Testing the Torsional Stiffness of Formula Student Chassis’

By Paul Revill

Acknowledgements:

Would like to thank Ryan Marks for his supervision during the project and Ian King for his aide in setting up of equipment. Finally David Williams and Paul Biggs for their guidance and assistance with the motion capture system.

1. Introduction

The aim of this paper was to determine the torsional stiffness of a Formula Student chassis and the most effective method to do so using a purpose built rig. The physical results were then used to prove an FEA model made in MSC Patran 2012.

1.1. Formula Student (FS)

The FS competition is a global inter-university competition organised (in the UK) by the Institution of Mechanical Engineers (IMechE). The FS website stated the basis of the competition was:

> "Your team is tasked to produce a prototype for a single-seat race car for autocross or sprint racing, and present it to a hypothetical manufacturing firm. The car must be low in cost, easy to maintain, and reliable, with high performance in terms of its acceleration, braking, and handling qualities." (IMechE, 2013)

Within the FS competition there were typically three types of chassis used by teams. These were a space frame chassis, a monocoque chassis or a half monocoque and half space frame chassis. The chassis' tested in this report were a half monocoque and half space frame, known as CR10 and a full monocoque known as CR07.

1.2. What is Stiff Enough?

The torsional stiffness of the chassis should be sufficient enough to not affect the suspension set up. Therefore a common practice was to state an objective torsional stiffness in terms of the roll stiffness of the suspension. Mihailidis et al. (2009) stated that the torsional stiffness of a FS car "should be at least ten times greater than the roll stiffness of the suspension". This is collaborated by Thompson et al. (1998). Bertera et al. (2014) calculated the roll stiffness of CR10 suspension as 1433Nm/deg. Therefore based upon

Mihailidis et al. (2009) statement the aim for CR10 was a torsional stiffness of 14,330Nm/deg. The set up for CR07 was not known

1.3. Aims

- Conduct physical testing using:
    - Digital Image Correlation
    - Linear Variable Differential Transformers
    - Motion Capture
- Determine most reliable method to use for future torsional stiffness testing.
- Determine torsional stiffness of CR10 & CR07.
- Create and validate an FEA model of CR10.

2. Previous Work

2.1. Previous work using Torsion Rig

The torsion rig used in the experimental testing was originally designed and it's manufacture overseen by Manning (2004). The rig has since been used by various FS teams within Cardiff University with varying amounts of success.

2.1.1. Torsional Stiffness of the 2004 Cardiff University Formula Student…

Figure 2-1 shows different scenarios of torsional testing. Manning (2004) chose Scenario B as this was more representative of what a race car would experience. Scenario A was more representative of going over a bump off-road.

Manning (2004) performed calculations to ensure the rig would be strong enough. The rig consisted of uprights to secure the front of the car and a centrally pivoted arm to be secured to the rear of the car. This arm had a manually operated jack on one side to apply the required load.

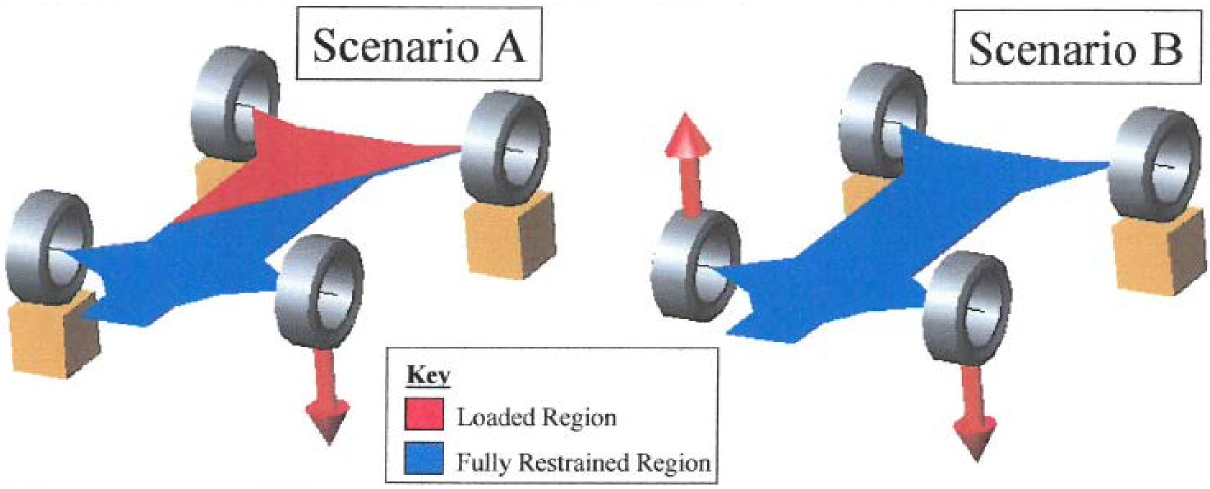

Figure 2-1 Torsion rig loading scenarios, figure from Manning (2004)

2.1.2. Previous Use of the Torsion Rig within Cardiff University

Within Cardiff University the torsion rig had previously been used for torsion testing three cars (CR01, CR03 & CR04). These cars all used motion capture to determine the torsional stiffness of the chassis'.

Featherstone et al. (2005) provided useful insight on modelling the engine as an infinitely stiff box structure. This was justified as "although obviously an approximation, [it] is reasonable due to the order of magnitude difference in stiffness between the tube of the space frame chassis and the solid engine." The FEA conducted by Featherstone et al. (2005) when compared to the physical results were significantly different. Featherstone et al. (2005) state a further publication would be done comparing the improvements however this has not been found.

