position on the ground 10 times per second and combines this data with that from its on-board two axis accelerometers to provide accurate data.

The two immediate practical advantages are that it does not need a lap beacon to trigger it and it is not necessary to install and calibrate wheel speed sensors. Lap beacons are easily forgotten and left on the pit wall and speed sensors tend to be expensive and are easily damaged in an accident. Calibrating them can be a pain because the rate of tyre growth at speed is not easily measured, so the rolling circumference is always a best guess.

Sadly although the maps are accurate enough to show what side of the road the car crossed the startline, without definitive maps of the boundaries of the track, they are not much use in evaluating racing lines. A race circuit is no place to do two slow laps surveying the inside and outside kerbs!

Because of the basic difference in concept, the software works somewhat differently to conventional packages, but still provides all of the functions needed for proper analysis. It also has a few bells and whistles not easily available to conventional systems, but that will be dealt with in a later article.

FIGURE 3 BRAKING

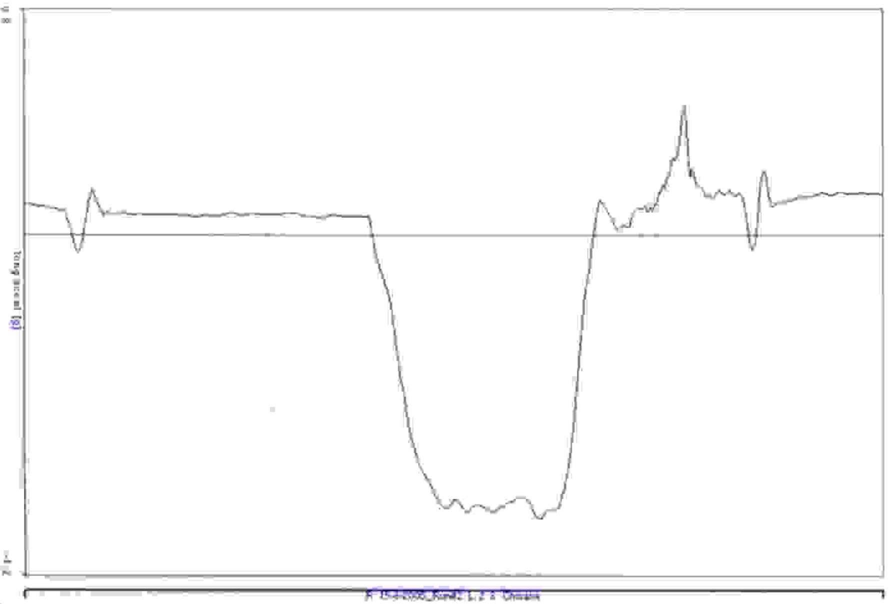

LEFT This trace shows a production-based car under heavy braking. The blip at the beginning of the trace is the last gearchange along the straight, the small positive long g figure shows that the car is running out of steam and when the brakes are hit, the downward slope shows the build up of braking effort. The flat portion is the time spent at peak braking and the steep climb shows the driver getting out of the brakes before turning in

The inconsistent corners are giving you a strong message that there is a problem there. For some reason it has not been possible for the driver to get through those corners consistently. It could be a driver problem, being unable to get to grips with some aspect of the corner, but it could be that the car is not allowing this to happen. Don't forget that one of the inconsistent approaches must be quicker than the rest and this is where the time slip trace can be useful.

Use the traces and some common sense. Look at the whole corner and try to sort out where the inconsistencies come from. Does it start under braking – does the car arrive at the corner in a nicely composed manner with an orderly transition from braking into cornering? Use the lateral, longitudinal and combined g charts to decide.

With the braking, what you are looking for is a long g trace which plunges sharply downwards, flattens off and then climbs upwards as the lateral g trace builds into the corner. This represents rapid application of the brakes (the sharp downward plunge), sustained heavy braking (the flat portion), and then a climb up to and over the zero line.

> "Usually, what a driver calls opposite lock is a reduction in the amount of lock applied"

The climbing long g trace does not indicate that the car is speeding up, only that the rate at which it is slowing down is reducing. Actual acceleration starts when the trace crosses the zero line. FIGURE 3 shows a typical braking trace – this time of a production-based car achieving almost 1g of braking effort.

The downward slope shows that the build up took about 1.3 seconds and 2.3 were spent at or near peak braking. The steep return slope indicates that this is not a car that encourages the driver to turn in while braking.

Once the car gets to the bend you need to consider what happens next. Don't abandon the traditional methods that you used before data logging. Split the corner into phases and consider each one. We've already considered the braking phase but what can we now learn about the turn-in, the mid-phase and the exit?

The lateral g curve is full of information but it contains two separate strands. Its general shape indicates the line that the car takes through the corner. A symmetrical trace shows the classic racing line through the bend – the one that starts wide, moves towards a central apex and then drifts out wide again. A trace that peaks early and then flattens out indicates a late apex. This is a bit counter-intuitive but it makes sense. To take a late apex, the driver has to turn in fairly sharply and then straighten the line and aim at an apex towards the end of the bend. The opposite – a gentle lead-in with a sharper curve towards the end of the corner – indicates an early apex: the driver has committed to an early clipping point that will eventually force a sharper curve towards the end of the bend.

The car's handling will disguise these shapes. Oversteer will generally make the trace ragged as the car grips, loses grip with opposite lock correction and then grips again. The effects of oversteer on the lateral g curve is to put rounded spikes into the curve at a rate that reflects how quickly the car reacts. A purpose-built race car will react more quickly than a modified production car and the frequency of the spikes will be shorter. If a steering trace has been logged, the symptoms are unmistakable and a throttle trace can often tie the oversteer in to the driver's right foot.

FIGURE 4 CORNERING TRACE

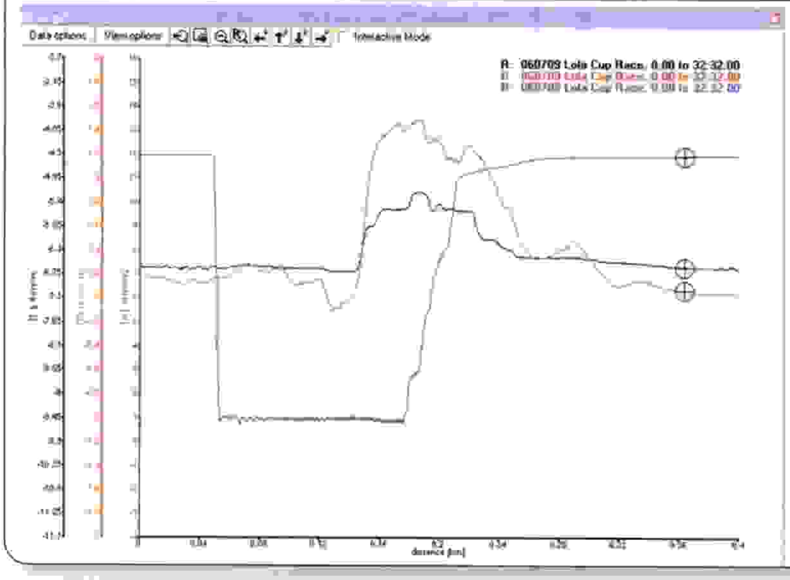

LEFT Here lateral g is in red, steering input in black and throttle position in purple. The general shape of the lateral g curve indicates a late apex (the highest g forces are seen at the entry to the corner so that most of the turning is done early and the driver can get on the throttle as soon as possible). The car's handling confuses the issue slightly because there is a trace of mid-corner understeer – the driver felt the need to use more steering after the throttle was applied and the lateral g forces decreased slightly. The front end then got some bite and built up cornering force again

FIGURE 5 OVERSTEER

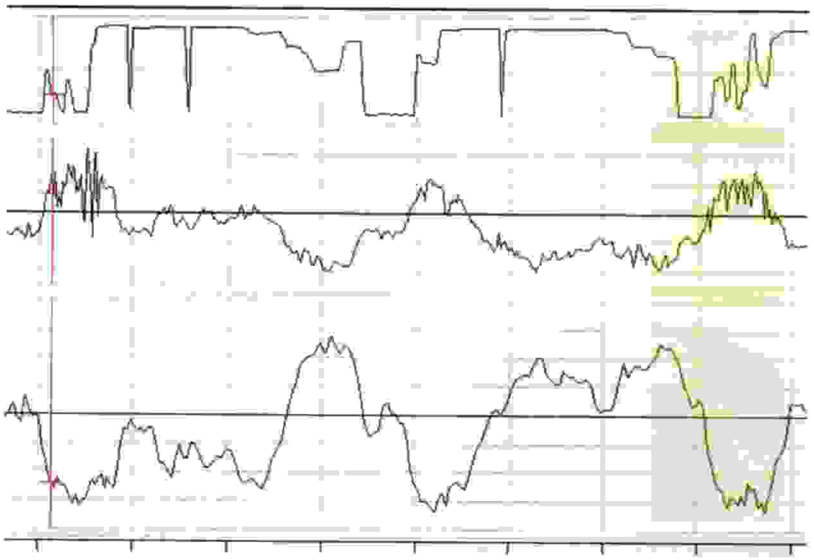

LEFT In this screenshot the top trace is throttle, the middle is steering and the bottom is lateral g. Looking at the green shaded area, the driver is very anxious to get into the throttle (top trace) but doing so necessitates lots of steering correction (middle trace) and the need to back out of the throttle again. This gives the lateral g (bottom trace) the characteristic jagged appearance associated with oversteer

Opposite lock, in the sense that the driver is steering to the right in a left-hand bend, rarely occurs. Usually, what a driver calls opposite lock is a reduction in the amount of lock applied. What normally happens is that the trace indicates the steering correction to be pretty rapid and limited to about 20% of the applied lock. Once the correction is made, the original steering angle is almost always put back on immediately.

ALL IS REVEALED

The sequencing of the traces is interesting. You can expect to see a steering correction *followed* by a reduction in lateral g. So what the driver feels is the car beginning to break away through a reduction in self-aligning torque of the steered wheels *before* the cornering force diminishes. The back steps out and reduces the slip angle of the front tyres, which are still gripping at this point, but this reduces the torque felt at the steering wheel and prompts our hero to wind back the steering and probably get out of the throttle as well.

Then when the crisis goes away, it's back on with the lock and the power. You can tell a lot about the driver and the car by working out what goes on when the front has more grip than the rear. All of this excitement serves to hide the overall shape of the curve from which we originally hoped to understand the line that the car took.

Understeer can be more subtle. You might see the lateral g forces fall away as the driver applies the throttle and this could also be tied in with an increase in steering angle. On the other hand, understeer is often shown up as a flat area on the g trace – often towards the end of the corner – or diminishing g levels through the corner.

But there are other variants. To confuse the issue, you might find that the steering trace reflects the driver's desperate attempt to kill understeer by rapidly putting on lots of lock at the entry to the corner and winding it off again in reaction to the slide that this has created. The trace climbs rapidly, heads momentarily in the opposite direction and then resumes a more normal pattern. Engineers of front wheel drive cars might recognise this, but it is not something that presents itself in Sports 2000 too often.

FIGURE 4 shows understeer and the challenge is to understand why. Clearly the driver influences the situation by the application of throttle but could the car be improved? Peak lateral g is just a touch down on what would be expected from a car of this type. The turn-in was clean enough (so the dampers were doing their job) and the throttle was applied early so the driver felt confident. Some change to the relative roll rates front and rear might kill the understeer but before embarking on this change a whole range of factors need to be considered.

