Real-Time Side-Slip Angle Measurements using Digital Image Correlation

Devin K. Johnson[a], Theunis R. Botha[a], P. Schalk Els[a,*]

[a] *Vehicle Dynamics Group - Dept. of Mechanical and Aeronautical Engineering - University of Pretoria - Private Bag X20, Hatfield, 0028 - SOUTH AFRICA*

Abstract

In vehicle dynamics there are many parameters that are desired for vehicle control and modelling. One of the most important parameters for handling and stability is the vehicle side-slip angle. The ability to directly measure the vehicle side-slip in real-time will aid and improve many driver assist systems such as stability control schemes and roll-over mitigation, especially over rough terrain.

Commercial side-slip angle solutions are available but they are prohibitively expensive and are only suitable for use during vehicle development and performance evaluation. They are also restricted to small side-slip angles and give unsatisfactory results at low speeds and over uneven terrain. Previous research has proven that digital image correlation can be used to accurately measure vehicle side-slip angle over rough off-road terrain using inexpensive, off-the-shelf cameras. However, side-slip angle calculations were performed in post processing from pre-recorded footage and not implemented in real time due to the large computational times of the novel algorithms developed.

This paper describes the improvements made to the algorithms that enable real-time implementation. The side-slip angle is measured using a single camera pointing downwards to the terrain and digital image correlation. The sensor is tested on a flat surface using a rig that allows for validation. The maximum sampling frequency and accuracy are investigated. The system is shown to measure accurately and in real-time up to 100km/h speeds.

Keywords: computer vision, side-slip angle, off road, image correlation

1. Introduction

Modern passenger vehicles rely on complex control systems to assist drivers in avoiding accidents or at least to limit damages in the event of a collision. These systems, known as Advanced Driver Assist Systems (ADAS) use features such as Anti-lock Brake Systems (ABS) and Electronic Stability Control schemes (ESC) to enhance the safety of the occupants. Typically, these systems require a large combination of sensors to determine the state of the vehicle. A vehicle model then uses the information to predict the behaviour of the vehicle within normal driving conditions. The predicted model is compared to the actual vehicle behaviour and then corrects any deviation using ADAS (Rajamani, 2005). One parameter that can greatly aid the performance of ADAS is the vehicle side-slip angle, β. Vehicle side-slip angle can be used as a measure of the vehicle's handling and stability and is, therefore, a valuable parameter in vehicle dynamics. Inagaki et al. (1995) demonstrated that β and it's derivative, $\dot{\beta}$, offer better insight to the vehicle's lateral stability as compared to yaw-rate. However, due to the difficulty in measuring the side-slip angle most control systems rely on the yaw rate. Chung and Yi (2006) however found that using a stability control scheme that was based on side-slip angle resulted in an overall improved vehicle performance on a virtual test track.

The tyre-road interface is one of the most important research areas in the field of vehicle dynamics, due to all vehicle excitation forces acting at this interface (excluding aerodynamic forces). Many parameters that govern the forces at this interface are required to fully understand the dynamics, one of which is the side-slip angle. Bakker et al. (1987) showed that there is a strong relationship between the lateral force generated by the tyre and the tyre side-slip angle. Most tyre models require accurate side-slip angle measurements as a necessity for the characterisation of tyres, especially during dynamic manoeuvres or when performing tests on tyre test tracks.

Measuring side-slip angle on off-road terrain presents unique challenges such as measuring

*Corresponding author
Email addresses: devin.kent.johnson@gmail.com (Devin K. Johnson), theunis.botha@up.co.za (Theunis R. Botha), schalk.els@up.ac.za (P. Schalk Els)

at low speeds and over rough, uneven terrain where current measuring solutions fail to accurately measure in these conditions.

Side-slip angle is notoriously difficult to measure, however, such sensors do exist. The Kistler Correvit S-HR (Kistler, 2016) is a commercially available side-slip angle sensor, developed for mainly smooth, hard roads where the motion is predominately planar with little body motion. This sensor uses a combination of the Doppler effect and an absolute measuring method for determining the side-slip angle. The sensor is limited to a maximum side-slip angle of $\pm$ 15deg, does not perform well below 20km/h and experiences difficulties when moving over uneven surfaces. Therefore, this sensor is not suitable for terramechanic applications which generally occurs on uneven terrain and low speeds. The sensor has a maximum sampling frequency of 250Hz. Due to prohibitive costs of such sensors, the field of vehicle dynamics has opted to instead estimate the vehicle side-slip angle rather than directly measure it. These estimation methods use sensor fusion techniques that combine sensors such as accelerometers, GPS, and rate gyroscopes to estimate the side-slip angle. These techniques proved successful and correlated well with the values of simulated vehicle side-slip angle (Bevly et al. (2006), Botha and Els (2012) and Hac and Simpson (2000)). These sensors are inherently noisy and require large sensor excitations, such as experienced high-speed dynamic manoeuvres, to provide accurate results and are also highly prone to integration drift. As a result, these estimation methods are unsuitable for off-road scenarios, where tests occur at low speeds and experience high levels of noise due to terrain roughness. Tyre side-slip angle is not typically measured in field tests, however, are measured during tyre characterisation. The tyres are typically mounted to a rig where the side-slip angle is accurately controlled (Dora et al., 2006).