The internal report by Azran et al. (2006) about testing of CR03 has significant errors in the report and thus was deemed unreliable. However a point that was carried over by Bryant et al. (2008) was that a pre-load was required to take up the slack in the system before data to determine the chassis stiffness could be taken. Bryant et al. (2008) determined the torsional stiffness of CR04 as 239.93Nm per degree, which was 14% of the FEA model. However the methods behind the FEA are not outlined and have since been lost. The issue with CR04 was because of a “flexure issue of rear suspension mounts.” Based upon photos of testing the suspension was not ‘locked out’. Therefore this would account for a significant amount of twist, as found by Wyatt et al. (2012).

An issue outlined by both Azran et al. (2006) and Bryant et al. (2008) were the chassis’ producing false markers due to their vinyl finish.

2.2. Other Work Based on Formula Student Cars

2.2.1. Design, Analysis and Testing of a Formula SAE Car Chassis

Riley & George (2002) outlined four main vehicle loading situations that should be considered when designing a Formula SAE[1] car. However Riley & George (2002) stated clearly that torsional stiffness was “generally thought to be the primary determinant of frame performance for a FSAE race car”. The report used a stick model of an example FSAE car with dampers being represented using springs and the frame being represented by a torsion spring. An equation to determine the total spring constant of the car in terms of frame, suspension and tyres was derived. The aim from this method was to be able to demonstrate the relationship between the wheel rate, the vehicles overall stiffness and the chassis stiffness. However there was no ideal or aim stated based upon these equations. Riley & George (2002) do state that based upon previous experience a chassis of at least 1600 ft-lbs/deg (≈2170Nm/deg) is needed.

3. Testing Apparatus

The testing apparatus used in the experiment were:

- Linear Variable Differential Transformers (LVDTs) along one side of the chassis connected to a data logger with a calibrated load cell to measure the displacement along the side of the car and load being applied to the pivot arm of the rig.

[1] Formula Student was known as Formula SAE (FSAE) outside of Europe.

- Digital Image Correlation (DIC) technique used to measure the strain and displacement along the same side of the chassis as the LVDTs.
- 'Qualisys' motion capture system to monitor the movement of the whole car at strategic points similar to those of the LVDTs.

3.1. Linear Variable Differential Transformers (LVDTs)

LVDTs are electromechanical transducers that convert a mechanical rectilinear movement into a corresponding electrical signal (Macro Sensors, 2014). LVDTs are reliable due to their friction free operation (which gives them a distinctively long lifespan) and have a fast dynamic response (Macro Sensors, 2014). The use of LVDTs over dial gauges reduced the time of testing (Featherstone (2005) estimates that the use of dial gauges would have extended the test time by a factor of 10) and the need for the jack to hold pressure reliably. Through the further use of a data logger with a calibrated load cell connected, the LVDT data directly corresponded to the load (and thus the torque) being applied. LVDTs were used to gain simple single axis displacement data to use as a comparison to the MCS and DIC results.

3.2. Digital Image Correlation (DIC)

Stereo DIC is a technique where two cameras which detect, through calibration, their position relative to each other and all imaging parameters (such as magnification). This allows the technique to be able to calculate the 3D coordinates for any surface point (Herbst & Splitthof, 2013). This enables the system to monitor the displacement of the surface and determine the strain on it. This technique allows the strain to be determined over a large surface area instead of using multiple individual strain gauges with Wheatstone bridges. DIC requires a completely random "speckle" pattern to enable the cameras to effectively track the images in relation to each other. The aim from studying the strain over the side of the chassis was to aid in determining weak areas of the chassis where improvements in the chassis design should be focused.

3.3. 'Qualisys' Motion Capture System (MCS)

MCS' track the position of highly retro-reflective markers mounted to fixed positions on a structure as it moves. The system use multiple motion capture units (MCUs) emitting infrared light to track the markers. A small residual error normally occurred from calibration. This error is determined by the accuracy of each MCUs ability to measure the known length of a calibration wand. The movement of these markers can then be converted using various techniques into useable data. The torsion rig has been designed to be used with this system and has been used reliably by Featherstone et al. (2005), Azran et al. (2006) and Bryant et al. (2008) in testing the torsional stiffness of previous FS cars.

4. Testing

4.1. Preparation & Methods

The rig was assembled and positioned to enable eight Motion Capture Units (MCU's) around the rig and also allow the DIC cameras to be positioned at a suitable distance. The cars were positioned centrally around the pivot point to ensure the torsion was applied through the centre of the car.

4.1.1. Jack, Load Cell and LVDTs

The jack and load cell were positioned 1.30m from the pivot point. The jack with load cell was raised before each test until there was minimal movement between the pressure plate and the load cell. The pressure plate was pivoted on a steel ball to ensure the force was applied over a constant area on the rig arm as it pivoted. The load cell was zeroed before each test began once it was in position. This method introduced a slight error in terms of the load cell accuracy, however this was minimal compared to the loads that were applied. The load cell was calibrated prior to testing for use up to 5kN.

Once the cars had been mounted within the rig and positioned centrally the LVDTs were attached at the pre-tapped locations. Various locations of interest were chosen, however the main locations were the front hub and rear hub.

The load cell and LVDTs were attached to a data logger that had a sample rate of 2Hz. The test data was exported from the data logger as a spread sheet with the resolution set to three decimal places.

4.1.2. 'Qualisys' Motion Capture System

Eight motion capture units (MCUs) were set up surrounding the rig in high & low positions ensuring that all markers placed on the cars were visible by at least 2 MCUs. Figure 4-1 show example marker placements used on the two cars. The areas where the markers were placed correspond to LVDT locations. The markers were paired on either side of the car to enable the results to be processed as vectors.