The decision would be easier on a test day than race day, and cannot be made on the basis of just one corner. It would need to be a repeated pattern to justify any changes and in any case mild understeer might be part of a set up strategy that makes the driver feel comfortable in a long race.

FIGURE 5 shows oversteer created by a greedy driver who is very impatient to get on the throttle. The early throttle application brings with it steering corrections and maybe a little more patience would benefit the lap time.

Reading the data adds considerably to the workload on race day but, done methodically, should improve the quality of decision making. It is very easy for egos to get in the way, with people adopting their own interpretation and defending them rather than admitting that there are other possible explanations.

It is very important that the data be re-examined as soon as possible after returning to base when heads are cooler and there is time to take a more leisurely and opportunistic look through the data. In the next instalment we will consider the range of other information offered by the various analysis packages. ■

PRECIOUS GAINS AT THE RAINBOW'S END

From rainbow maps and histograms, to XY charts and maths channels, Graham Templeman explores what we can do with data logger software

EVERY data logger comes with its own software package and, luckily for us, the manufacturers keep an eye on what their competitors are up to. This means that when one company comes up with an innovation, the rest soon follow. So the range of facilities is pretty standard and it all becomes fairly intuitive with a bit of practice.

HISTOGRAMS

All the packages will provide the facility to produce histograms showing the extent to which a particular variable was present. They are a useful starting point to review the broad general behaviour of the car. The usual set includes engine revs, wheel speed and throttle position. Damper histograms are also useful as a car sorting aid and will be dealt with in a future issue.

An example, taken from the Race Technology DL1, is shown as Figure 1. The three charts show data from the fastest four laps. The left-hand chart is the RPM channel which shows that for about 75% of the time the driver was able to keep the revs in the 5,000 to 6,500 range which means that it stays pretty well in the best part of a Sports 2000's engine power curve.

> “Track maps were originally nothing more than a means of navigating through the data. Now they are capable of more”

The speed chart is on the right hand side and shows the range of speeds for which we have to provide gear ratios. On this day, the range was between 40 and 120 so that gives us a starting point for first and top gears. There is a plateau in the 45 to 70 range, indicating some time spent

FIGURE 1 HISTOGRAMS: A USEFUL STARTING POINT

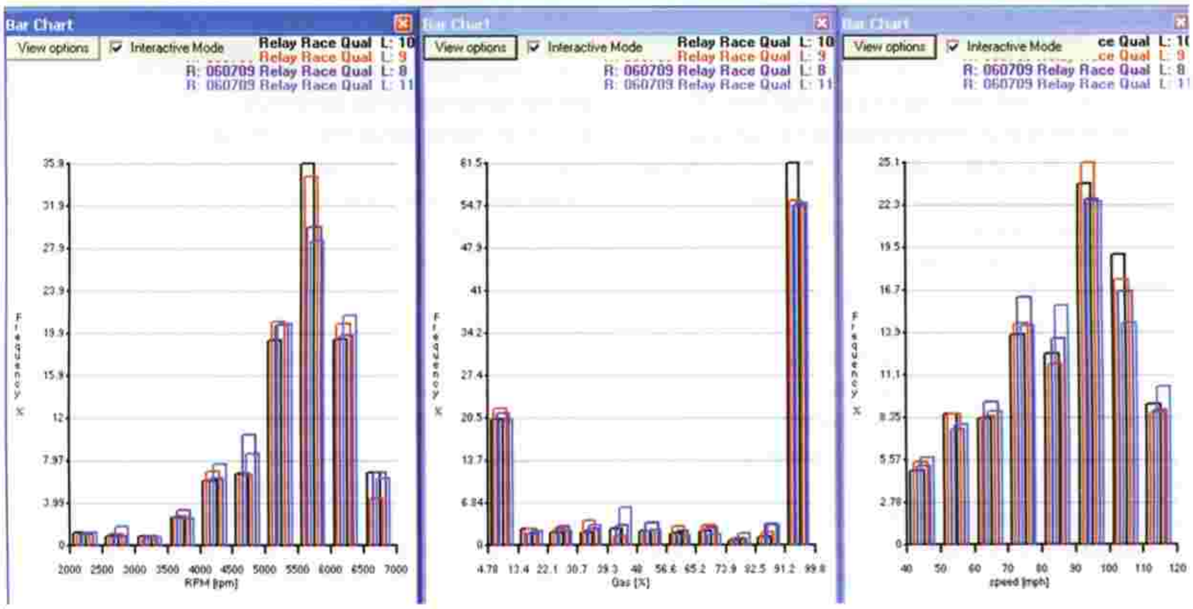

FIGURE 2 BRAKE TEMPERATURE RAINBOW MAP

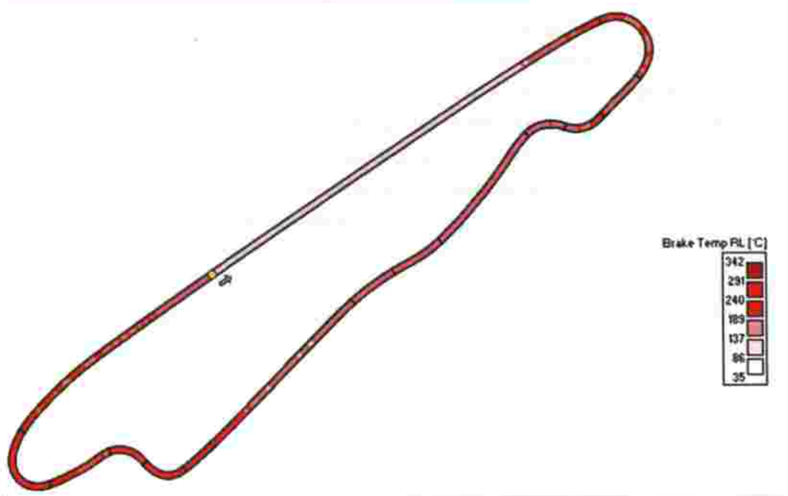

FIGURE 3 RAINBOW MAP – TIME SLIP

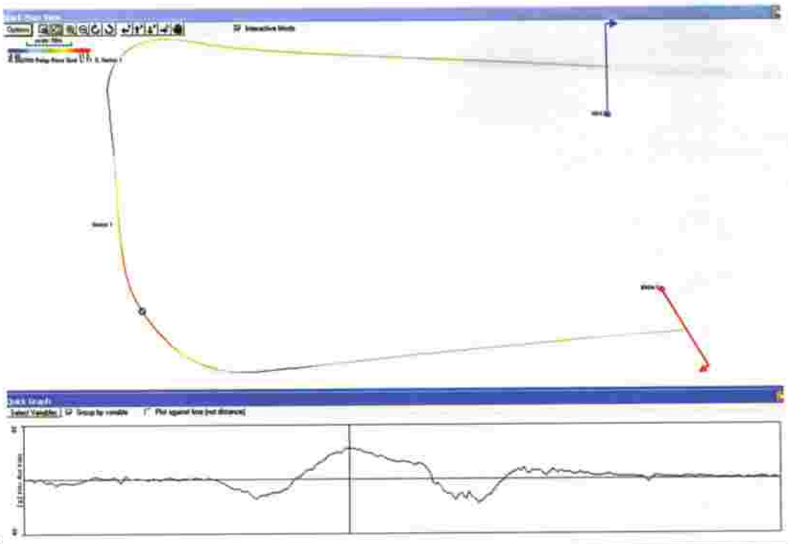

in corners at those speeds, so we need to look at the gear charts to find a ratio that would be in the power in that range, and the same goes for the 70 to 90 range.

The tail end of this histogram ties in with the upper ranges of the RPM chart and perhaps indicates that the current top gear is on the low side since the engine is being forced to run beyond its peak power. It would be foolish to rely entirely on the histograms to decide the gearing when strip charts of RPM and speed will confirm the speeds through the turns and the revs and speed at the end of the main straight. But the histogram gives an instant feel for what is going on with the car.

The throttle chart (the middle of the three) is fairly typical in that it's either full on or full off with little in between. Perhaps this one is a bit more point and squirt than you might expect given the ratio of power to grip, but that in turn raises the question of whether this is related to the chassis set up or the driver.

The fact that the driver is not spending time on part throttle, mid-corner, could indicate that the chassis is either so nervous that the driver feels inhibited or that it is so good that there is no messing about with half measures. Back to the strip charts and the driver for further evidence. In this case the fastest lap is the black lap and, as might be expected, this is the one where significantly more time is spent on the gas.

MAPS & RAINBOW MAPS

Track maps were originally nothing more than a means of navigating through the data. At a given point on the circuit these characteristics were present. Now though, they are capable of more than this. The track animations with the fast blob and the slow blob racing each other round the circuit map are all very well for impressing a potential buyer, but don't impart any information that cannot be gained elsewhere.

More significant is the ability to shade the map according to the data. Two rainbow maps are shown. The first (Figure 2) is taken from the sample data provided with the free download of MoTeC's Interpreter software. It shows the range of brake temperatures as the car goes round the circuit and indicates that it stays within a safe working limit even in the heaviest braking zones.

The second map (Figure 3), this time from the DL1, shows the rate of time slip in one sector of the circuit. This refers to the speed difference at any point and compares the current lap with the fastest *sector* recorded that day. A positive number is bad and on the map this is shown at the red end of the spectrum with blues showing time saved. The map shows clearly the cardinal sin of in fast, out slow. The sector shown involves time saved under braking, which then compromised the corner and resulted in losses in time through and out of the bend. The strip chart shows where the time is lost or gained, the rainbow map illustrates it starkly. ▶

XY CHARTS

The remaining graphical presentation is the XY chart, which allows the user to plot one variable against another. It's easier to show than to explain. For instance, you can move the friction circle from the textbook into reality simply by creating an XY graph of lateral against longitudinal g. Figure 4 is one such chart. The top half of the chart is acceleration, the bottom half is braking, and right turns are to the right of the chart.

This chart is fairly satisfactory, showing a driver who is quite prepared to turn in under braking and to accelerate while still under strong lateral g. It is equally easy to create a gear chart (Figure 5) by comparing RPM with Speed to mimic the sort of chart provided by the gearbox manufacturer.

MATHS CHANNELS

The data that we can display is not limited by what is physically collected by the logger. Most software packages have some sort of additional maths facility that can be used to create extra channels of data. So if we logged suspension travel and were interested in the amount of body roll, it would be possible to create an extra channel that took the relative damper movements left and right and used them to calculate the amount of body roll. It all depends on what we see as issues needing attention and how much time we have to create the data that would help our understanding.

