

Academy Report ACA022

Next Generation Monitoring Systems

N. Dhillon, A. Hannay, R. Workman

Report details

Report prepared for:	TRL Academy		
Project/customer reference:	SIP20009		
Copyright:	© TRL Limited		
Report date:	January 2022		
Report status/version:	Final		
Quality approval:			
Indi Gill (Project Manager)	Jan 2022	Alex Wright (Technical Reviewer)	Jan 2022

Disclaimer

This report has been produced by TRL Limited (TRL) .

The information contained herein is the property of TRL Limited. Whilst every effort has been made to ensure that the matter presented in this report is relevant, accurate and up-to-date, TRL Limited cannot accept any liability for any error or omission, or reliance on part or all of the content in another context.

Table of Contents

Executive Summary	ii
1 Introduction	1
2 Traffic-speed survey vehicles	2
3 An alternative, lower cost survey system	4
4 Components of an alternative, lower cost survey system	6
4.1 Data acquisition unit (DAU)	6
4.2 Wireless communications	6
4.3 Location referencing	7
4.4 Camera	7
4.5 Inertial sensors	8
4.6 Laser sensors	8
5 Investigation of individual components	10
5.1 Approach	10
5.2 Raspberry Pi Data Acquisition unit (DAU)	11
5.3 Camera	12
5.4 IMU	17
5.5 GPS	17
5.6 Wireless Communications	27
6 Integration and testing	29
6.1 Power management	30
6.2 Data storage	30
6.3 Housing and mounting the system	31
6.4 Control and processing software	31
7 Network trials	33
7.1 UK local roads	33
7.2 Network trial in the Global South	39
8 Discussion and Conclusions	48
9 Recommendations and Future Use	50
Appendix A World Bank IRI Visual Scale	52

Executive Summary

In the UK, survey vehicles, operating at traffic-speed, are deployed across the road network to assess the condition of road pavements. These apply high-quality (and high cost) equipment to measure surface shape and visual condition to provide a high-resolution assessment of the surface. Due to their cost and complexity these surveys are undertaken infrequently and using a small fleet of specialist vehicles. The data provided is well suited to strategic asset management, but does not provide up to date information on emerging defects, for which road authorities tend to rely on driven or walked manual inspections.

Over the last few years significant progress has been made in the development of low-cost sensors and data collection units that have potential for application in highways. This project aimed to understand the capabilities of this emerging technology, and whether it could provide useful information on road condition. The project explores technologies for: data acquisition and management of sensors; data delivery, both to on board storage and wirelessly to the cloud; GNSS for location referencing; digital imagery for visual condition; inertial measurements for pavement roughness; and compact laser systems for pavement shape at high resolution. A broad review of options led to the selection of a Raspberry-Pi based Data Acquisition System. This is complemented by a compact camera, GPS and inertial measurement system, which together form a portable low-cost system. Data delivery is achieved using on-board storage, Wifi and 4G GSM. Finally, an external low-cost Solid State LiDAR system provides high resolution measurements of shape. The total cost is a few hundred pounds.

Laboratory trials are used to characterise and refine a prototype system deploying the above. The trials highlight the limitations of current solid state LiDAR sensors in this application, and suggest that further development will be needed before this can be routinely deployed. However, the remaining sensors show strong potential for use in road condition assessment, providing reasonably accurate location referencing and good quality images in which road defects can be identified. The inertial system is shown to be capable of measuring road roughness to a level that would support broad condition assessment.

A wider trial of the prototype system in a potential application – the measurement of roughness (IRI) on global south road networks – was carried out in El Salvador. The trials compared the prototype with other devices based on smartphones. The prototype delivered IRI to a similar range and performance to the smartphone-based approach, with the added ability to collect images, the potential to add additional sensors and provide real-time reporting via the cloud. The prototype therefore showed comparable performance, combined with higher levels of practicality and capability, and the potential for higher levels of consistency through a common low-cost measurement platform.

In the light of this research it is felt that the initial application for the device would be aimed at condition surveys in global south nations (which are unable to deploy the complex equipment required to carry out network level condition assessments, and hence rely on incomplete or inaccurate data to make decisions on maintenance). It is suggested that, following refinements to the prototype, it could offer a network assessment tool that could be deployed to these nations as multiple low-cost surveying units providing semi-autonomous, real time, network wide measurements of roughness, supported by imagery for assessment of visual condition.