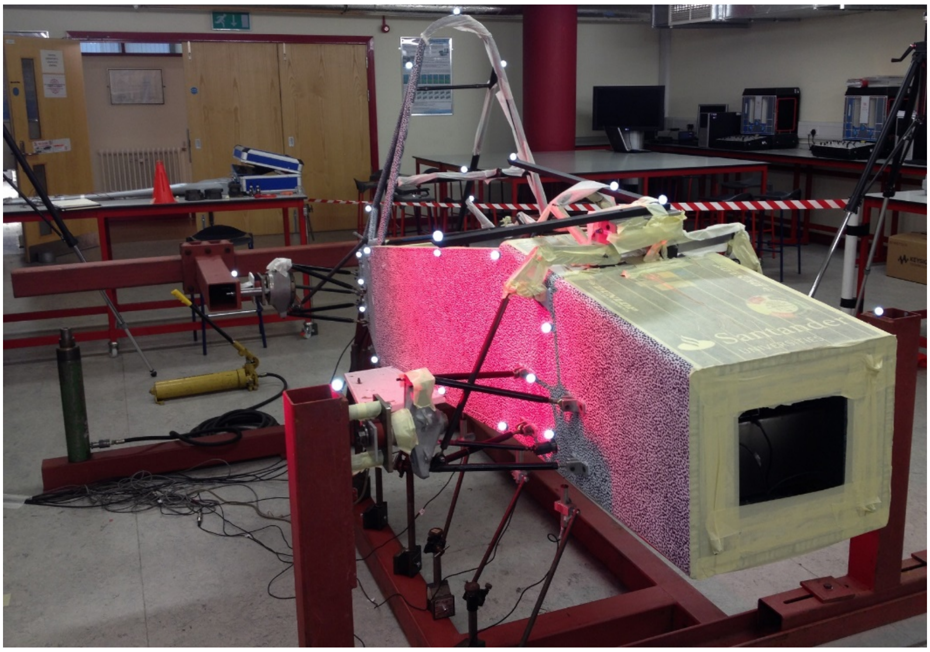

Figure 4-1 Example marker placements on CR07

4.1.3. DIC

The quality of the spatial resolution was tested on sample speckle patterns and on the cars to determine the appropriate speckle size and camera distance. The ideal speckle size was determined to be approximately 5mm as they correspond to approximately 15 pixels when the DIC cameras were at 1.6m. This was deemed as an appropriate distance as it fitted half of the chassis into the DIC images. Only half of the chassis' were selected to be monitored during the testing to ensure that most of the DIC image was filled by the chassis'.

The chassis' were spray painted white and then the speckle pattern was applied by the FS volunteer team using marker pens. Based upon the 1.6m distance the DIC was calibrated prior to testing using the calibration target. The DIC system was moved after two tests to capture the two halves of the cars. When the DIC system was moved re-calibration was carried out, following exactly the same procedure each time. In addition the DIC system logged the load cell output. This enabled the system to depict the voltage from the load cell which was proportional to the load applied. Images were taken with the DIC system with no load applied (as error checks), then at approximately every 200N as the load was applied and finally when the load had been fully removed.

4.2. Results & Discussions

4.2.1. LVDTs

The LVDT data had to be processed to gain torsional stiffness values for the various stages. The values from the LVDTs were zeroed based upon their initial readings. The next step was to convert the linear values into angles. The measurements were converted into angles based upon their distance from the theoretical roll centre of the chassis. The roll centre was taken as the line where the plane of the centre of the chassis intersected with the plane dictated by the centre of the wheels. Therefore the distance from this line was calculated for each LVDT position using basic trigonometry. Then again using basic trigonometry the angles were determined.

Although the front of the chassis was intended to be fixed in position there was movement detected there. Therefore the angles calculated were zeroed based upon this movement. The aim of the project was to calculate a chassis stiffness not including the suspension, justifying this method of zeroing. The final step in the processing was to calculate the torsional stiffness values. The loads logged were multiplied by 1.3m (the constant position of the jack) to calculate the torque value being applied by the rig. These values were then divided by the angles calculated to thus give torsional stiffness values.

4.2.1.1. CR07

From studying Figure 4-2 the LVDT movements were not stable until a load of over 400N was applied. Azran et al. (2006) stated that the application of a settling load of 0.5kN was used and so readings below this amount were disregarded. The term settling refers to the slack being taken up due to unpredictable movements such as tolerance of hole and bolt sizes and movement in the suspension arms.