Botha and Els (2015) proposed an alternative to current methods whereby the side-slip angle could be accurately measured using inexpensive, off-the-shelf cameras and digital image correlation. This technique was developed to overcome the hurdles faced with measuring the side-slip angle in off-road conditions. It measures side-slip angle robustly at low speed and over varied terrain and does not require large dynamic excitation. This technique was successfully tested on smooth concrete, rough cobblestone paving, snow, ice and various

other mixed off-road conditions at various speeds and on several different vehicles. Various methods were proposed, either using a single camera or a calibrated stereo-vision rig, comprising of two cameras at a fixed distance apart. The sensor could be mounted on the vehicle for vehicle side-slip angle or on the tyre for tyre side-slip angle. For both cases, the sensors must be facing downwards towards the terrain at all times. The extraction of side-slip angle from the digital images was performed in post-processing from pre-recorded footage due to the large computational expense of image processing. This technique proved to give excellent results that are valuable for vehicle and tyre testing under typical rough, low speed, off-road conditions where terramechanic aspects are important.

This study builds on Botha and Els (2015) by implementing the single camera technique in real-time. The sensor can be used for vehicle and tyre testing but also as a direct input to ADAS systems. The result is a camera-based sensor that uses off-the-shelf cameras and lenses to accurately measure the side-slip angle in real-time and at low cost.

2. Side-slip angle with digital image correlation

Three algorithms were developed by Botha and Els (2015). Algorithm 1 uses a simple and efficient planar motion algorithm to track the motion of a single camera. Identifiable points on the image, known as features, are tracked across sequential images, meaning the location of that point is found on the corresponding image. Feature tracking will be discussed in Section 3.2. The locations of the features are represented in pixel coordinates, and not in real-world coordinates. The direction that the features has moved in, is considered the direction of travel of the camera. The camera determines the lateral (Δx) and longitudinal (Δy) motion of features in pixels. Considering that the lateral and longitudinal movement is known, the angle can be calculated using Equation 1.

$$\beta = \arctan(\frac{\Delta x}{\Delta y}) \tag{1}$$

By mounting the camera on the vehicle so that it is facing downwards towards the terrain, the motion of the camera is then considered the planar motion of the vehicle. If the camera is mounted such that the feature motion has an angle of zero when the vehicle undergoes

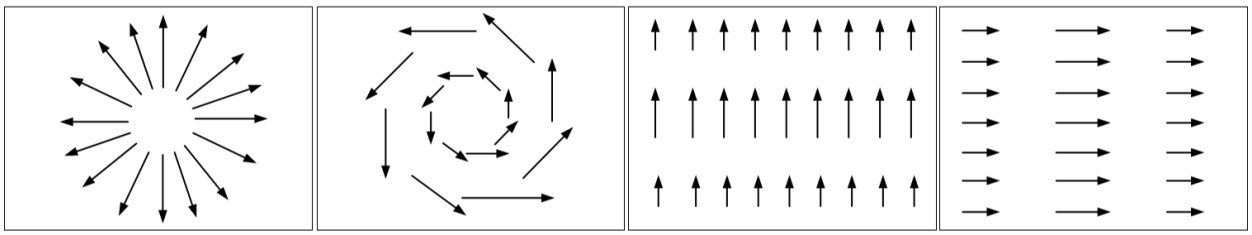

Figure 1: Observed effect of vertical motion, yaw, roll, and pitch on feature vector motion (Botha and Els, 2015)

no side-slip, such as when driving in a straight line, then the angle of the feature motion vector is equal to the side-slip angle. This method is the least computationally expensive of the algorithms, making it the most suitable candidate for implementing in real-time.

This algorithm makes the assumption that the vehicle undergoes purely planar motion. This is true for testing on smooth flat surfaces, however, the assumption does not hold when the vehicle is travelling over rough terrains where the terrain induces other vehicle motions such as vertical displacement, roll, pitch and yaw. Figure 1 demonstrates the effect of various vehicle motions on the feature vector motion.

To combat this, algorithms 2 and 3 were developed by Botha and Els (2015). These algorithms attempted to determine the rigid body motion that the vehicle experiences. By doing so, the pure lateral and longitudinal motion can be distinguished from the other motions. Algorithm 2 uses pose estimation which determines the rigid body motion, using a single camera. As before, features are tracked between sequential images that are used to calculate the essential matrix (Stewenius et al., 2005), which maps the relative rotation and translation motion between views, whereby all rotations can be determined. The translational velocities can also be estimated although these are in scale, meaning that they need to be multiplied by a scale factor to convert them to real-world coordinates. The scale is dependent on the motion and is therefore not consistent throughout the measuring process. The result is that the pure planar motion can be distinguished from the rigid body motion experienced when travelling on off-road terrain. The side-slip angle can be calculated by taking the angle between the lateral and longitudinal movement, as shown in Equation 1. Since the angle is not affected by the scaling the scaling does not pose a problem. This

algorithm requires many complex iterations and error checking before converging to a suitable answer, making it computationally expensive and potentially difficult to implement in real-time.

The third algorithm uses a stereo-vision rig that obtains all measurements directly in unscaled real-world units. This method uses two cameras that view the same scene at the same time which allows for 3D measurements. Once a 3D map of the scene has been generated, features on the map can be tracked between sequential maps. As before, the rotation and translation matrices can be determined from the tracked points. However, unlike algorithm 2, these are unscaled and represent the motions in SI units. Similar to algorithm 2, the pure planar motion underwent by the vehicle can be distinguished from other vehicle motions and the side-slip angle calculated by taking the angle between the lateral and longitudinal movement, as shown in Equation 1. The 3D map generation is computationally expensive and is the bottleneck for computation time in this algorithm, making it difficult to implement in real-time time. One major advantage of this algorithm is that many other vehicle parameters can be extracted simultaneously such as roll and pitch angles and the vehicle velocity.