1 Introduction

Currently, in the UK, surveys are carried out over the strategic and local road networks to assess the condition of road pavements using traffic speed survey vehicles, according to the TRACS and SCANNER specifications¹. These surveys deploy high-quality (and high cost) equipment (lasers, cameras, inertial units etc.), which measure surface shape and visual condition to provide a high-resolution assessment of the surface. Due to the cost and complexity of these surveys they are undertaken infrequently (typically annually) and using a small fleet of specialist vehicles. Whilst the data provided is well suited to the application in strategic asset management (e.g. to obtain condition indicators and to identify locations in need of further investigation) it does not support operational needs – for example providing up to date information to highway engineers on defects that emerge over a short period of time and may need immediate attention. For this application road authorities tend to rely on driven or walked manual inspections.

Over the last few years there has been significant progress in the development of low-cost, high-speed, sensors and data collection units that have potential for application in Highways. This project has aimed to understand whether these systems could offer a viable alternative to the higher cost tools. If so, such small and low-cost units could be attached to vehicles to collect and report condition data on a more frequent basis, providing highway engineers with information to support both network asset management and operational requirements.

The aim this project has therefore been to understand the capabilities of emerging low-cost sensors and data collection units, and whether they could provide information on road condition. This has been used to support the further development and assessment of a potential low-cost modular surveying system and their potential uses.

¹ TRAFFIC Speed Condition Surveys and the Surface Condition Assessment of the National NETWORK of Roads

2 Traffic-speed survey vehicles

On the strategic and local road networks in the UK, survey vehicles are utilised to provide data in accordance with the TRACS and SCANNER specifications (for strategic and local roads respectively). These are bespoke devices that integrate several sensors mounted permanently to the vehicle (Figure 1). The raw data is collected at traffic-speed and passed through bespoke algorithms to deliver condition parameters relevant to the assessment of road surface condition (e.g. measuring texture depth, roughness, rutting, cracking).



Figure 1: HARRIS 3 – Highways England TRACS audit vehicle

The sensors fitted to these devices typically include:

- Location referencing equipment consisting of a wheel encoder (Distance Measurement Instrument, DMI) that generates pulses at set distance intervals (for accurate linear referencing (i.e. along the road) and to trigger the sensors) and a GPS for geographical location (inertially assisted differential GPS is commonly applied to achieve accuracies of better than two metres)
- Cameras (a combination of those pointing down towards the pavement surface to carry out highly detailed defect analysis, and those to view the surrounding assets)
- Inertial Measurement Units (IMU - accelerometers and gyros which can enhance the GPS position while also providing information about the vehicle's movement to support the measurement of road shape (profile))
- Displacement Lasers to measure the surface shape (profile) at high resolution transversely and longitudinally, including pavement texture. The measurement of shape can also be achieved using 3D imaging sensors. Note that scanning laser systems deploying LiDAR could also be applied for the measurement of shape. However, most LIDAR systems are unable to meet the accuracy and/or resolution

requirements of the TRACS and SCANNER specifications and therefore LiDAR is not deployed for the measurement of shape in current vehicles. However, LiDAR can be used to providing data on the area surrounding the vehicle (e.g. distance to bridges, barriers etc., but this is not a compulsory requirement of SCANNER or TRACS.

The sensor package installed on these vehicles typically costs several £100k. The sensors pass their data to a Data Acquisition Unit (DAU), which controls the sensors, undertakes any pre-processing and delivers the data to the on-board storage. The collected raw data is then post processed to produce the condition parameters, such as the level of roughness, depth of rutting or the extent (area) of cracking present.