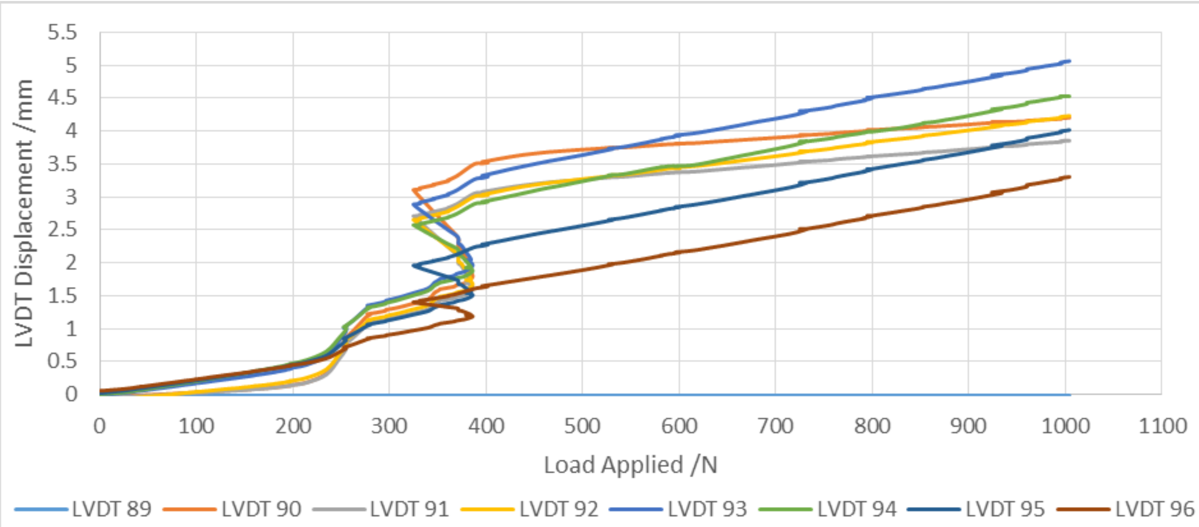

Figure 4-2 LVDT movement against load applied for first test of CR07

The LVDT data shows that CR07 appears to increase in torsional stiffness with the applied torque. This would imply that the chassis became stiffer with the higher applied torque. This was not the expected result from the torsion testing. The torsion test was expected to produce a reasonably constant torsion stiffness independent of the applied torque. This unpredicted result could be down to the assumptions made when processing the LVDT data, such as the assumption that the chassis rotates around a horizontal centre line that was the intersection of the plane dictated by the wheel hubs and the central vertical plane of the chassis. If this assumption was incorrect then the calculated angles of twist were incorrect and thus the calculation of torsional stiffness for the chassis would be incorrect.

A possible method that could be employed in the future would be to have LVDTs on both sides of the chassis. This would allow a comparison of whether the roll centre is in line with the central vertical plane of the chassis, which would be indicated by the LVDT readings being equal but in opposite directions from each set of LVDTs opposite each other on the chassis. Another strong consideration that could be the cause for the unexpected results was that the LVDT data only considered a single axis of movement which was thought to be reasonable as the movement was so small. The other consideration was that the LVDT data considered solely the bottom of the chassis and had no representation of how the sides of the chassis moved due to the applied torque.

4.2.1.2. CR10

CR10 data showed the same trends as the CR07 data, that the LVDT movement was not stable until a load of 500N was applied. Therefore again following Azran et al. (2006) only the readings recorded above a 500N load were considered.

Figure 4-3 shows the overall calculated torsional stiffness' for CR10. The graph is not fully representative of what was predicted for a torsional stiffness graph. The first and second test of CR10 shows a large amount of movement when a steady torque of approximately 1050Nm was applied. This was an indication of the angle of twist increasing at the steady torque. This was likely due to something on the car settling at this new torque. From studying the data from all the LVDTs it was shown that the settling is measured in all of the LVDT movements with decreasing influence. Therefore the settling was taking place at the front fixed end of the car.

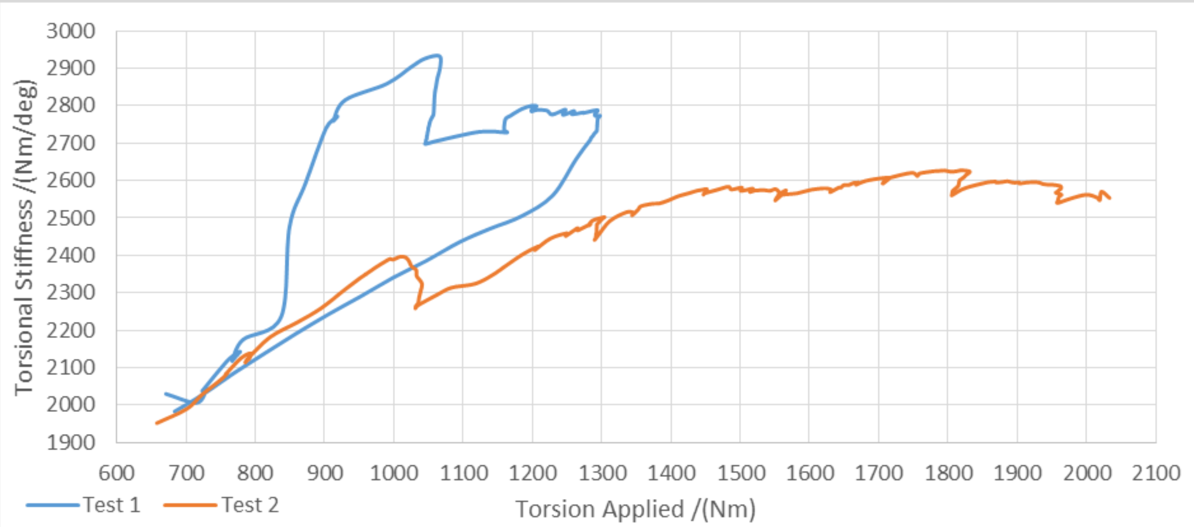

Figure 4-3 Calculated overall torsional stiffness against torsion applied for both tests of CR10

The second test of CR10 as shown in Figure 4-3 shows similar values up to the applied torque of 1050Nm. At this load it was noted that there was a reduction in torsional stiffness value. It was theorised that there was a joint which moved at this loading.

From 1400Nm onwards in the second test there was the expected result of a flat line value for torsional stiffness. Taking an average of all these values gave a torsional stiffness value of 2579Nm/deg.

The same issues as outlined in CR07 are likely to have occurred. There are many assumptions here which lead to LVDTs by far not being the best possible method for determining the torsional stiffness of a FS chassis.