Botha and Els (2015) validated the algorithms by comparing the measured side-slip angles to known side-slip angles. The cameras and sensor were mounted on a gimbal that allowed them to be rotated in one-degree increments relative to the vehicle. This allows for the side-slip angle to be set at a fixed angle that forces a known artificial side-slip angle to be measured by the sensor while the vehicle travels in a straight line. This method was chosen as opposed to dynamic manoeuvres that would induce vehicle side-slip as it is not a repeatable and controlled process. Although measurements were compared to the Kistler Correvit S-HR sensor, there is also no accurate comparative measure that would validate the measurements. The gimbal system was mounted close to the vehicle's centre of gravity, or CG, to reduce the effect of other vehicle motions. These tests were conducted at various side-slip angles and over two terrains, a flat concrete surface and a rough Belgian paving to simulate off-road terrains. The vehicle velocity did not exceed 20km/h during these tests. The results of the tests are summarised in Table 1 that show the mean values, the standard

deviation (STD) and the root-mean-square (RMS) values. It is evident that all three the algorithms outperformed the Correvit S-HR sensor with smaller errors and lower noise levels proving the use of cameras as a viable alternative to the sensor.

Table 1: Results obtained for all techniques at various side-slip angles and over two terrains

	Correvit S-HR			2D Planar			2D Pose			3D		
Flat[deg]	RMS	Mean	STD	RMS	Mean	STD	RMS	Mean	STD	RMS	Mean	STD
0	2.30	0.42	2.26	0.91	-0.36	0.84	0.91	-0.31	0.85	0.83	0.06	0.83
2	2.76	2.45	1.28	1.87	1.75	0.66	1.89	1.76	0.69	1.81	1.70	0.60
5	5.86	5.48	2.09	4.98	4.90	0.89	5.02	4.92	0.98	4.93	4.85	0.87
10	10.39	10.10	2.48	9.96	9.95	0.53	9.97	9.95	0.60	9.81	9.77	0.87
Belgian[deg]	RMS	Mean	STD	RMS	Mean	STD	RMS	Mean	STD	RMS	Mean	STD
0	1.71	0.65	1.58	0.91	-0.09	0.91	0.94	-0.02	0.93	0.92	-0.09	0.92
2	2.98	2.52	1.59	2.10	1.85	0.99	2.18	1.92	1.01	2.06	1.83	0.95
5	5.60	5.42	1.42	4.92	4.85	0.83	4.97	4.90	0.86	4.89	4.79	0.96
10	10.56	10.35	2.12	10.03	10.00	0.75	10.07	10.04	0.78	9.90	9.87	0.87

Comparing the results from the flat surface to the Belgian paving the STD of all three algorithms increased. However, they only increased (on average) by 10% and the error never exceeded 1deg. The mean values stayed relatively constant with a low error. Even though the assumption of pure planar motion was clearly violated over the Belgian paving, the 2D Planar algorithm performed better than expected with comparable accuracies to the other algorithms. This leads to the conclusion that using a more computationally expensive algorithm to distinguish between the planar motion and the rigid body motion does not significantly improve the accuracy of the results. The computationally efficient, 2D Planar method, is therefore sufficient and the best candidate for real-time implementation and will be used during this investigation.

3. Computer Vision Techniques

In this paper an open source computer vision library called OpenCV (OpenCV, 2017) was used and all computer vision techniques were implemented in C++. It was chosen due to the large library of optimised algorithms that allows for fast execution, as required for real-time implementation.

3.1. Camera calibration

Digital images require pre-processing to remove distortion created by the lens before they can be used for accurate computer vision measurements. The lens distortion can be mathematically modelled and therefore corrected. The required parameters can be determined through calibration. The most popular calibration method was developed by Zhang (2000) and uses images of a known pattern and size, at various angles and distances from the camera to solve the lens model parameters in a least squares manner.

3.2. Feature tracking algorithms

The 2D planar method requires features to be tracked between images. This is achieved through optical flow techniques. Sparse optical flow algorithms are computationally less expensive as they track the motion of small subsets of the image instead of the motion of the entire image, as is the case in dense optical flow, making it more suitable for real-time implementation.

The most common and widely used sparse optical flow algorithm is the Lucas-Kanade algorithm (Lucas and Kanade, 1981). However, many other feature tracking algorithms exist such as the Scale-Invariant Feature Transform (SIFT) (Lowe, 2004), Speeded-Up Robust Features (SURF) (Bay et al., 2008), Oriented FAST and Rotated BRIEF (ORB) (Rublee et al., 2011) and Binary Robust Invariant Scalable Keypoints (BRISK) (Leutenegger et al., 2011). More algorithms do exist but will not be considered during this investigation. These algorithms differ mainly by their feature descriptor and typically use a simple brute force matching strategy that matches features by comparing them against all other possibilities using a suitable metric.

A comparative study was conducted between the above-mentioned algorithms to determine which algorithm would be the best candidate for real-time applications. These algorithms were selected as they are included in the OpenCV library.