As an aside, if two maths channels are created, one for front body roll and one for rear, they will almost invariably ▶

FIGURE 4 USING AN XY CHART TO PROVIDE A FRICTION CIRCLE

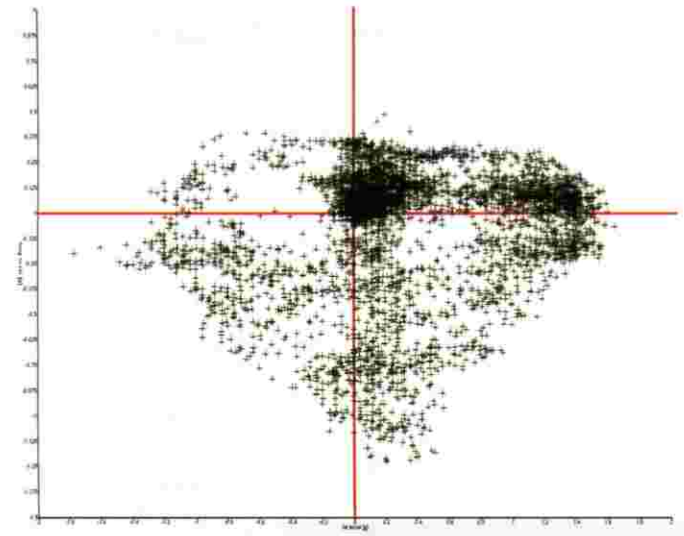

> “You can move the friction circle from the textbook into reality simply by creating an XY graph of lateral against longitudinal g”

FIGURE 5 USING AN XY CHART AS A GEAR CHART

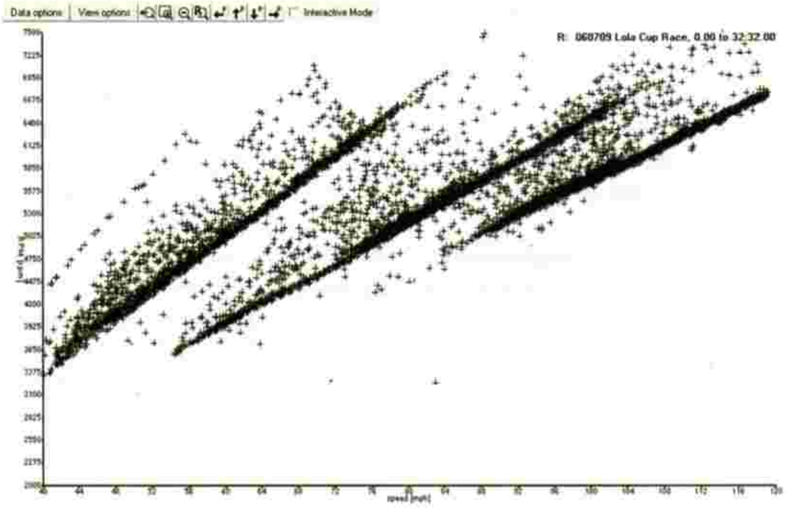

FIGURE 6 SIMPLE MATHS CHANNEL

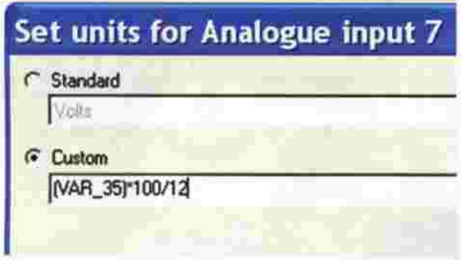

Set units for Analogue input 7

○ Standard

Volts

● Custom

(VAR_35)*100/12

show different rates of roll. One explanation for this is that the chassis is twisting, but the reality is much more likely to be that the car is working one set of tyres much harder than the other and this appears in the data as differential body roll. If suspension adjustments then bring the calculated rates closer together and lap times improve, we have happened upon an heuristic method of sorting the handling.

Most of the time though, maths channels are not that complicated. One very useful one is a gearing channel and this is easily done by calculating the miles per hour per thousand RPM. This then gives a data set that can show the driver what gear was being used at any part of the circuit and shows the engineer what ratios were in the car when the data was taken. The formula would be: MPH/1000rpm = Speed (in mph) ÷ (RPM ÷ 1000).

> “The maths channel will enable us to configure the data into whatever values we want and make good some of the inadequacies of the physical installation”

Maths channels can also be used for calibrating sensors and for putting the raw data into a form that is easy to work with. For example a throttle sensor will merely output a voltage related to position. We can use the maths channel to convert from a voltage to a percentage (in the case of a 12 volt signal) by simply making the variable equal to the: Percentage value = logged value x 100 ÷ 12.

The maths channel will enable us to configure the data into whatever values we want so that it can make good some of the inadequacies of the physical installation. If, for example, our installation does not make full use of the available movement of the sensor, the range of readings captured by the logger might only be from 0 to 9.5V rather than 0 to 12V. Obviously a more careful installation would be the best bet, but the maths channel can be used to restore the full 100% range in the data by modifying the previous formula to: Percentage value = logged value x 100 ÷ 9.5.

Figure 6 shows a screenshot from the Race Technology analysis program that does just this for the DL1. As always with computers, you have to use the right syntax and in this case VAR_35 refers to analogue channel 7 which was the input from the throttle potentiometer. You need also to be careful to use the correct evaluation order. Remember that there are two answers to 2 x 3 + 1, depending on the order in which you do the calculations.

The maths channel will also help adjust the data where we think it is necessary. In the case of steering, the actual angular displacement of the steering wheel will vary for a given cornering force according to the radius of the corner. A fast bend needs much less steering to reach 1.5g than a hairpin would. This then makes the steering traces different when comparing high speed and low speed corners. Some people dislike this inconsistency and use the maths channel to create a 'speed adjusted steering angle' to compensate. The usual formula for this calculation is: Speed Adjusted Steering = Steering x Speed x √Speed.

LOOK-UP TABLES

Another facility provided in some of the software packages is the look-up table. If your software provides this facility, it is not necessary to use a mathematical expression to define the relationship between the signal and its meaning as discussed earlier in relation to the throttle position sensor. Instead you simply measure the values directly and enter them into a look-up table.

This is not too significant in a simple case, but consider a linear potentiometer used to measure steering movement. It might be that you are using only 110mm of the available 150mm of travel and, due to the linkage involved, the movement of the steered wheels might not be in a linear relationship to rack travel. Using a look-up table greatly simplifies the setting up of the sensor.

In this case we can cope easily with the non-linear relationship and also the fact that we want zero steering to be shown at 55mm of travel and we want to have negative and positive values depending on whether the steering is to the left or right. A typical look-up table (this one is taken from the AIM software and wants its inputs as millivolts) is shown as Figure 7.

The benefits of this are obvious, but the down side is that the software still has to interpolate between the values that you enter and does this by establishing a mathematical relationship between the input and the required output. Figure 7 shows the software has created a function that smoothes out our measured data in a way that may not always be what we wanted. The smoothing means that it

FIGURE 7 TYPICAL LOOK-UP TABLE

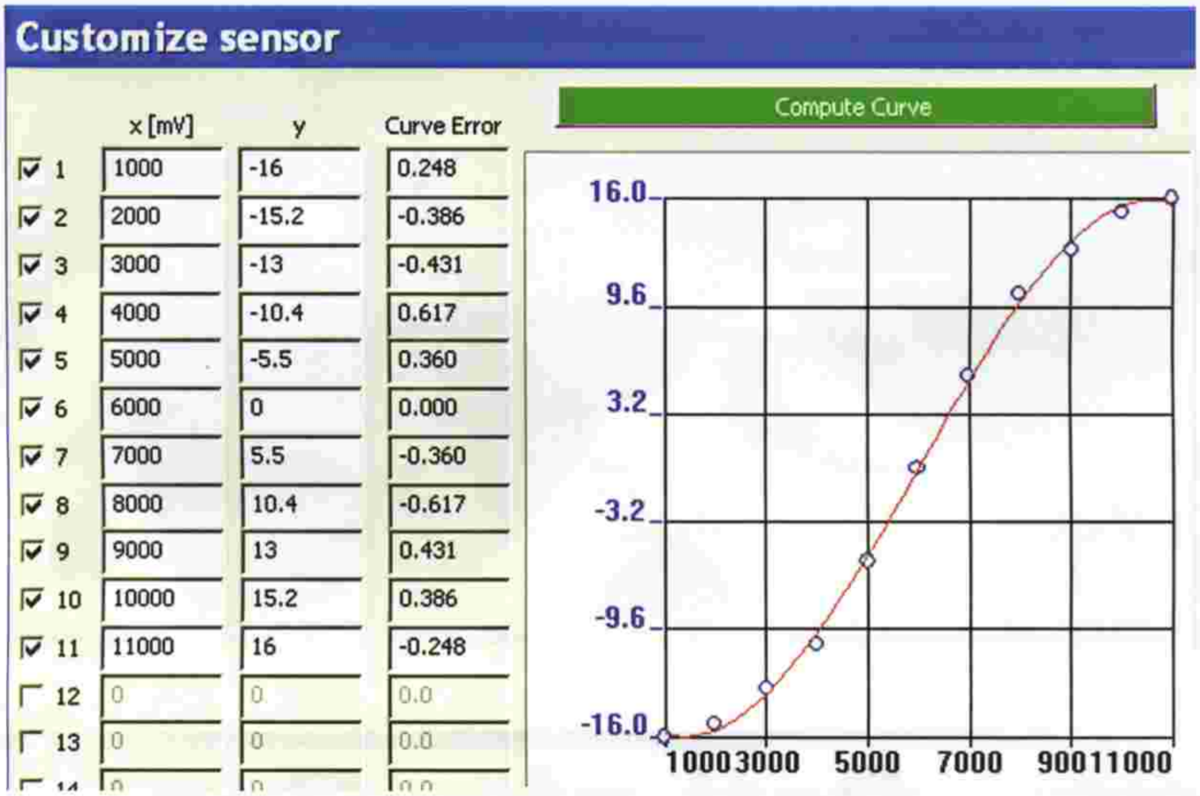

Customize sensor

	x [mV]	y	Curve Error
☑ 1	1000	-16	0.248
☑ 2	2000	-15.2	-0.386
☑ 3	3000	-13	-0.431
☑ 4	4000	-10.4	0.617
☑ 5	5000	-5.5	0.360
☑ 6	6000	0	0.000
☑ 7	7000	5.5	-0.360
☑ 8	8000	10.4	-0.617
☑ 9	9000	13	0.431
☑ 10	10000	15.2	0.386
☑ 11	11000	16	-0.248
☐ 12	0	0	0.0
☐ 13	0	0	0.0

will interpret a signal of 4.25V as 9.4 degrees of left steering when in fact physical measurement put it at 8.8 degrees.

AVOID BUILDING IN INACCURACY

This is not of any consequence if all we are looking for is nice values on the strip chart axis. But if we are going to take this number and create another channel with it there might be problems. For example, we might create a channel that shows actual steering vs theoretically required steering (to assess understeer or oversteer) and we would then be building in an inaccuracy that could stack up with others to compromise the information we provide.

At the entry level, the choice of logger commits you to using the manufacturer's own software but that turns out not to be too much of a problem. The format of the data files written by each system is pretty well unique and few people possess the skills required to translate data from one format into another. Software invariably comes with the ability to export the data into a spreadsheet so if you need to look at the data in some way not provided for by the logger software it is possible.

One shortcoming of the maths channels, generally, is the inability to compare one data point with the others. For example, if your software shows time slip but not the rate of time slip there is no way of comparing the speed on this lap to the speed at this point on the fastest lap.

> “The export facility enables the true data geek to import data into a spreadsheet and work on it in whatever way is required”

Maths channels are just that – if you can formulate an equation, it's OK, but there is no ability to carry out logical tests (eg is this greater than or less than... or use maximum values in your calculation). You can merely operate on a specific data point to provide the data in a new form. This is where the export facility comes into its own because the true data geek can import data into a spreadsheet and work on it in whatever way is required. ■

STAYING THE COURSE

Like the idea of data logging but worried about the hassle involved? Graham Templeman offers practical tips on how to make life easier

IT IS a sad fact that a significant number of data loggers are installed, used for a while and then abandoned. The reasons are often more to do with hardware problems rather than any dissatisfaction with the process of logging.