4.2.2. ‘Qualisys’ Motion Capture System

The motion capture data processing began with using the MCS to track and designate all the markers. This involved ensuring that the markers movements were smooth and that no false markers had been tracked. The system had tracked them at a sampling rate of 60Hz thus ensuring all movements were tracked continuously. The data was exported from ‘Qualisys’ and imported into Matlab. Several codes in Matlab were developed to process the large amount of data. One code calculated average values over steady periods when the chassis’ had set loads applied. The second code created vectors in 3D space that could be used to calculate the angle of twist the chassis went under. Where possible multiple vectors at the same locations were used to calculate the angle of twist.

The loads that have been used for calculating the torsional stiffness values are ideal loads and the actual load applied could have been approximately ±10N due to the jack not holding a constant load. The LVDT data has shown that loads below 500N should be disregarded and this was further supported by the MCS data. Therefore the effect of the error margin was up to ±100Nm/deg on the lightest load considered however this reduced to ±10Nm/deg on the highest load considered.

Figure 4-4 shows the torsional stiffness values calculated from the MCS data against the applied torque. The dotted line in the figure represents the 500N load and so any torsional stiffness values calculated with a load beneath this value were not being included in average calculations or other figures. The values were included to show that for CR10 the two tests produce very similar trends. This can be contributed to the hysteresis that has been shown in the LVDT data. Due to the load being taken off uniformly it was not possible to create values of torsional stiffness values whilst the load was being removed for the motion capture data which would have shown hysteresis. The first CR10 test showed the start of levelling out that was reflected in the second test. The second test was taken to a higher load and the figure shows that at these higher loads the torsional stiffness was more constant. This could be due to the first test causing all deformations that were then permanent. On the second test these deformations had already occurred and as such the angle of twist values would be greater causing the torsional stiffness values to be lower. As such if the CR10 test had been repeated again the values would be expected to match or be similar to those of the second test.

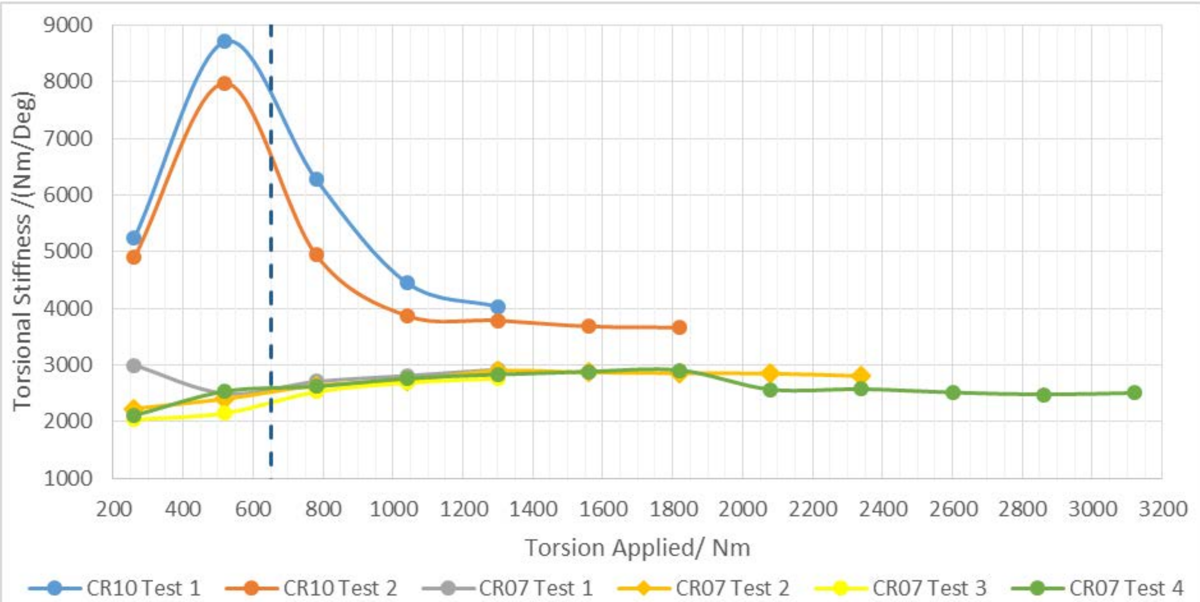

Figure 4-4 Overall torsional stiffness' calculated from motion capture data against applied torque

Figure 4-4 shows a high torsional stiffness at 520Nm torque, for CR10, which then decreases and levels out. This was not an expected result and is believed to be the result of movement at the front of the chassis which should have been fixed in place and as such was compensated for.

Figure 4-4 shows fairly consistent and linear lines of torsional stiffness above the 500N pre-load. Test 4 does shows a drop above 1800Nm torque which is likely due to the chassis being broken in test 3.

Figure 4-5 (a) & Figure 4-5 (b) show the torque applied against angle of twist for CR07 and CR10 respectively. The data plotted is above the 500N allowance for stabilising the rig. The torsional stiffness of both chassis should be constant and therefore a linear line of best fit was plotted on the figures. This line was specified to pass through the origin, as the torsional stiffness is constant it is reasonable to assume that when no torque is applied there is no angle of twist. Also when torsion was applied there would be a twist angle generated straight away, thus meaning that there would not be a y-intercept.

Figure 4-5 (a) gives a value of 2,637Nm/deg for the torsional stiffness of CR07.

Figure 4-5 (b) gives a value of 3,176Nm/deg for the torsional stiffness of CR10.