3 An alternative, lower cost survey system

As an alternative to the bespoke arrangements provided on the “high-end” systems deployed for network condition surveys, this work has considered whether low-cost (under £200 per sensor) sensing and data acquisition systems could be implemented to provide information on the condition of highways assets. An outline set of requirements was established for the system, including:

- A small and portable form factor providing the ability for the system to be fitted cheaply and quickly to vehicles and transferred between vehicles
- The use of compact sensors providing data to support the initial target application (e.g. understanding condition), but with the ability to attach additional sensors for future upgradeability
- Simple operation by the user, requiring minimal specialist skills
- Ease of data delivery. For example the deployment of simple processes to extract the data after the survey, and ideally the ability to send data to the cloud in almost real time so that the system is essentially automatic (minimum operator work)
- The ability to post process data into usable outputs without the requirement for complex third-party software

Based on the above the work focussed on the following items of technology as components of the low-cost survey system:

- A Data Acquisition Unit (DAU) to control the system and store the data from the sensors
- Data delivery to both on board storage (SD card, USB) and to the cloud via wireless communication
- Geographical Location referencing provided via GPS. For linear referencing and sensor triggering it was deemed impractical to deploy a physical Distance Measurement Instrument (DMI) as any system developed within this project would need to be easy to set up and attach to a vehicle – which is complicated if the system must rely on an external wheel-mounted DMI. As it is possible to obtain distance pulses by interfacing with a vehicle’s CANbus, this was also considered. However, it was concluded that interfacing with the CANbus on different types and manufacturers of vehicles can be difficult, and potentially invasive. Therefore, this was not found to be a practical method of getting distance information in this work.
- The collection of images of the forward view (for visual condition / inventory), provided by compact cameras
- Inertial measurement sensors to support the general measurement of pavement shape (roughness)
- Compact laser systems to measure pavement shape at high resolution – to support detailed assessments of pavement condition such as pothole depths

A further constraint placed on the technology was that any system should be small enough to not require a bespoke mechanical framework to attach sensors and should be powered from the vehicle battery without “overloading” the system.

4 Components of an alternative, lower cost survey system

Even within constraints discussed in Section 0, a range of possible low-cost sensors were available for the project. These were researched to understand the current capabilities and costs.

4.1 Data acquisition unit (DAU)

For the required system it was necessary for the DAU to:

- Be physically small and robust
- Have the ability to receive raw data from different sensors
- Have low power consumption
- Have wireless Wi-Fi capabilities

The Data Acquisition Units used in the devices discussed in Section 0 are large and costly. However, in recent years a number of small microcontrollers have evolved, which have operating systems allowing them to be programmed and operated like a standard computer, whilst also having the capability for connections via a pin header (digital I/O, UART, serial, SPI, I²C etc.) allowing interface to other electronic devices. These were the focus of the DAU technology review for this project.

The technology was shortlisted to Arduino and Raspberry Pi, as these are widely used, low cost and with a large community of support. However, there are other similar products available, whose capabilities match the Raspberry Pi (i.e. Banana Pi, Odroid N2, Pine 64, RockPro 64). The Raspberry Pi was chosen for testing as, it has:

- Low power consumption (5V)
- Low cost - around £40 per unit
- A pin header allowing multiple sensors to be attached
- Small physical form
- A flexible Linux operating system
- Can be programmed in C and Python
- Direct storage capability (to “memory cards”)
- Ethernet and Wi-Fi capabilities

At the time this research was undertaken the RPi 3B+ was the most recent edition of Raspberry Pi and was hence the version selected. However, since then a new generation Raspberry Pi (4) has been released.

4.2 Wireless communications

Current surveys (Section 0) do not employ real time data delivery. This is because (i) the data size is large and (ii) for an annual survey real time delivery of data is not a priority. For these systems, data is stored on board, and then taken off the vehicle for processing (which

can take many hours or days for a single survey) before delivery. However, for operational use of the data, real time delivery would be attractive. When designing the system, it was considered that the ability to directly transmit data from the network in real time would be a desirable capability. Wi-Fi does not deliver this requirement, therefore the addition of a 4G dongle was considered.

Standalone 4G boards are now commonplace and SIM cards with good data packages are relatively cheap. However, the system needs to have low power consumption. As there was little difference in the specification of GSM dongles, USB powered 4G systems were selected. Such systems should be able to provide upload speeds of 8 Mbps on average and 50 Mbps peak. 5G was considered, which could provide upload speeds of up to 1 Gbps, but this is not widespread across the UK and was hence not practical.

In addition to this, the Raspberry Pi also can connect directly to Wi-Fi. There is hence the potential to remove/upload survey data when parked and connected to local Wi-Fi, which in general will have higher upload rates than 4G.