These are complex systems relying on exposed and delicate sensing devices and lots of external wiring and removable connectors. Small wonder, then, that in a harsh motorsport environment they cease to function.

Many competitors then abandon the whole business of data logging because it just seems like too much trouble. This is a pity, because with proper installation and reasonable care, there should be no problem keeping the thing alive.

THE SYSTEMS UNIT

Manufacturers give ample instructions on how and where to mount the unit but a couple of things are often overlooked. In the case of a modern lightweight logger, the anti-vibration mounts have to be much softer than you think, since the mass of the unit is so low that conventional rubber mounts would just be too stiff. If using bolts, be sure that the bolts really do isolate the logger from vibration (see Figure 1).

One solution that is often offered is that of Velcro. It meets the anti-vibration criterion (at least for a lightweight unit) but if it is used for accelerometers it can cause a calibration problem. Leaning on the unit shifts the Velcro, changes the installed angle of the device and influences the readings that it will give.

In cars with open cockpits, rain can be a nuisance. One cheap and effective solution is to mount the systems unit in an ordinary plastic lunch box. The locking lid keeps out the rain and the box itself can be isolated from vibration. It causes great hilarity in the pit lane – 'Ho, ho. Will you be stopping for lunch?' – but it gets the job done.

FIGURE 1 ANTI-VIBRATION MOUNTS

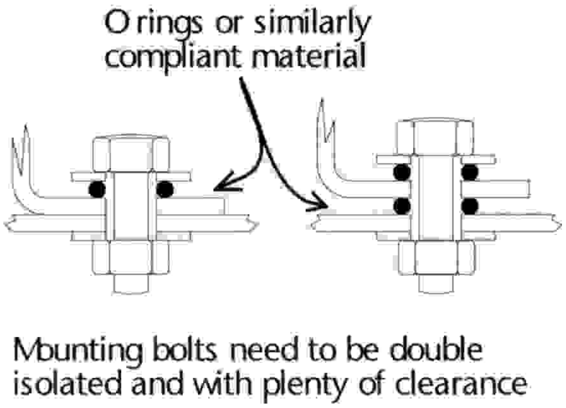

BELOW This shows the spare cable that seems inevitable with pre-fab connectors

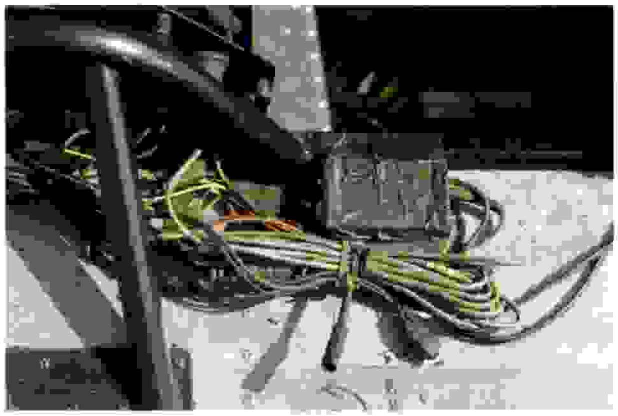

WIRING

The cabling of the unit also needs attention. There is a lot of electrical noise in the environment and care should be taken to minimise its effects. Runs of cables, tied neatly together, may look great but there is a danger of electrical noise creating spurious signals.

Although it doesn't look as good, let each cable take its own route and only meet the others close to the logger. Using shielded twisted-pair cable will minimise the noise effects although it has to be said at least one logger manufacturer uses ordinary 5 Amp household flex – brown for the feed, blue for the signal and yellow/green for 0 Volts.

When it comes to connection methods, manufacturers are split between plug-in sensors and terminal blocks. For the sort of home-brewed system discussed below, terminal blocks seem to win hands-down.

Both have advantages and disadvantages. Terminals are user-friendly if you are developing and experimenting but might lead you into having to create an external terminal block to tidy things up. The problem is that the logger will provide just one terminal for the system voltage and you end up trying to fit several wires into the same point.

Systems which use pre-fabricated connector cables with their own plug and socket systems avoid this problem but can be messy. The cables never seem to be the right length and you end up with coils of over-generously provided cable hidden all around the car.

An additional problem with plugs and sockets can be self-inflicted by clumsy handling that can push the pins out of register. This can be difficult to trace. Just because your manufacturer uses plugs, doesn't mean that you can't go down the add-your-own sensor route. The plugs are proprietary items and a few minutes on the appropriate electronics supplier's website will generally find you what you need. In the UK, companies like Farnell, Rapid Electronics or RS all provide detailed catalogues on their websites and are happy to deal in small quantities.

SENSORS

Many of the sensors are vulnerable to accidental damage so it seems like a good idea to protect them as much as possible. The first essential is to keep the sensor out of the way of potential physical harm. Anything mounted on to suspension is inevitably susceptible to flying debris and to accident damage. Not only are the sensors at risk, so is the wiring and it also needs to be protected.

Vibration and misalignment hurt sensors. For example, a throttle position sensor (TPS) needs careful set up to ensure that it is not taken beyond its mechanical range. A typical rotary TPS has a travel of 120 degrees to measure 90 degrees of movement so it needs to be carefully positioned in the middle of its range of travel. Damage occurs easily if the pot is forced to travel further than its physical limits.

Vibration is inevitable and invariably shortens lifespan. Wherever possible, some comfort should be provided through linkages that allow for a degree of misalignment and vibration absorption. For example, a rotary TPS can be connected using a short length of small-bore rubber hose as a flexible coupling.

The range of dedicated sensors is enormous and is aimed at professional teams or those with deep pockets. For instance, all sorts of technologies are available for sensing position.

There are LVDTs (linear voltage displacement transducers) which are highly accurate but require the signal to be processed further before it can be used. There are also non-contact Hall effect devices that measure either rotary or linear movement and are capable of millions of cycles before failure. You can also buy beautiful linear potentiometers that come in virtually any stroke length and are equipped with exquisite miniature rod end joints to make installation simple.

> “With proper installation and reasonable care, there should be no problem keeping the thing alive”

As this is the Practical Racer section, we ought to look at ways of making our own. The cheapest and most popular method to measure displacement is by using a variable resistor configured as a potential divider. A supply voltage (12 or 5 volts depending on your logging system) is fed into one end, the other end is connected to ground and the signal is available by measuring the voltage between the ▸

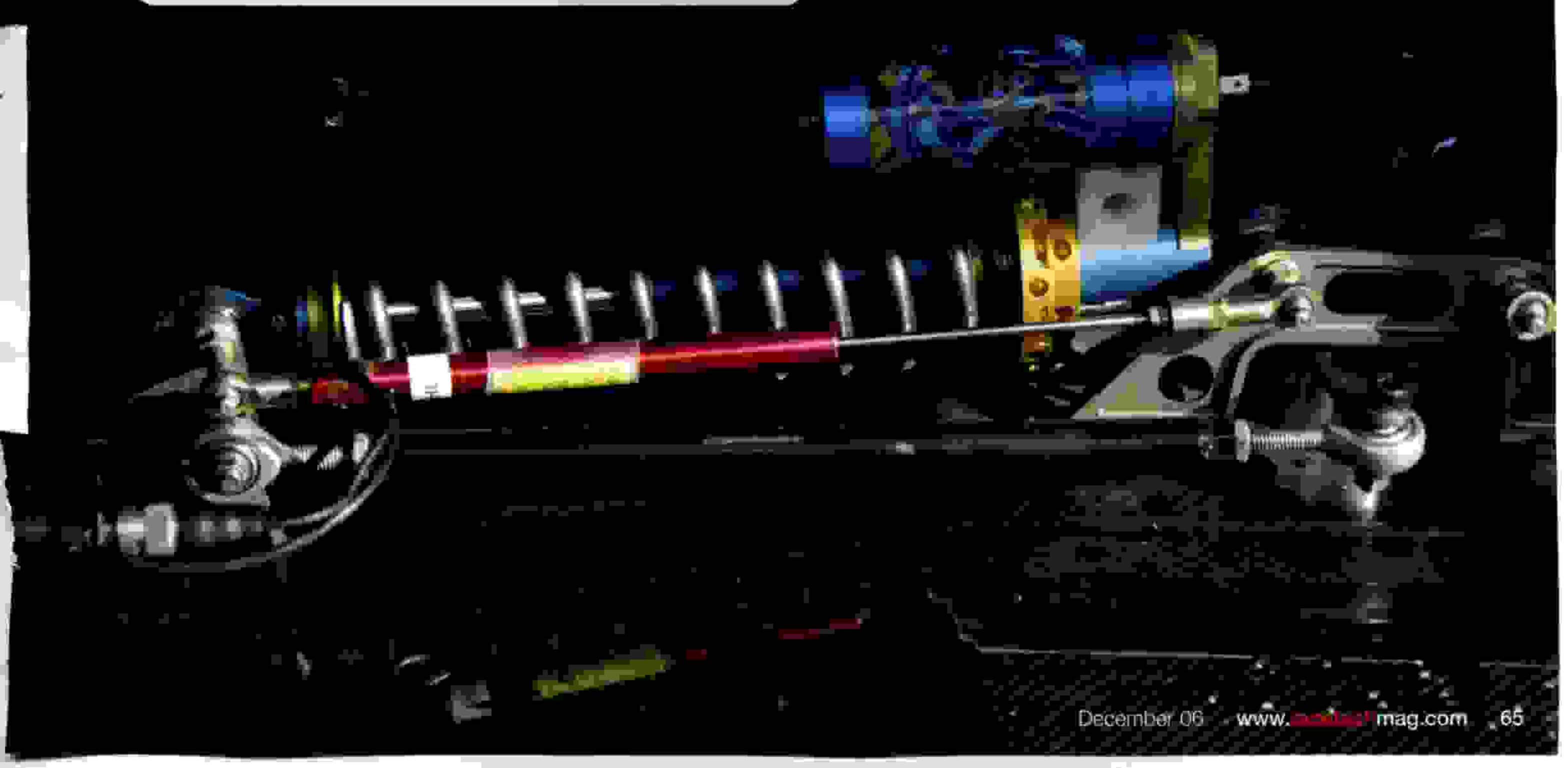

BELOW This is the professional way, as demonstrated by a displacement transducer on the rear suspension of a DTM car – accurate, easy to fit and, if your mechanics earn a wage, not too expensive

LEFT Cheap, high-quality rotary units are available at a breakers yard near you

BOTTOM A flexible coupling and plenty of adjustment mean that this recycled TPS can be properly aligned and should last for more than just a few races

third (wiper) terminal and ground.

Finding out how the system is wired is straightforward enough. If your system uses pre-fab cables, there will be a three pin plug and a multimeter can safely be used to identify the power supply pin and the 0 Volts pin. When the meter hits the correct pair of pins, the supply voltage will show on the meter. If it has a positive sign, the meter's leads are connected the right way round and the red lead is on the supply pin. The signal line will be the one that's left and will ultimately need to be connected to the wiper of your potentiometer.

If your system uses screw-in terminals, you are looking for a power supply (+12 V or +5 V), a ground (0 V) and a signal terminal. They should be clearly marked.