Table 2 shows the computation time required to identify and track 50 features across a 640x480 sized image on a desktop computer with an i5 quad core processor 2.8GHz, using a single core. Pre-recorded footage of a vehicle travelling over an asphalt road was used

for this investigation. From this, it is evident that the Lucas-Kanade (LK) algorithm is the least computationally expensive and was able to track the features in 15ms, which is less than half it's closest competitor. The accuracy of the algorithms was not tested during the analysis as it is outweighed by the importance of computation time.

Table 2: Various feature tracking algorithms performance comparison

Algorithm	Average Computation time [ms]
Lucas-Kanade	15
SIFT	570
SURF	95
ORB	35
BRISK	550

3.3. Lucas-Kanade feature tracking

The Lucas-Kanade optical flow algorithm tracks features across images by solving the optical flow equation in a least squares approach. The algorithm makes three important assumptions to simplify the optical flow equation. The first assumption is that brightness across the tracked images are consistent. The first assumption would be invalid if sudden shadows are generated but can be overcome by using adjustable camera parameters such as sensor gain, aperture and contrast. The second assumption is that feature motion is small. The assumption of small motion is partially overcome by introducing a modification to the algorithm developed by Bouguet (2000) that tracks over a pyramid of smaller scaled images which reduces the motion on the image. The motion on the smallest image is then tracked and the solution used on the next largest image. This process is repeated for the larger scaled images until the solution of the motion of the original image is found. The third assumption is that the windowed region around the feature has the same motion as the feature itself. This is valid when planar motion is observed from a stationary scene although may not hold if multiple objects move or if the motion is predominantly rotational. These assumptions reduce the optical flow equation to Equation 2, known as the Lucas-Kanade

equation, where I_x, I_y and I_t represent the partial derivative of the pixel intensity in the x and y directions as well as with respect to time t. The algorithm solves this equation to determine the velocity in the x-direction, V_x, and the velocity in the y-direction, V_y. The exact location is found by iteratively minimizing the cost function in Equation 3 by stepping in the direction of the velocity vector at each iteration, where $I(x, y, t)$ represents the pixel intensity at pixel coordinate (x, y) at time t and $I(x + u, y + v, t + 1)$ represents the pixel intensity in the comparative image.

$$\begin{bmatrix} V_x \\ V_y \end{bmatrix} = - \begin{bmatrix} \sum I_x I_x & \sum I_y I_x \\ \sum I_x I_y & \sum I_y I_y \end{bmatrix}^{-1} \begin{bmatrix} \sum I_x I_t \\ \sum I_y I_t \end{bmatrix} \tag{2}$$

$$\epsilon = \sum_{i}^{m} \sum_{j}^{m} [I(x, y, t) - I(x + u, y + v, t + 1)] \tag{3}$$

The Lucas-Kanade algorithm firstly requires features that can easily be tracked. The algorithm developed by Shi and Tomasi (1994) has been mathematically selected to be optimal in order to accurately and reliably solve the feature motion using the Lucas-Kanade tracking algorithm and is used during this investigation. The features are selected by first taking the second derivative of the image pixel intensities to obtain the Hessian matrix $H(x, y)$ in a window around a pixel:

$$H(x, y) = \begin{bmatrix} \frac{\delta^2 I}{\delta x^2} & \frac{\delta^2 I}{\delta x \delta y} \\ \frac{\delta^2 I}{\delta y \delta x} & \frac{\delta^2 I}{\delta y^2} \end{bmatrix} \tag{4}$$

Considering the eigenvalues of the $H(x, y)$, if the smallest eigenvalue is above a set threshold and the ratio between the eigenvalues are below another set threshold, the point is classed as unique and easy to track. These are typically corner points that have high contrast in multiple directions.

4. Algorithm

The algorithm starts by obtaining images from a single camera that is pointed downwards towards the terrain. Before features are identified, the distortion is removed from

the images. Features are identified and then subsequently tracked using the Lucas-Kanade algorithm. Since the tracked features may have outliers present, the Random SAmple Consensus (RANSAC) (Fischler and Bolles, 1981) algorithm is used. RANSAC is an iterative algorithm that is used to estimate parameters of a model from data which contains outliers and was used to filter the tracked features to ensure no outliers obscure the side-slip angle measurements. A regression model is then generated from only inliers to calculate the side-slip angle as per Equation 1. If the magnitude of the feature motion vector is less than 2 $pixels$, the side-slip angle measurement is rejected to ensure no loss of accuracy. Figure 2 shows a typical output of the algorithm and shows the feature motion vectors. Green feature motion vectors have been classed as inliers, whereas blue feature motion vectors are outliers and are not considered in the calculations. The algorithm was set to track 50 features.