4.3 Location referencing

Location referencing is very important to condition surveys and the GPS unit is an essential component of the system. Indeed, as the decision had been made to not deploy a wheel based DMI, the GPS would be the primary device for both geographical location referencing and distance measurement. On current survey systems either the GPS or the wheel encoder can be used as a synchronisation device and to trigger the sensors. It was decided that this would be impractical with the Raspberry Pi. Instead, the GPS would be treated as a separate stream and used to calculate point to point distance. This distance would be fed into post processing algorithms to support linear (distance based) processing of data from other sensors.

A low-cost GPS sensor is generally only accurate to within tens of metres, and it is important to understand knock-on effect of this on the use of the data. Smart phones currently deploy reasonably accurate, low-cost, small GPS receivers as standard. Therefore, these types of receivers are widely available. A number of GPS receivers were investigated, all with the capability of simple data communication with the raspberry Pi. These generally cost around £5-£20 and claim a similar level of accuracy (around 5-15 metres). Low-cost, compact devices capable of producing higher accuracy location data were not readily available.

However, a higher cost external GPS (around £200) could offer great improvement, as this moves the sensitive antenna away from the processing electronics, and the antenna could be placed on the top of the survey vehicle with a better view of the sky. This would also provide WAAS enabled circuitry to (probably) increase locational accuracy to better than 3m.

4.4 Camera

High-end network survey devices collect forward facing images approximately every 5 metres travelled (at around 50 mph this is usually 4 or 5 images a second). High quality, high-resolution, ethernet and USB cameras can be sourced for under £1000 (including lenses). However, to store high quality images at a fast rate there is generally a need to have a very high-powered computer or frame grabber, significantly increasing the cost and size of

the system. There is therefore a trade-off between resolution, capture frequency and data acquisition system cost and size.

One of the benefits of the Raspberry Pi is that it has a dedicated camera port which allows a relatively high-speed capture. There is a limited number of cameras available that use the protocol the port requires (Jetson Hardware). However, Sony have developed Jetson compatible cameras for the Raspberry Pi V2 which cost as little as £20 each.

4.5 Inertial sensors

As noted above, an Inertial Measurement Unit (IMU) collects information about vehicle motion and can be used to understand the road profile (roughness). As with technologies such as GPS, there has been several advances in the development of miniature IMUs due to their deployment in smart phones. Therefore, there is a large range of very low-cost inertial units available, with broadly similar levels of accuracy. The requirement was to ensure that any selected IMU could be easily connected to the data acquisition unit. Due to the range of hobbyists using the Raspberry Pi it was possible to source a proven IMU board with source code. The IMUs were purchased for under £50.

4.6 Laser sensors

Laser sensors offer the potential to measure pavement shape at high resolution, which could provide detailed assessments of pavement condition such as pothole depths. There has been significant innovation in this area in recent years, with the biggest steps being made in low-cost solid-state LiDAR, possibly stimulated by the need for localised scanning systems in automated vehicles. Therefore, it was felt that laser sensing via a LiDAR would offer great potential to a low-cost pavement condition system.

The requirements for a LiDAR to measure pavement condition were that it could be small and easily powered and mounted, whilst also having a practical stand off from the road surface (i.e. the distance from the road at which it could be mounted) and an accuracy suitable for the detection of defects such as potholes. Generally, TRACS and SCANNER measurement systems can measure over a width of around 4 metres and measure up to 1/100 mm accuracy. However, the data that algorithms are applied to is usually reported in 1/10ths of mm. 1 mm accuracy was deemed a more practical target in this work.

It was found that LiDAR units deployed in automated vehicles, although low cost and high quality, have a large measurement range (sometimes up to 100 metres) and a large stand-off (generally not measuring objects closer than 1 metre), and are designed to map the surrounding area to detect large objects – i.e. the measurement of any single data point does not need to be to accurate to the millimetre level. Commercial availability was also unclear (i.e. the ability for the non-vehicle sector to procure in low volume).

An investigation into alternative low-cost systems identified the need to balance the (claimed) number of transverse points collected in each ‘scan’, the width of scan, the stand-off, the accuracy of measurement points and the ability to overcome sunlight. This led to the selection of a Time of Flight (ToF) solid state LiDAR available for under £200, with a claimed specification of:

- 86 degrees view for each scan (one distance measurement per degree)
- Distance range of 0.1 to 0.4 m
- Distance accuracy of 1.5% of distance
- Scan frequency of 10 Hz
- Serial port data output

As seen in Figure 2, the form factor is very compact.