> “One cheap and effective solution is to mount the systems unit in an ordinary plastic lunch box”

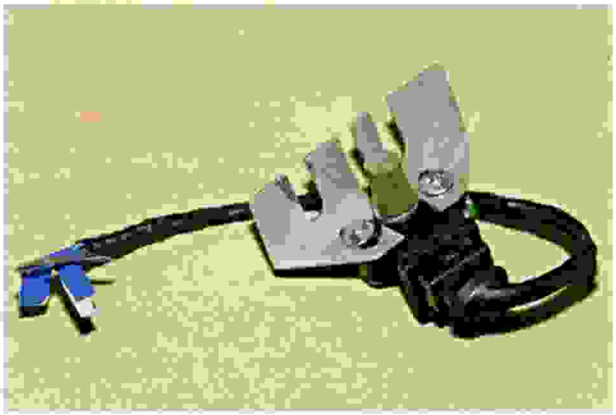

Variable resistors come in all sorts of shapes and sizes, so you should be able to find pretty much what you want. An obvious starting point is the throttle position sensor found on just about any post-1990 fuel injected road car. They are top quality (expected to last the life of the car without trouble) and readily available.

A recent visit to my local car breakers yard provided me with four for less than a fiver. The photograph shows the collection that I turned up in half an hour of ferreting around. The haul offered enough variety of shapes and sizes to fit any installation.

If this doesn't appeal, the usual electronic supply houses will provide good quality variable resistors (potentiometers) at reasonable prices. Avoid the wire-wound type – they lack accuracy – and focus on the resistive film variety.

Slide type potentiometers are available but are not really suitable since their mechanical construction is not robust enough for motorsport. Rotary potentiometers seem to be the better bet. They are usually available as 90 and 270 degrees or three or 10 turns. The level of resistance is not too important but 5000 Ohms is okay.

USING YOUR CAR'S OWN SENSORS

The manufacturers are understandably cagey about the use of sensors that don't come out of their own catalogue. They market their own sensors because doing so reduces the chance of accidental electrical damage to the logger and enables them to make provision in the software for easy and accurate calibration. Just what the customer wants. ▶

Sometimes, though, for reasons of expense and convenience the end user might want to go it alone. A good example would be a temperature sender that once drove a now-defunct dashboard gauge but now becomes a very tempting sensor for a logging system.

Borrowing sensors in this way is okay, but be careful when sharing sensors between two or more systems. You are likely to find that the host system (especially in the case of an ECU) already has a second resistor built in so that you can shortcut much of what follows about deciding on the value of a second resistor and fitting it. It will simply be a matter of connecting the signal voltage that the sensor provides, directly to the logger.

Whatever happens, you need to be wary of interactions between the two systems and be convinced that allowing the logger to freeload off the input for another device will not affect the safe and efficient operation of the host equipment. That's easy to test if you are sharing a temperature sensor between the logger and the dashboard, but it needs a bit more care if you are borrowing from the anti-lock braking system or fly-by-wire throttle.

RESISTIVE SENSORS

Temperature and pressure sensors are normally resistive. That is, their resistance varies with changes in whatever is being measured and calibrating this type of sensor is a bit of a fiddle. Their variable resistance acts on the current in the circuit, not the voltage. But we need to provide the logger with a variable voltage, which we can do by using the same potential divider concept as for measuring displacement.

In this case, instead of having a variable resistor that changes as it is moved mechanically, we can add a second resistor into the system so that we now have a home made potential divider that can provide us with the necessary signal. Figure 2 shows this using a thermistor in conjunction with a second resistor.

FIGURE 2 THERMISTOR AND RESISTOR

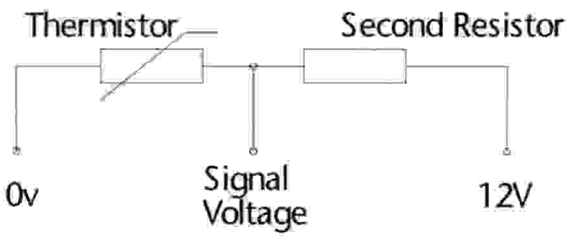

In practical terms, the starting point is measuring the resistance of the sensor over its working range. The first two columns in the fragment of an Excel spreadsheet (top right) show the data that was collected using a multimeter, a thermometer and a pan of hot water. The third column was created using the formula below to calculate what the output voltage would be, assuming a second resistor value of 1000 Ohms.

$$\text{Voltage Out} = \text{Voltage In} \times \frac{R2}{R1 + R2}$$

The value chosen for the second resistor will depend on the measured value of the first. The figures in the fragment of the spreadsheet (Figure 3) show the working range to be from 270 to 1790 Ohms, so a resistor in this range will give a reasonable signal. The 1000 Ohms used in this case gives a range of about 5 Volts out of 12 and this would give an adequate resolution on even an 8 bit system.

If your software has lookup tables, the job is almost complete.

LEFT This professional Penny and Giles linear displacement transducer has the added benefit of being waterproof

FIGURE 3 SPREADSHEET FRAGMENT

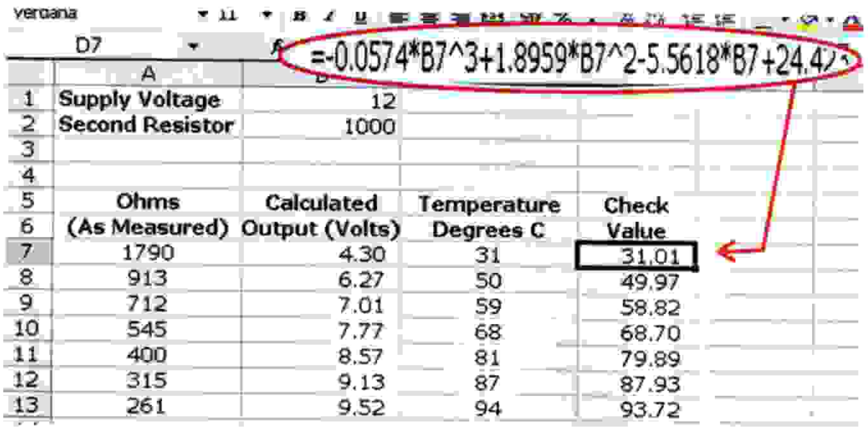

```
=-0.0574*B7^3+1.8959*B7^2-5.5618*B7+24.4
```

	A	B	C	D
1	Supply Voltage	12		
2	Second Resistor	1000		
3				
4				
5	Ohms	Calculated	Temperature	Check
6	(As Measured)	Output (Volts)	Degrees C	Value
7	1790	4.30	31	31.01
8	913	6.27	50	49.97
9	712	7.01	59	58.82
10	545	7.77	68	68.70
11	400	8.57	81	79.89
12	315	9.13	87	87.93
13	261	9.52	94	93.72

You simply need to wire in a second resistor of an appropriate value and arrange a feed at the appropriate voltage in order to be able to check the signal output over the operating range. These can then be entered into a look-up table for the software to read and your system will translate the output from Volts into degrees.

Without a look-up table, there are a couple more steps. Excel is used to create a mathematical function that the software can read. The process is to draw a scatter diagram of the recorded values (the blue data points on the chart in Figure 4) and to use Excel to show the trend line.

This trend line has several formatting options, one of which is to print out the mathematical expression that describes the shape of the line on the graph. This is highlighted in the red ellipse on the chart.

> “Sometimes, for reasons of expense and convenience, the end user might want to go it alone”

This expression can then be typed into the software to describe the behaviour of the sensor. This goes a long way to overcoming the problem of not having the look-up table facility.

In the case of the graph shown, the expression is:

$y = -0.0574x^3 + 1.8959x^2 - 5.5611x + 24.423$

This can be seen in the edit line of the spreadsheet where an extra check column was created.

For the technically-minded, a third order polynomial expression defines the line best. For those of you that are not Excel fans, e-mail me (gtempleman@gmail.com) for a copy of the spreadsheet and more detailed instructions.

FIGURE 4 VOLTAGE v TEMPERATURE

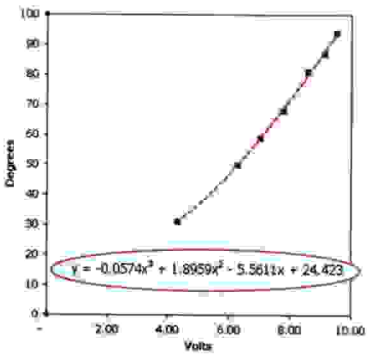

All that remains now is to use this hard-won formula in the logger software. You can now create a user-defined variable, measuring degrees Celsius, where the temperature is calculated as:

-0.06* Voltage cubed + 1.9* Voltage Squared - 5.56* Voltage +24.4

The rounding off is permissible, given that we do not need enormous precision in this case.

If you don't have look-up tables, this fairly drawn-out process is important because it provides a method for calibrating any type of sensor. You can measure the variable that interests you and the output that your home-brewed sensor provides and use Excel to identify a maths function that best describes the relationship that you can use to create a user-defined variable.

For example, it provides a method to calibrate steering in degrees left and right through a process of taking voltage readings every one or two degrees of steering displacement and creating a mathematical expression in exactly the same way.

Armed with a decent set of data for steering, we could be well on the way to measuring tyre slip angles v grip to feed our simulation program better information or move on to further insights about understeer or oversteer. Stay tuned. ■

THE PROFESSIONALS!

The big boys have more time and more toys than the weekend warriors to help them understand what makes a car tick. But the basic principles of data logging remain the same. By Graham Templeman

SO FAR in this series, we have been looking at how to make the most of an entry-level data logging system which might cost not much more than a couple of sets of tyres. But what happens when the system costs the same as a small motor home and has a full-time expert to look after it?

To find out how different things can be, we spent some time with two professional teams. As you might expect, things are both very similar and very different.

First, Lee Penn of LNT Motorsport and Albert Lau and Jean-Patrice Loof of West Surrey Racing deserve a huge thank-you for being so helpful in providing the material for this feature. Lee let me watch and ask innumerable questions during a test session at Silverstone and Albert and J-P gave me a really thorough run through the data engineering of an A1GP car at the WSR headquarters.

Starting with the similarities, amateurs and professionals are all trying to minimise lap times by understanding and reacting to what the car is doing. The difference comes in the resources available to them. A well-funded professional team will log as many parameters as it is interested in and will have the manpower to install and maintain the system and to manage the resultant data.

The number of channels is not an issue – if you need more, you use an expansion box – and what seem like astronomical prices for sensors to the clubman become rational choices for a professional team. Laser ride height sensors fitted each corner would cost about £600 per corner range and the choice is made on the basis of whether the extra information will help make the car go faster.

In this case, knowing the actual dynamic ride height and the damper movement allows the engineer to calculate how much the tyre deflects under load. This can be measured as either a distance or a frequency. If there are strain gauges fitted to dampers, this data can be combined with the tyre deflection to calculate the actual tyre spring rate. If the car is being run in a category where the choice of rubber is free, then the team can talk with some confidence to the tyre engineers about the appropriate tyre specification for this car.

Apart from the inevitable budgetary constraints, professional teams also have to work within the rules imposed by the series organisers. British GTs, for instance, are not allowed to use telemetry (the transmission of data back to the pits while the car is running) and A1GP only allows telemetry at 1 Hz. Other series specify the type of logger to be used and impose further constraints by limiting the available memory.