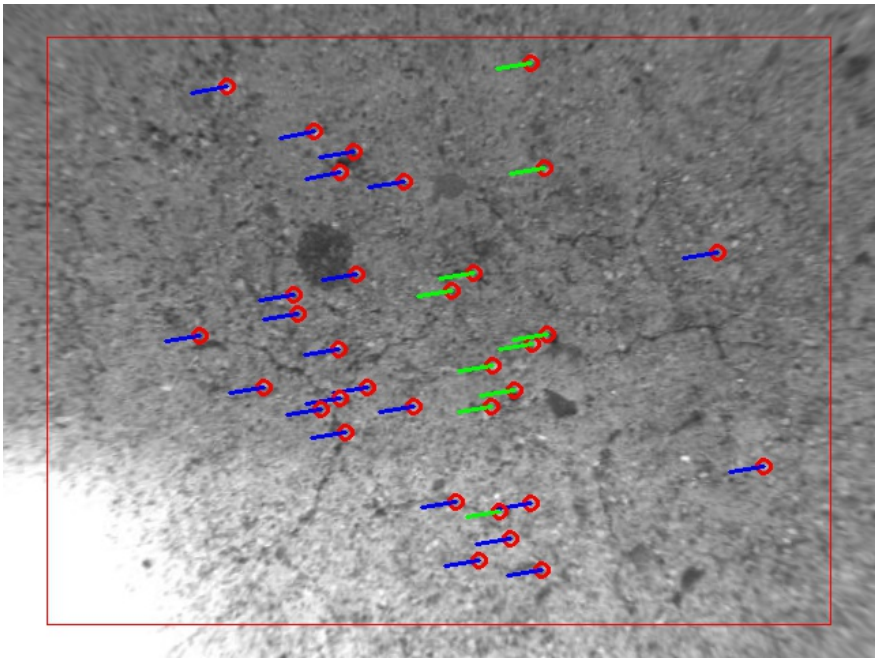

Figure 2: Features tracked across sequential images

The algorithm was threaded across all cores (thus four threads were created) to boost the computational efficiency. This results in each core running its own local version of the algorithm. This increases throughput (the overall sampling frequency of the system) but does not decrease latency (the time it takes to process a set of images). The main thread would

continuously check whether a thread was busy and then supply it with a set of images for processing once it became available. For clarification, Figure 3 shows a flow diagram of the side-slip angle algorithm.

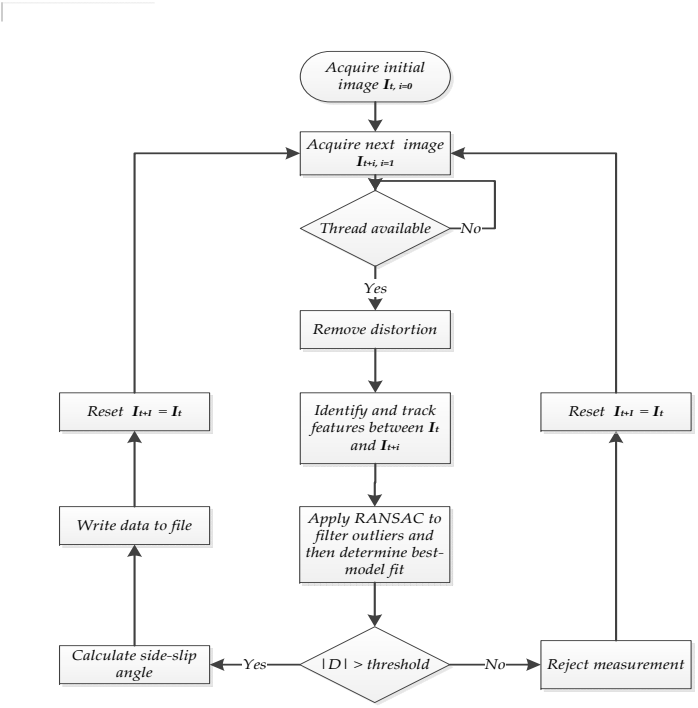

Figure 3: Flow diagram of side-slip angle algorithm

5. Testing

The side-slip angle measurement was tested using the same test setup as Botha and Els (2015) described in Section 2 and is shown in Figure 4. The camera was mounted 450mm from the testing surface, resulting in a field-of-view (FOV) of 480mm x 350mm. Ideally, the camera should be mounted on or close to the centre of gravity of the vehicle to minimise the effect of vehicle motion, however, for testing purposes, it was mounted on the rear of the vehicle. The vehicle was driven in a straight line on a flat concrete road and the side-slip angle set to 0, 5 and 10 deg. It should be noted however that the system can theoretically

measure any angle but is limited because of the gimbal setup only allowing up to 10deg rotation.