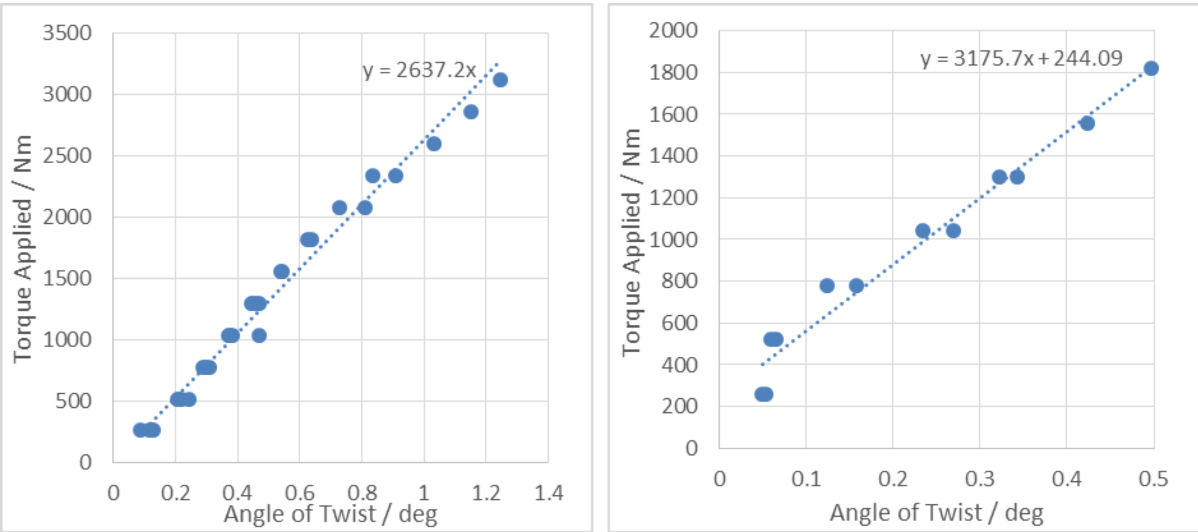

Figure 4-5 (a) Applied torque against angle of twist for CR07

(b) Applied torque against angle of twist for CR10

It should be noted however that CR07 was tested without an engine or drivetrain. The CR07 chassis had also been modified for use as a simulator. These modifications include locking out of the rear suspension and a reduction in the size of the front roll bar. Therefore these results were not directly comparable.

4.2.3. DIC

The DIC data was processed using a dedicated program. Whilst the main aim of the DIC was to study the strain over the side of the chassis. The DIC did not measure any distinguishable strain. The DIC most likely did not work due to it being stereo DIC. Stereo DIC works in a single plane however the chassis' were thought to be experiencing strain in multiple planes. The DIC pattern on the space frame did not work at all. This was likely again due to such a large area being studied and the space frame tubes were very thin in comparison. The space frame is also a circular cross section and as such the DIC cameras would have focused only on a thin line of the space frame. This meant that even if the DIC cameras had been focused solely on the space frame it was unknown if the DIC would have picked out all the bars.

4.3. Comparison

The MCS and DIC data had different co-ordinate systems therefore to enable comparison of data the magnitude of the movement vector from no load was calculated. Figure 4-6 shows a selection of the magnitudes plotted against load applied for all methods used.

The MCS and DIC data have a very strong correlation showing that these methods are most likely the most reliable. The slight discrepancies between the data sets could be due to the gauge positions for the DIC data not being in exactly the same position as the markers. The LVDT data was consistently lower than the magnitudes of the MCS and DIC measurements. This is due to it measuring a single axis instead of in 3D space.