> “Joe Clubman tries hard but just how careful the teams are to manage the data properly was a revelation”

In the eyes of a club racer, the range of sensors in regular use is enormous. In addition to the usual

RIGHT West Surrey Racing's meticulous approach to gathering and storing information is legendary on the pit lane. It means that when data is compared from its A1GP cars, for instance, one car can be superimposed upon another with minimum fuss and in the knowledge that the figures are valid

FIGURE 1 VITAL SIGNS

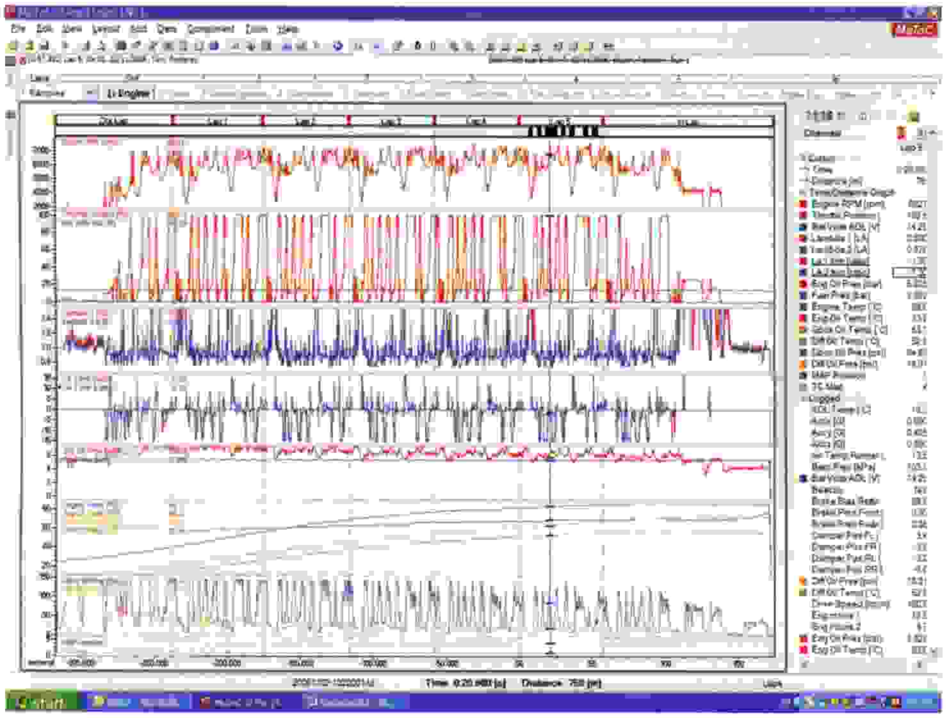

LEFT The first step is for the data to be dumped onto a couple of screens in order for the vital signs to be checked. This screen includes engine and gearbox pressures and temperatures and lambda measurements from each bank of cylinders

suspects, the speeds of all four wheels, brake line pressures, damper travel and pushrod or damper load sensing all seem pretty much routine and four-corner laser ride height sensing, infra red brake temperature monitoring and three axis accelerometers are also used. There did not seem to be much interest in the direct measurement of yaw.

WSR's A1GP uses a pitot tube to measure airspeed since this can vary significantly from road speed and can help to explain anomalies in the data. Sepang has two long straights that run at almost 180 degrees to each other and airspeed and downforce can vary significantly between the two. The effects of a tow from another car can be seen very clearly because the pitot tube data fluctuates wildly.

Logging rates are what you might expect with most parameters being logged at 20 to 50 samples per second. The less important data, like battery voltage, is logged at much lower rates. Data relating to suspension movement is logged at about 200 Hz. Although faster sampling rates are possible, 200 Hz works and avoids the need to filter the noisier data that faster rates would provide.

Having access to information about damper travel means that the professionals are less interested in the actual value of lateral g generated by the car. The assumption is that a professional driver is capable of getting the best from the car and there is very little soul-searching about just how fast the car is cornering. Lateral g might merit a quick check to see that values are within the expected range, but very little more.

Of more interest is the attitude of the car as indicated by damper movement. This gives insights into the extent to which the car pitches, heaves and rolls and the effects that this is having on handling.

Downloading the data and having at least a quick look at it after every run is seen as imperative. A starting point is to dump all of the data on to one or two screens just to check that all the sensors are operational and giving reasonable data. Then follows a quick look at the drive train's vital signs – revs, temperatures and pressures – just to give a swift overview of trends over the run.

Pressures may be expected to fall slightly as temperatures increase and an eye is kept on fuel pressure because drops in pressure can be indicative of fuel surge. Figure 1 shows one such screen that includes engine and gearbox pressures and temperatures and lambda measurements from each bank of cylinders.

The needs of the driver come next. Both teams used a driver activity screen showing the usual things; the only difference at this level is that there is a routine monitoring of brake line pressure. This gives an indication of the level of braking asked for by the driver and if front and rear circuits are both measured, even more information can be extracted.

A maths channel showing the relationship between front and rear pressures is an immediate and unarguable record of brake balance. The brake balance should show a consistent percentage but occasionally the trace deviates from a straight line when the pedal is pressed or released. This is strong evidence of some sort of stiction in the balance bar linkage.

Figure 2 shows a typical driver comparison screen and the shaded area on the left shows the loss of time through turn 1 due to an over-ambitious entry ▶

FIGURE 2 DRIVER COMPARISON SCREEN

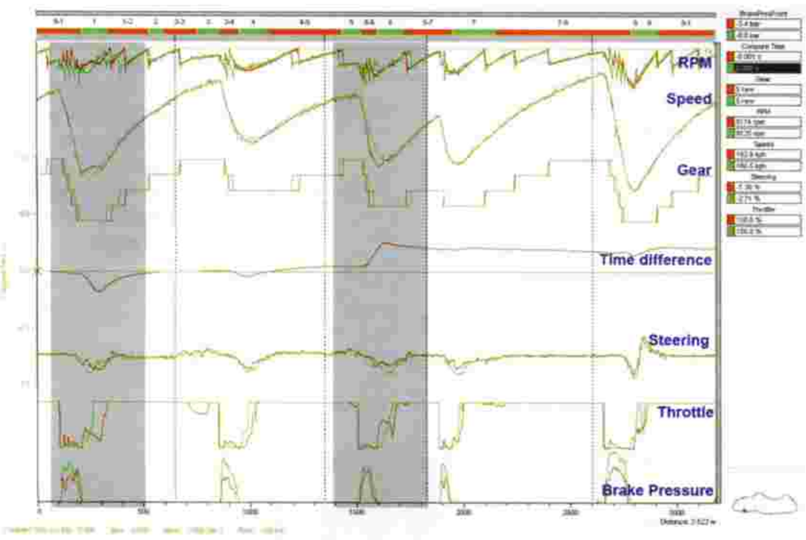

LEFT The shaded area on the left shows the loss of time through turn 1 due to an over-ambitious entry speed. The right hand shading indicates time gained by firm braking and an orderly progress through turn 6

speed. The right hand shading indicates time gained by firm braking and an orderly progress through turn 6.

There is a presumption that the driver will be getting the best out of the car and that all that is required is to identify which technique yields the best results as shown on the time comparison trace.

Figure 3 shows a fragment of a brake pressure trace where the red line is the professional driver and the blue line belongs to a club championship winner having his prize drive with a first taste of carbon brakes and high downforce. The lack of belief in the system's power is obvious both in terms of braking early and gently. By the end of the day, the traces were much more comparable.

One aspect at the pro level that came as a surprise was the inter-relationship between the dash logger and the engine ECU. Dash loggers are designed to take advantage of the ECU's sensor data but communication is very much two-way. For example, when ▶

FIGURE 3 BRAKE PRESSURE TRACE

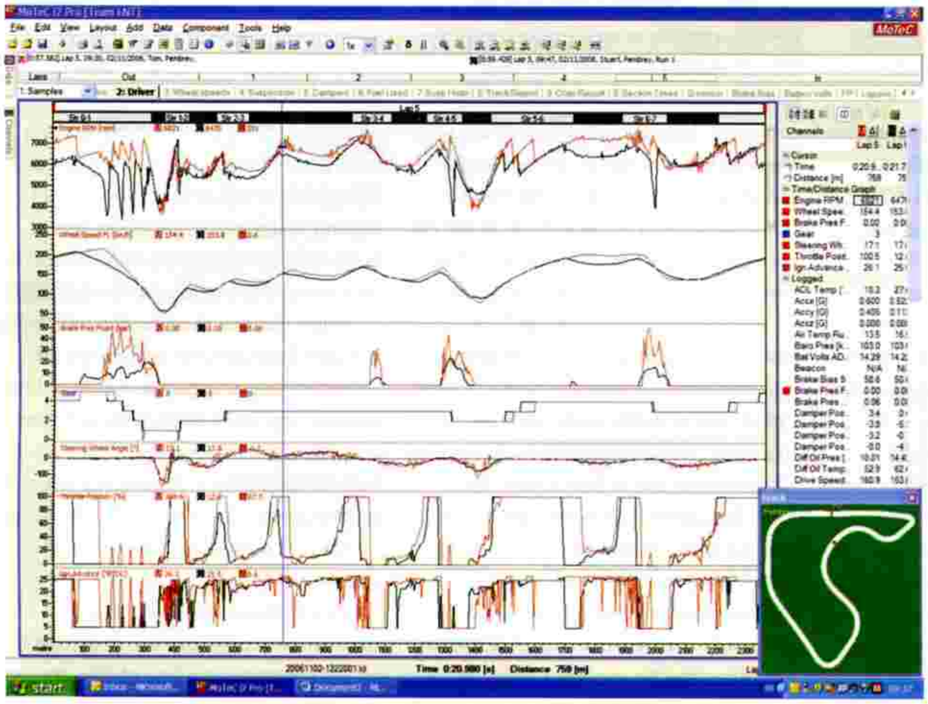

LEFT The red line is the professional driver, the blue a club championship winner having his prize drive. The latter's first experience of carbon brakes and high downforce is reflected in his braking early and gently