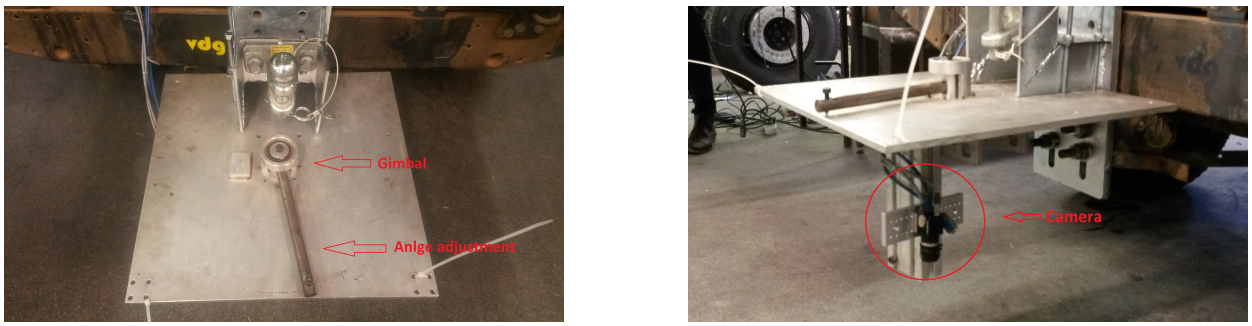

a) Top view of gimbal setup

b) Side view of gimbal setup

Figure 4: Camera Gimbal Setup

The camera used in during this investigation is a Point Grey Flea3 USB 3.0 camera with Fujinon DF6HA-1B lenses. The camera can be mounted closer to the surface and still yield a large enough FOV by using a smaller lens focal length. The camera was set to a resolution of 640x480 pixels with which a maximum frame rate of 450Hz can be achieved. An external synchronising pulse was used to trigger the camera to capture an image. This also enables the synchronization of other measurement devices if needed. Due to the high sampling frequency, it is required that the camera has a short shutter time. This reduces the exposure the camera sensor has to the scene and therefore sufficient light is required. The aperture of the camera was set relatively large to allow high light exposure to the sensor. This decreases the reduces the depth of field. However, since the distance between the camera and the testing surface is constant, a large depth of field is not required. The short shutter time also reduces motion blur on images.

The processing unit used during this investigation was an Intel NUC with a 6th generation i7 quad core/8 thread processor that runs at 2.5GHz (Intel Corporation, 2016) with solid state hard drive and 16Gb DDR4 memory.

6. Results

The performance of the sensor is evaluated by considering the processing time per measurement and the accuracy of the measurement.

6.1. Processing time

Initial performance evaluation was focussed on the maximum obtainable sampling frequency. This was accomplished by increasing the sampling frequency until the computation time became excessive and the algorithm started skipping frames. As a result, it was found that 250Hz was the maximum obtainable sampling frequency with no frame skipping. Setting the sampling frequency to 300Hz resulted in an average sampling frequency of 283Hz and increasing it further to 350Hz resulted in an actual sampling frequency of 302Hz. Although higher sampling frequencies were obtained, they are not consistent and the skipped frames create ambiguities in the measurements. This sensor matches the 250Hz sampling frequency of the commercial sensor, the Correvit S-HR (Kistler, 2016).

6.2. Accuracy

The accuracy of the measurement is determined by setting the side-slip angle to a predefined value on the gimbal setup and then driving in a straight line. The measured side-slip angle is then compared to the known side-slip angle. The sampling frequency was set to 250Hz which is the maximum obtainable sampling frequency. In these tests, the vehicle was accelerated from rest up to 60km/h. Table 3 shows the results where the side-slip angle was set and compared to the average side-slip angle that was measured for that run. The STD for each run is also added as an indication of the noise levels obtained. The results show that a maximum of 0.46deg error in the side-slip angle was obtained throughout testing, although most measurements were considerably closer with a mean error of 0.25deg. The mean STD is 0.56deg, and did not exceed 0.7deg. It should be noted that the testing surface was not completely flat and at speed, caused motions on the vehicle which could explain some of the noise experienced.

Table 3: Side-slip angle results on flat concrete at 250Hz

Side-slip angle [deg]	Run	Mean [deg]	Standard deviation
0	1	-0.0671	0.4621
	2	0.0598	0.5030
	3	0.3620	0.4990
5	1	5.053	0.6144
	2	5.399	0.5024
	3	5.10	0.6914
10	1	10.46	0.5873
	2	10.43	0.5742
	3	10.06	0.6507

Figure 5 shows a test where the side-slip angle was set to 10deg and (a) shows the size of the motion vector (D) in *pixels*. Coincidently D correlates well with the actual vehicle velocity, in km/h, that was obtained from the vehicle‘s speedometer and was used as an approximation towards the vehicle velocity. This assumption is only valid for this specific setup and does not generally hold. For future investigations other means of vehicle velocity are required. Figure 5 (b) shows the number of features that were successfully tracked. The decrease in this number will be discussed in the following section. Figure 5 (c) shows the side-slip angle measurement which measures 10deg as expected.

6.3. Discussion

A commercial sensor, such as the Correvit S-HR, was not available during this investigation and no direct comparison could be made. However, these results are comparable to that of Botha and Els (2015) who used a similar experimental setup on the same test surface (refer to Table 1) for both the DIC methods and the Correvit S-HR sensor. Comparing the results to those measurements, it is evident that similar accuracies were obtained proving that the camera based sensor outperforms the commercial sensor. Lower noise bands are

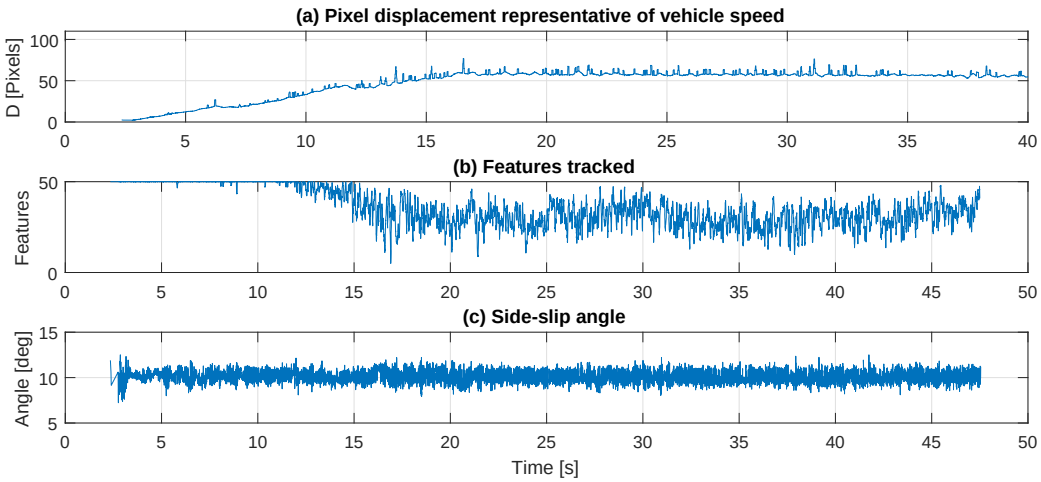

Figure 5: Results obtained with side-slip angle set at 10 deg

also observed.