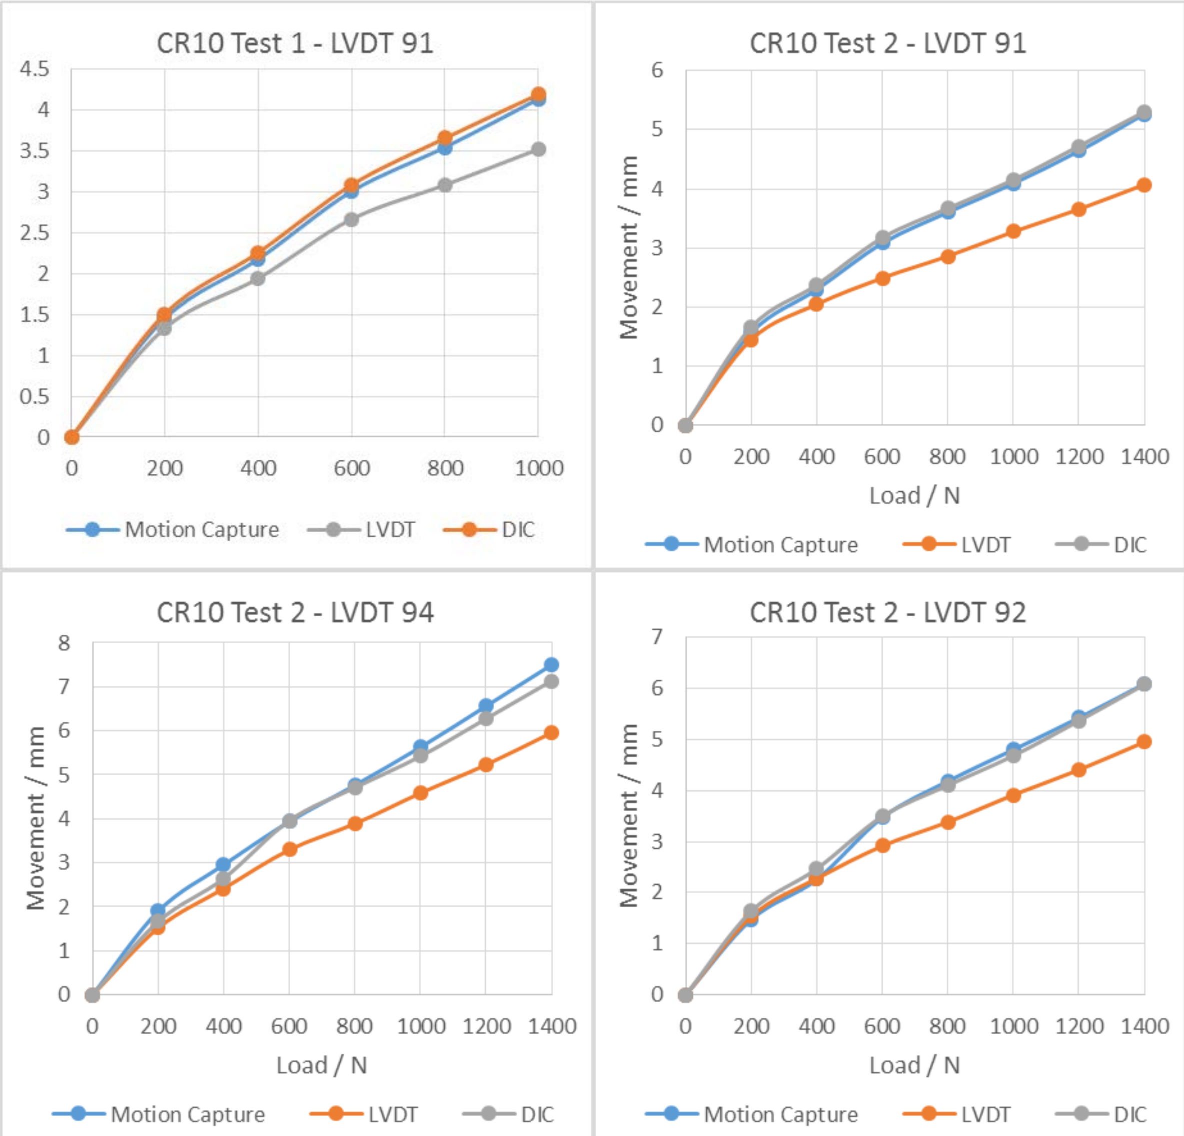

Figure 4-6 Representative examples of MCS, DIC and LVDT data comparison

5. FEA Validation

The FEA model of the chassis was compared with the MCS physical test results as these were proven as the most reliable and comprehensive. The FEA model of CR10 was produced using a space frame model created from a model by *Evans* as the basis. The monocoque was then created from this model following guidelines set out by Griffiths et al. (2014). The model developed was tested against the same constant loads as achieved through the MCS data. Figure 5-1 shows the FEA data and the MCS data.

Quad8 mesh was used for the monocoque. The monocoque was restrained in line with the centre of the front hubs. This reflected the physical tests when processing the MCS data. The torque was applied at the rear at the location of the centre of the rear hubs. These locations were attached to the space frame through infinitely stiff members. The members attached to the space frame at the locations the suspension members would have been attached. The honeycomb was modelled as a laminate based upon FEA models made by Rigby (2014).

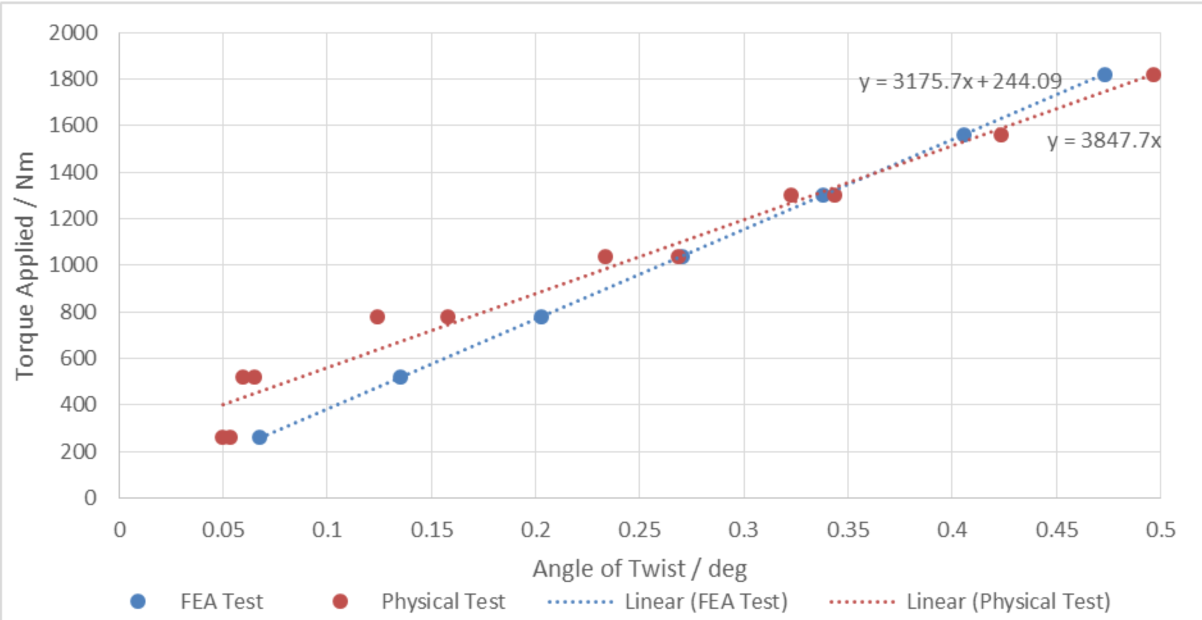

Figure 5-1 Applied torque against angle of twist for CR10, physical testing data and FEA data

The FEA gave a torsional stiffness value of 3176Nm/deg. There was correlation with only 17.5% difference between the FEA and physical testing results. It was expected that the FEA would be greater than the physical testing due to the ideal situation present. The FEA model matched the test data because of the high stiffness given to the engine block modelled. Featherstone et al. (2005) stated that the engine can be modelled as an infinitely stiff object as it was considerably stiffer than the space frame material. The mesh size used on the monocoque was based upon guidelines set out by Griffiths et al. (2014) and the model supplied by *Evans*. However performing studies on the mesh properties (such as on size, aspect ratio and skewness) to gauge an ideal mesh size would be beneficial for the future of Cardiff Racing.