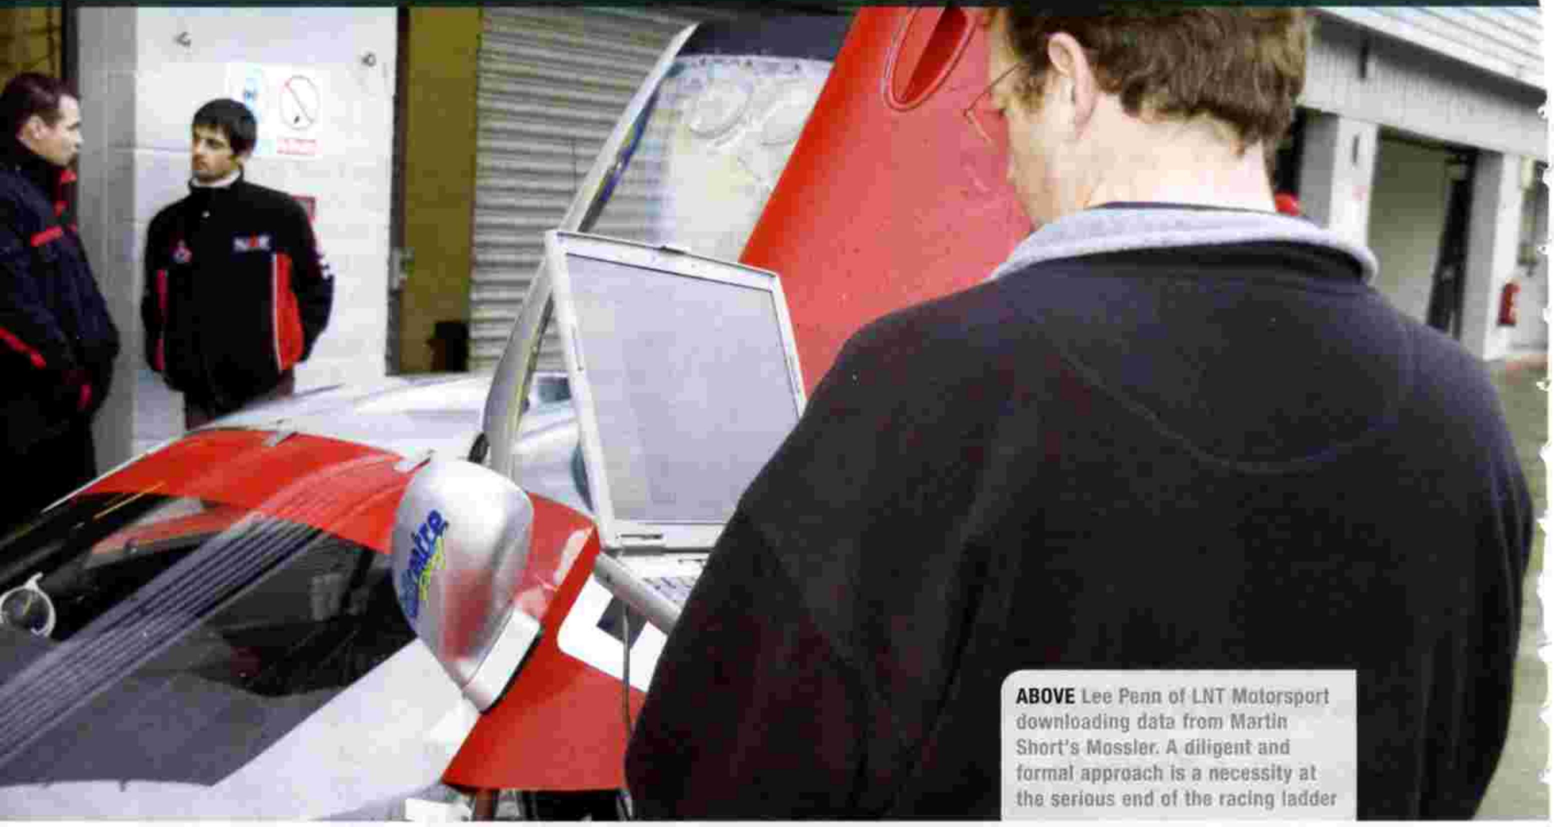

ABOVE Lee Penn of LNT Motorsport downloading data from Martin Short's Mossler. A diligent and formal approach is a necessity at the serious end of the racing ladder

four wheel speeds are logged, the percentage speed difference front to rear can be calculated as a logger maths channel and the dash can then send signals to the ECU to use some strategy for traction control. To do this, the ECU can cut sparks or fuel or can retard the ignition slightly to reduce torque.

Working with a client of LNT Motorsport, one of the items on Penn's agenda in the morning's test was to refine a strategy for traction control. When an engine is mapped on the dyno, it is possible to collect data about the extent to which retarding the ignition reduces the torque output. The data for various engine speeds and degrees of retardation provides the raw material for the traction control mapping.

While there is moderate wheelspin, the ignition is retarded and torque can be reduced by a known percentage. If the wheelspin gets severe, the power is cut completely. The challenge was to create a strategy that assisted the driver through the turns and rescued him in the event of severe changes in grip. A traction control map that worked fine in the dry had proved far too sudden in the wet.

Another revelation was just how careful the teams are to manage the data properly. Whereas Joe Clubman tries hard to be sensible and rational, teams absolutely need to have a formal and properly-documented approach. They have to use well-designed and well-understood protocols for collecting and storing data.

At WSR, with two-car teams in both BTCC and A1GP, there are carefully laid down data conventions for things like how the driver is identified (first name only) the car (by the chassis number) and the file structure (year, circuit name, date, driver).

Channels are given standardised names (Load FL, Load FR) and measurement conventions are set (downforce is positive, load is measured in kgF). These standards then apply to both the cars. The team can then not only superimpose one lap upon another, it can superimpose one car upon another with minimum fuss and know that the data is valid.

The set up of the data logger sensors becomes part of the formal operational processes. The mechanics set up the car before each run with a given amount of fuel and driver ballast. While the car is in this state all accelerometers, damper pots and strain gauges are also zeroed, as is the steering with the steering wheel pointing straight ahead. It is by this sort of attention to detail that championships are won.

The logger is even used to provide the data for the management of the components. Engine run time is recorded as a matter of course and the gear histogram (that shows the percentage of time spent in each gear) can be used to life individual ratios.

Many club racers will be familiar with the routine of draining the fuel tank before an event, putting in a known quantity prior to qualification, draining again and calculating fuel consumption from the amount consumed. This is then used for deciding on the race fuel load with fudge factors to allow for lower consumption on warm up laps and the possible intervention of the safety car.

The professionals use logged fuel injection data to calculate fuel consumption. They create a maths channel that combines engine activity, the flow rate of the injectors and the injection duty cycle to calculate the actual amount of fuel consumed. The dash can display another maths channel calculation, this time fuel load at start minus fuel consumed, to indicate the amount of fuel remaining in the tank.

The canny engineer can also relate data showing fluctuating fuel pressure to amount of fuel left in the tank and be able to cope with, or take steps to cure, any fuel surge problems that present themselves.

Having data for all four wheel speeds gives an insight into braking and traction problems. Club racers tend to rely on a single sensor on the outside front wheel and any lock up can affect the distance calculations and hence the presentation

FIGURE 4 DAMPER TRAVEL ON AN A1GP CAR

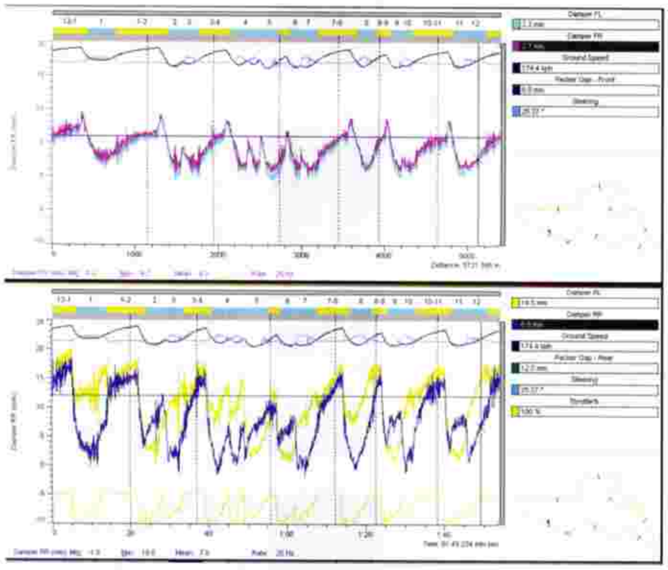

LEFT The straight line at 6 mm on the top trace shows the stage at which the car goes on to the bump stops; above that line, the bump stops are being compressed. As the road speed increases, the dampers are being compressed by the aero load and as the driver steers through the long right-hand Turn 1 the speed and the downforce decrease. There is a small amount of roll, but only to the extent that the outer wheel compresses the damper 1 mm more than the inner. The lower set of traces shows that there is much more movement at the rear but still limited to less than 10 mm

of the data. Logging two front wheels allows the speed calculation to be done on the maximum of these two, i.e. automatically compensating for locked wheels.

With four wheel speed sensors, locking up under braking can relate either to an incorrect balance (one end locking before the other) or driver technique (braking while turning in). Poor gearchange technique is visible with rear wheels locking in downshifts or spinning in upshifts.

A professional team is also likely to try to find the time for creating an aero map for the car. This process involves driving the car in a straight line at steady speeds and with differing aero settings. This is yet another opportunity to spend money because piezo-electric pressure sensors are available that measure the relatively low numbers associated with airflow over surfaces. By measuring the load on the dampers or pushrods, the downforce and the centre of pressure can be calculated and by measuring pressures the interactions between front and rear wings and tunnels can be explored.

> “Engine run time is recorded as a matter of course and the gear histogram can be used to life individual ratios”

The effect of high levels of downforce is obvious from the damper trace. Figure 4 shows front and rear damper travel on the A1GP car and the aero load is almost enough to wipe out body roll completely.

Looking at the top trace, the straight line at 6 mm shows the stage at which the car goes on to the bump stops; above that line, the bump stops are being compressed. As the road speed increases, the dampers are being compressed by the aero load and as the driver steers through the long right-hand Turn 1 the speed and the downforce decrease.

What does not happen, though, is that the left front compresses and the right front extends. Without downforce, you would expect to see the traces heading in different directions as the body rolled. There is a small amount of roll, but only to the extent that the outer wheel compresses the damper 1 mm more than the inner. Only on three occasions during the lap does the inside wheel manage to extend the damper, and then only by very small amounts.

The lower set of traces shows that there is much more movement at the rear but still limited to less than 10 mm. No wonder the top teams own damper dynamometers. Knowing the damper travel means that it is possible to calculate damper speed and to prepare histograms of the percentage of time spent at each speed. This then gives an idea of the relative importance of high speed and low speed damper setting and the extent to which a reasonable balance has been achieved.

Again, you have to envy the professional teams with all the extra toys that they have to play with and their ability to measure what is happening.

Overall, that sums up the pro/am difference. More time, more toys and more chance to understand what makes a car tick. The principles – an emphasis on accuracy and a methodical approach – remain the same.

One sound bit of advice sticks in the memory: people try to read too much into the data and not realise that it is just one part of a complicated bigger picture. The driver, the team and many other sources of information need to be consulted before thinking that you know all the answers. ■

THE COCKPIT COACH

Data logging can play a key role in driver development. Graham Templeman finds out how from a team working hard to develop the stars of tomorrow

A TRADITIONAL comment made about motor racing is that it is one of the few sports where the competitor doesn't have the benefits of proper coaching.

Some of us will remember amazing photos of Jim Russell standing at the apex of a bend indicating the clipping point to students driving past at apparently racing speeds. But we also knew that, publicity shots apart, once the car was on-track, the driver was on his own.

Not any more. With data logging, the effects of every action taken by the driver can be examined in detail later on the screen of a laptop.

To find out how data logging contributes to driver development, we talked to Mike Rose, team principal of PR Motorsport. Mike has been running young drivers (14 to 17-year-olds) in T-Cars for the past six years and he's obviously very good at it. Both of his 2006 drivers, Adrian Quaife-Hobbs and Ollie Webb, have qualified for lucrative Formula BMW Scholarship awards for 2007.

In case it has passed you by, T-Cars have a space frame chassis with composite touring car bodywork and are powered by sealed Mountune-prepared two-litre tuned Ford Duratec engines. They run on street-legal Avon ACB10 tyres and use their 145 bhp to propel them to a top speed of over 120 mph. Lucky 14-year-olds.

The data logging is controlled by the regulations, and a Stack system is required that logs revs at 20Hz, left front wheel speed and throttle position at 10Hz, and oil and water temperature and oil pressure at 1Hz. Oddly, lateral g is also required at 1Hz. Brake line pressure and fuel pressure are optionally measured at 10Hz. The use of an expansion module to go beyond the basic eight channels is permitted and the must-have here is steering wheel position.