For the accuracy measurements, the vehicle velocity was limited to 60km/h. However, to further test the capabilities of the sensor tests were conducted at higher speeds. It was found that measurements could still be made up to 100km/h which was the maximum obtainable speed for the test vehicle on the test surface. This is also far better than required for most off-road testing on rough terrain. Figure 6 shows a test where the side-slip angle was set to 5deg and the vehicle accelerated from rest to 100km/h. Figure 6 (a) shows the D and therefore an estimate of the velocity as explained. Figure 6 (b) shows the number of features tracked during the measurement. As the vehicle velocity increases, the ability to track features also decreases. At 100km/h it appears only a small portion of the 50 features are correctly tracked. Although only one point is required to calculate the side-slip angle, the robustness is drastically reduced.

Even at the lower speed of 60km/h, where the accuracy tests were conducted, a reduction in features tracked occurs although not to the extent of 100km/h. There are three possible causes for this phenomenon. The first is blurring on the images caused by the high velocity of the vehicle. This could be reduced by decreasing the shutter time and decreasing the aperture but will require better lighting. The second is the assumption of small motion made by the LK tracking algorithm being violated. The motion can be reduced by increasing the

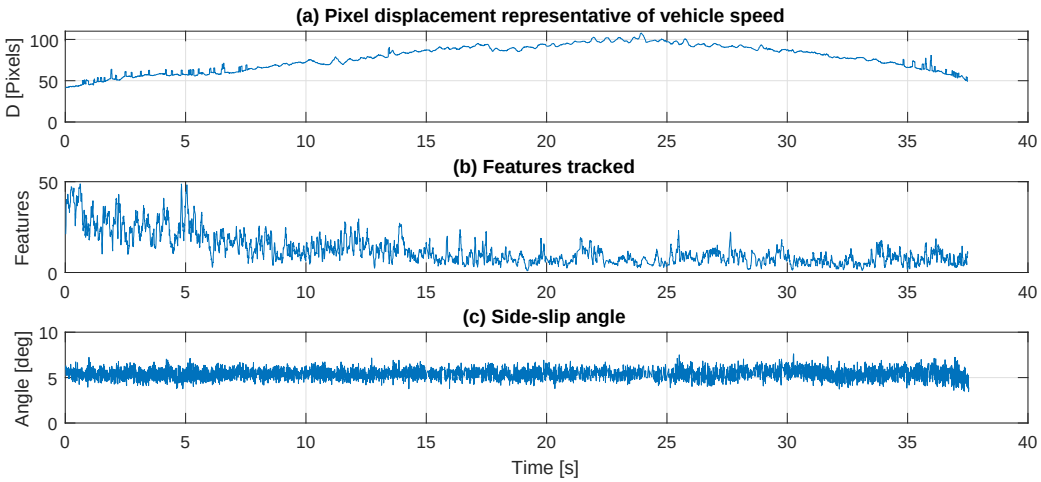

Figure 6: (a) Pixel displacement which represents velocity (b) the number of features that were successfully tracked (c) the measured side-slip angle which was set to 5 deg

distance between the camera and the ground. This would also reduce blurring caused by the motion, however, the accuracy at lower speeds will be reduced due to the larger field of view. Variable focal length can also be used to achieve the same effect. The focal length can be adjusted either manually or automatically based on the velocity at which the test occurs. The third possibility is the lack of overlap between frames. At 100km/h, there is 110mm displacement between samples, resulting in a 70% overlap between images. By rotating the camera 90deg such that the features are tracked in the width rather than the length, the overlap is increased to 78%. The overlap is important as only features in this region can be tracked. If features are chosen where there is no overlap, the feature cannot be tracked, and the feature is therefore lost.

As mentioned before, the pixel displacement was used as an approximation towards the vehicle velocity. If a calibration is done to find the correlation between the vehicle velocity and D, the exact vehicle velocity can be measured by the sensor in addition to the side-slip angle. This could supply ADAS with an alternative vehicle velocity measurement during dynamic manoeuvres where the velocity determined by the wheel speed is complex and generally unreliable such as during ABS braking.

7. Conclusion

From the results obtained it is clear that the real-time implementation of the camera-based side-slip angle sensor was successful. A maximum sampling frequency of 250Hz was obtained. The mean error across measurements was 0.25 deg with a mean STD of 0.56 deg. Although no direct comparison can be made between the camera based sensor and the commercial sensor, a comparison can be made to previous results obtained that used the same experimental setup. This proved that the camera-based sensor outperforms the commercial sensor in accuracy and with lower noise bands. It also shows that, although no tests were conducted on rough off-road terrains, the sensor will still outperform the commercial sensor on these rough terrains since the algorithms are the same. Rough terrains also allow better features to be identified and subsequently are easier to track. Supplementary tests showed that the sensor could measure up to 100km/h which is close to highway speeds. This is also far better than required for most off-road testing on rough terrain, showing excellent capabilities. In conclusion, the camera based sensor successfully provides a cheaper alternative to current methods by using inexpensive, off-the-shelf cameras with dedicated software.