5.1. Comparison

The FEA model of CR07 produced a torsional stiffness value of 3,431Nm/deg and the CR10 FEA model produced a torsional stiffness value of 3,848Nm/deg. However, CR07 was modelled without an engine which was theorised by Wyatt et al. (2012) as having an impact on the torsional stiffness. The amount this would have affected the torsional stiffness is unknown. Therefore it is likely the cars were of similar torsional stiffness when fully operational (as there was only a 13% difference in FEA values). Section 1.2 showed that the torsional stiffness value that Cardiff racing should be aiming for was approaching 14,330Nm/deg. The values determined through physical testing and FEA models showed that the Cardiff Racing cars were not approaching this value. This would explain why the cars always suffered from understeer. The chassis' were making the suspension set up unreliable. Sampo et al. (2010) contradict the commonly accepted aim for torsional stiffness and state that a torsional stiffness higher than five times the roll stiffness was enough. This would mean a target of 7,165Nm/deg or higher. The torsional stiffness values acquired were half of this aim still and thus even with a lower (more achievable target) the cars were still not torsionally stiff enough.

6. Conclusion

Physical testing has been successfully conducted on CR07 and CR10 with values of 2,637Nm/deg and 3,176Nm/deg respectively for their torsional stiffness. The FEA models have produced torsional stiffness values for CR07 and CR10 of 3431Nm/deg and 3,848Nm/deg respectively. Whilst these values exceed the stated aim of Cornell University (Riley & George, 2002) of 1600ft.lbs/deg (2169Nm/deg), they do not meet the stated aim of

14,330Nm/deg as calculated based upon Mihailidis (2009) specifications of torsional stiffness in terms of roll stiffness. The physical test results were based upon the MCS data as the LVDT data was found to be unreliable due to measurement in only one axis. No strain measurements were found with the DIC due to the strain occurring in multiple planes. The DIC displacement was compared with and found to have a strong correlation with the MCS measurements. Therefore the MCS has been determined as the most appropriate method for determing the torsional stiffness of FS chassis'. Validation of FEA models has been performed with the CR10 model being within 17.5 % of the physical testing results. The CR07 model was not fully validated due to the discrepancies between it and the physical car tested. Some improvements have been recommended, for both types of car, based upon the findings and future work has been suggested.

7. References

Azran, A. et al., 2006. *CR03 Report - Torsion Testing,* Cardiff: Cardiff University - 4th Year Group Report.

Bertera, B. et al., 2014. *Suspension Team - Interim Report,* Cardiff: Cardiff university - 4th Year Project.

Bryant, M. et al., 2008. *Section, CR04 Report - Torsion Testing,* Cardiff: Cardiff University - 4th Year Group Report.

Featherstone, C. et al., 2005. *Measuring the Torsional Stiffness of a Space Frame Chassis using 3D Motion Capture Techniques.* London, Trans-Tech Publications.

Griffiths, D. et al., 2014. *CR10 - Chassis & Suspension Team,* Cardiff: Cardiff Unviersity - 4th Year Project.

Herbst, C. & Splitthof, K., 2013. *Basics of 3D Digital Image Correlation.* [Online]
Available at: http://www.dantecdynamics.com/docs/products-and-services/dic/T-Q-400-Basics-3DCORR-002a-EN.pdf
[Accessed 12th February 2015].

IMechE, 2013. *Formula Student - The Challenge.* [Online]
Available at: http://events.imeche.org/formula-student/about-us/thechallenge
[Accessed 15th February 2015].

Macro Sensors, 2014. *LVDT Basics.* [Online]
Available at: http://www.macrosensors.com/downloads/misc/LVDT-Basics.pdf
[Accessed 12th February 2015].

Manning, D., 2004. *Torsional Stiffness of the 2004 Cardiff University Formula Student Chassis,* Cardiff: Cardiff University - 3rd Year Project.

Mihailidis, A. et al., 2009. The Design of a Formula Student Race Car: A Case Study. *Journal of Automobile Engineering,* Volume Part D, pp. 805-818.

Rigby, J., 2014. *Finite Element Analysis of Aluminium Honeycomb Panels,* Cardiff: Cardiff University - 3rd Year Project.

Riley, W. B. & George, A. R., 2002. *Design, Analysis and Testing of a Formula SAE Car Chassis.* Indiana, SAE International.

Sampo, E., Sorniotti, A. & Crocombe, A., 2010. *Chassis Torsional Stiffness: Analysis of the Influence on Vehicle Dynamics.* Warrendale, SAE International.

Thompson, L. L., Lampert, J. K. & Law, H. E., 1998. *The Effects of Chassis Flexibility on Roll Stiffness of a Winston Cup Race Car.* Michigan, Society of Automotive Engineers.

Wyatt, B. et al., 2012. *CR08 Report - Torsion Testing,* Cardiff: Cardiff University - 4th Year Group Report.