There are two unexpected gems hidden in the T-car regulations. The first is that testing by the under 16s has to be carried out under supervision. These are powerful two-seaters driven by very young drivers and so the second seat is occupied by a driver coach. This is great for training and safety, but from the viewpoint of the race engineer there is up to 100 kg of ballast on one side of the car's centre line. This assault on the corner-weights makes detailed data

RIGHT This screenshot shows the sort of information that is made available to each competitor at the end of a session. The time line in the blue panel at the bottom of the screen shows the time differences between the two laps and provides food for thought for the team and the driver

LEFT Featuring 145 bhp and capable of 120 mph, T-Cars are playing an important role in developing young talent. The driver's every action can be examined in detail later

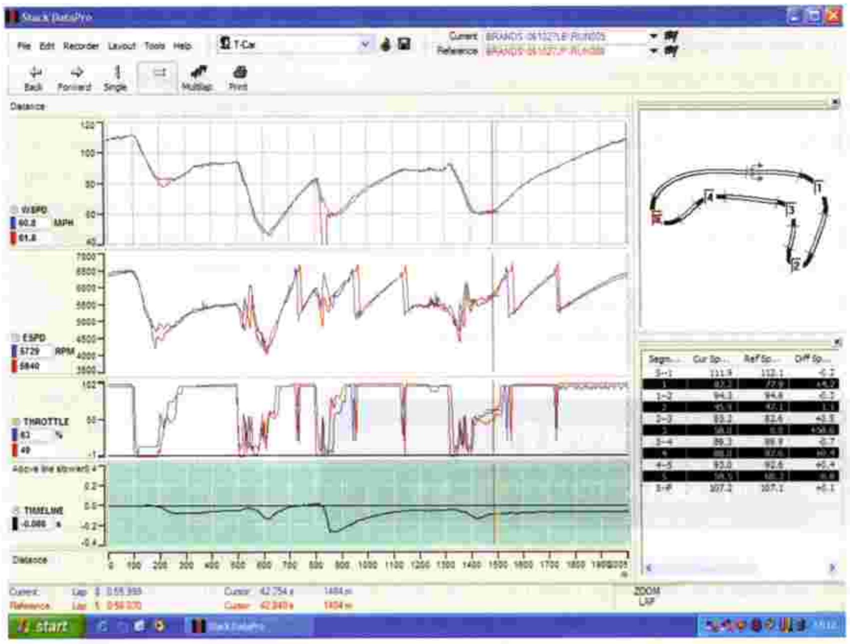

interpretation difficult because conditions are so different between testing and competition.

The other unexpected item is that the regulations require that all recorded data must be made available to the Championship Co-ordinator and Eligibility Scrutineer. After every official session, each competitor is provided with a printout that compares his or her time to the fastest in the session. The person putting up fastest time gets to see his own data compared with the data of the second fastest. This must be one of the few real cases of freedom of information.

As a training formula, enabling competitors to compare their data with the fastest is a great step along the way to seeing how things should be done. Whilst this is a real benefit for someone who is off the pace, it naturally rankles with the front runners like PR Motorsport when they are forced to disclose where they are making time. Understandably, it can be a bone of contention.

PR Motorsport's tactical response to this, bearing in mind that all *recorded* data is to be made available, is to log only the minimum data required by the rules during official events. This means that the teams do not log things such as brake pressures and steering inputs during competitive running.

One positive side issue has been the ability to refute allegations of cheating. The data shows without argument that the power levels are legal with acceleration figures that are consistent across all the cars and that the speed comes from the usual large number of small improvements that, individually, make little difference but all add together to make real gains.

The ethical issue comes with a corner such as Turn 8 at Knockhill (Taylor's Bend) where an incorrect approach speed can cost not tenths but whole seconds on the long straight back to Turn 1. The sharing of data immediately negates the advantages that come from talent intelligently applied.

Ethics apart, how does a top team bring on its young charges? Firstly, there are very few absolute beginners, with most youngsters having had some karting experience. Apparently, karting does not breed a single style of driver. As in any form of motorsport, each driver is different and ex-karters might be aggressive or gentle with the machinery and brave or careful in their driving.

They have enough experience to know what they want in the way of vehicle set-ups. The quality of the feedback also varies and the team has to be careful to check the feedback against the data. It has come across the schoolboy error confusing understeer with oversteer – well the car ran wide and the driver had to over steer to get it back on line.

A lot of this sheer inexperience is revealed by the data. There is a tendency for some new drivers either to be 'over the top' or slightly overawed by their new surroundings. A screen that shows speeds, revs and time slip against a team-mate is a simple and unarguable way of demonstrating where mistakes are being made and where time can be found. ▶

> “ Linking video footage to data develops the skill of being able to 'see' the data as representing a particular part of the track ”

This is one reason why PR Motorsport tries to keep the set ups similar for the two cars that it runs. It helps in the very early stages of a driver's career although as they gain more experience, it is possible for different set ups to evolve to suit driver preference. This in itself can be a positive, because the benefits of two different approaches reveal themselves in the data. Mike has found that there his drivers have always been prepared to share data knowing that it will help both themselves and the team.

From the team's viewpoint, the first stage of data interpretation is to look at the car's vital signs – how are the pressures and temperatures behaving? Gear ratios are fixed, so engine revs becomes a driving rather than a car preparation issue. Once satisfied that the car is okay, attention can turn to driver performance. As always, the basics of Speed, Revs and Time Slip reveal much of what is going on.

> “The approach of chasing the time slip is inevitably backward-looking”

An important first step is to use the data to identify simple driving faults, if any. For most 14-year-olds simply making quick clean gear changes is an art that still needs to be mastered. Showing time on the X axis allows the data engineer to measure the blip in the rpm trace and decide whether a young driver is managing to change in the requisite three or four tenths of a second. Difficulty in gear changing shows up as the engine revs going all over the place as the driver tries to find a cog.

Braking faults show up as well. The long g trace not only shows the maximum deceleration under braking but also how well the driver is using the pedal. It is human nature to want to get the braking done early and this will show up as a long g trace that oscillates up and down as the driver over-brakes for the corner, realises the mistake and gets off the pedal and then back on again. This is not the classically smooth two-sided valley from the perfect world.

The braking trace will also show whether braking is being done in a straight line. In this case, getting off the brakes will show as a relatively fast rise at the end of the trace. If the driver is trail braking into the corner the trace will show that he or she is coming off the pedal gradually as lateral g forces build up.

Once the basics have been covered, it's possible to move on to finding the more subtle increments of speed. This is where the time slip line comes into its own. Comparing two team-mates and having the benefit of the organiser-provided fastest lap data shows up where time can be made up. Parts of the circuit where there is room for improvement can be identified

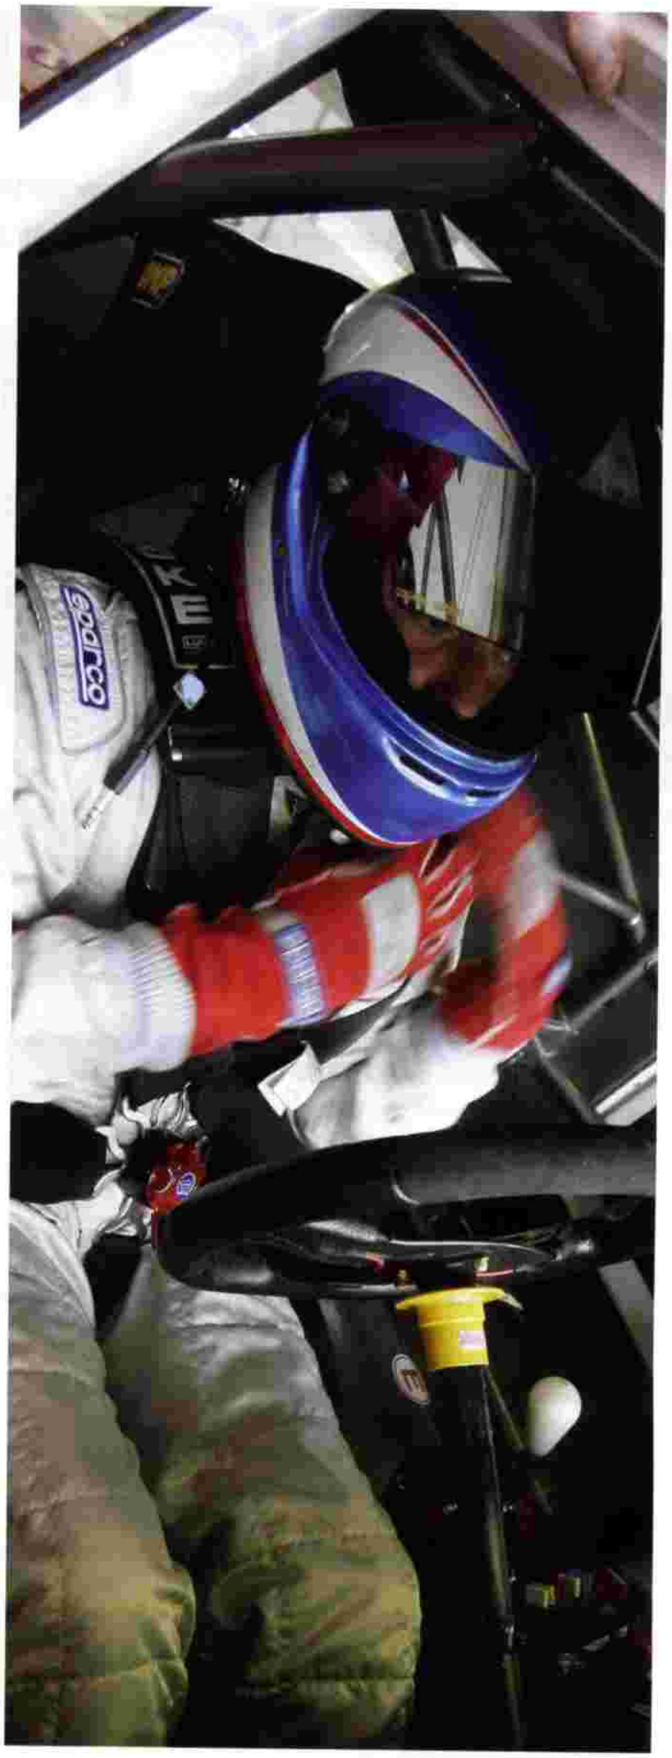

and it is possible to delve deeper into the data. Throttle position, brake pressure and steering wheel angle serve to show precisely what the driver is doing and the lateral g trace will show the impact that that has on the car's speed.

At this stage it is important not to concentrate too much on one limited part of the corner, but to realise that braking, corner entry, apex, exit and acceleration must be treated as a whole.

It is also important to have the right attitude to the data. The approach of chasing the time slip is inevitably backward-looking – trying to catch up with the best so far. Used slavishly, it can distract from the real objective, which is to improve on what has happened up to now. What must always be uppermost in everyone's mind is how things can be pushed forward.

Data interpretation is important, but it's not the only evaluation available to the team. The driver, although still learning the trade, is also closely involved and the only person capable of describing what the car felt like at a particular point.

In-car cameras are used and much use is made of Stack's ability to link the video footage to the data. This acts as an invaluable reminder to the driver about what was going on at a particular point and helps develop the skill of being able to 'see' the data as representing a particular part of the race track. Finally, this brings us back to the starting point: this is one of the few areas of motorsport where there really *can* be proper driver coaching. ■

LEFT Freedom of information: each competitor in T-Cars is able to compare their best efforts with those of the fastest driver. The rules require a Stack data logging system to be used

ABOVE Stack's ability to link video footage to data, seen here in a club test, helps develop the skill of being able to 'see' the data as representing a particular part of the track

RIGHT The driver may be alone in the race car, but the logging information enables the youngsters to be coached from the pit lane afterwards