Although the investigation was deemed successful, various improvements can be made. Firstly, the cause for the reduction in features tracked at higher speeds can be identified. This will improve the accuracy of the measurements at high speeds and possibly allow for even higher operating speeds. The overlap can be increased by rotating the camera as to allow the features to move along the width of the camera which has a larger field of view. A variable focal length lens can be used to obtain better accuracy at both low and high speed. The focal length can change depending on the vehicle velocity obtaining FOV that is optimal for each vehicle speed. The sensor can be calibrated to determine the exact relationship between the actual vehicle velocity and D (making the assumption that the sensor is fixed to the vehicle and its position will not change), the system could provide the velocity of the vehicle independent of other measurements.

Acknowledgements

The research leading to these results has received funding from the European Union Horizon 2020 Framework Program, Marie Sklodowska-Curie actions, under grant agreement no. 645736.

References

Bakker, E., Nyborg, L., Pacejka, H., 1987. Tyre modelling for use in vehicle dynamics. In: SAE Technical Paper. SAE International, Detroit, MI, USA.

Bay, H., Ess, A., Tuytelaars, T., Gool, L. V., 2008. Speeded-Up Robust Features (SURF). Computer Vision and Image Understanding 110 (3), 346–359.

Bevly, D. M., Ryu, J., Gerdes, J. C., 2006. Integrating INS Sensors With GPS Measurements for Continuous Estimation of Vehicle Sideslip, Roll, and Tire Cornering Stiffness. IEEE Transactions on Intelligent Transportation Systems 7 (4), 483–493.

Botha, T. R., Els, P. S., 2012. Vehicle Sideslip Estimation Using Unscented Kalman Filter, AHRS and GPS. In: ASME 2012 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference. pp. 651–659.

Botha, T. R., Els, P. S., 2015. Digital image correlation techniques for measuring tyre-road interface parameters: Part 1 - Side-slip angle measurement on rough terrain. Journal of Terramechanics 61, 87–100.

Bouguet, J., 2000. Pyramidal implementation of the Lucas Kanade feature tracker. Intel Corporation, Microprocessor Research Labs.

Chung, T., Yi, K., 2006. Design and Evaluation of Side Slip Angle-Based Vehicle Stability Control Scheme on a Virtual Test Track. IEEE Transactions on Control Systems Technology 14 (2), 224–234.

Dora, K. B., Karthikeyan, S., Rajasekaran, T., Pacejka, H. B., 11 2006. Design and development of innovative tyre test facilities for measuring tyre characteristics. In: SAE Technical Paper. SAE International. URL https://doi.org/10.4271/2006-32-0028

Fischler, M. A., Bolles, R. C., 1981. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. Communications of the ACM 24 (6), 381–395.

Hac, A., Simpson, M. D., 2000. Estimation of Vehicle Side Slip Angle and Yaw Rate. In: SAE Technical Paper. SAE International.

Inagaki, S., Kushiro, I., Yamamoto, M., 1995. Analysis on vehicle stability in critical cornering using phase-plane method. JSAE Review 16 (2), 216.

Intel Corporation, 2016. Intel NUC Kit NUC6i7KYK Features and Configurations . [Accessed: 2016-09-10].
URL https://www.intel.com/content/www/us/en/nuc/nuc-kit-nuc6i7kyk-features-configurations.html

Kistler, 2016. Correvit©s-hr sensors. [Accessed: 2016-09-10].
URL https://www.kistler.com/?type=669&fid=67238

Leutenegger, S., Chli, M., Siegwart, R. Y., 2011. BRISK: Binary Robust Invariant Scalable Keypoints. In: Proceedings of the 2011 International Conference on Computer Vision. ICCV '11. IEEE Computer Society, Washington, DC, USA, pp. 2548–2555.

Lowe, D. G., 2004. Distinctive Image Features from Scale-Invariant Keypoints. International Journal for Computer Vision 60 (2), 91–110.

Lucas, B. D., Kanade, T., 1981. An Iterative Image Registration Technique with an Application to Stereo Vision (DARPA). In: Proceedings of the 1981 DARPA Image Understanding Workshop. pp. 121–130.

OpenCV, 2017. OpenCV library. [Accessed: 2017-09-01].
URL https://opencv.org/about.html

Rajamani, R., 2005. Vehicle Dynamics and Control. Springer.

Rublee, E., Rabaud, V., Konolige, K., Bradski, G., 2011. ORB: An Efficient Alternative to SIFT or SURF. In: Proceedings of the 2011 International Conference on Computer Vision. ICCV '11. IEEE Computer Society, Washington, DC, USA, pp. 2564–2571.

Shi, J., Tomasi, C., 1994. Good features to track. In: 1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition. pp. 593–600.

Stewenius, H., Nister, D., Kahl, F., Schaffalitzky, F., 2005. A minimal solution for relative pose with unknown focal length. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). Vol. 2. pp. 789–794.

Zhang, Z., 2000. A Flexible New Technique for Camera Calibration. IEEE Transactions On Pattern Analysis And Machine Intelligence 22 (11), 1330–1334.