Figure 2: LiDAR system

5 Investigation of individual components

Research was carried out to investigate the capabilities of each system component, to support the selection of the sensors that would be taken forward to the final system, which will be referred to as the *TRL prototype*.

5.1 Approach

5.1.1 Camera

A selection of cameras were trialled, from high-end ethernet cameras (requiring connection to a laptop) to low-cost cameras specifically designed for the Raspberry Pi. The method of testing involved mounting the camera to a vehicle dashboard and recording images (either via a laptop or the Raspberry Pi), which were then visually assessed for their ability to allow pavement defects to be located. In the case of the Raspberry Pi camera, this camera was connected to its dedicated port on the Raspberry Pi.

5.1.2 GPS

Due to the similarity in cost and claimed performance of low-cost GPS units, only two units were purchased (although comparisons were drawn with other work undertaken using tablet and smart phones). These GPS plugged onto the GPIO pins of the Raspberry Pi, requiring only the pins that supply the 5V and the UART RX pin. The GPS data received could be read in a terminal program on the Raspberry Pi, with the coordinates being extracted from the GPGLL string containing latitude and longitude. These were then converted to OSGR (Ordinance Survey Grid Reference) and distance, to assess locational accuracy against reference data such as those collected by HARRIS 3.

5.1.3 IMU

As with the GPS units, the similarity in cost and claimed performance of low-cost IMU boards resulted in two being purchased, although it was later found that the electronics on the two boards were identical.

As with the GPS the X, Y and Z acceleration data could be retrieved via the I²C GPIO pins. The basic method of checking the validity of this data was to sit the Raspberry Pi on each of its sides – top, bottom, left, right, front and back and check that the 3-axis accelerometer data was recording 1 *g* or -1 *g* in each of the positions on one of the axis and zero on the other two.

Both the GPS module and the IMU module were wired onto a single board that could be connected to the Raspberry Pi GPIO so both sets of data could be recorded.

5.1.4 LiDAR

Due to the relatively high cost (between £100 and £1000) there was only scope within this project to purchase and trial one LiDAR. The LiDAR output was provided as serial data when connected to a laptop via a UART to USB interface. Software was written on the PC to

receive the LiDAR data and display it on the PC screen. The LiDAR was tested indoors by measuring (with a ruler) the distance between the LiDAR and a range of objects, and also mounting the LiDAR to on the front of a vehicle pointing down at the road surface and driven at slow speeds along a pavement surface.

5.2 Raspberry Pi Data Acquisition unit (DAU)

5.2.1 Design

The main function of the Raspberry Pi as a DAU was to control the collection of data from the sensors, which were physically attached to its ports, and to store this data in real time.

At the commencement of this work Raspberry Pi Control software was written in C, the language with which the team had most experience from its development of high-resolution survey vehicles such as HARRIS. However, the operating system on the Raspberry Pi is a version of LINUX and it was found after initial C software was developed, that it would be more efficient to use Python, which has an in-built development environment and support source code for the Raspberry Pi (especially in the area of the native Pi Camera).

The aim was to develop a simple control system that would minimise the need for a user to interact with the device (so that it could effectively operate without user interaction). This differs significantly from systems such as HARRIS where an operator starts and stops data collection and continuously monitors the survey. A decision was made to effectively make the device operate autonomously, commencing recording data whenever the DAU is powered on. Data would be split up automatically into runs based on time, i.e. each dataset would last x minutes and a new set of data would start recording as soon as another finished. This would create a more robust system in case any specific dataset failed to save. Hence it would require no operator input at all – just turn on the system and go. This would again set the system apart as even a smartphone-based device would require a user to start up an App in the vehicle. However, as a result of this, the user would not be able to interact with the system at all (e.g. via a screen or keyboard). However, the user does need to know whether the system is functioning. Therefore, a number of LEDs were added to alert the user to the status of the system.

The next question to address was the location and alignment of data. As noted above, it was not feasible to use a DMI for linear location referencing and synchronisation. This would have to be achieved using time – i.e. the DAU's internal clock. Therefore, the data collection component of the control system was designed such that each packet of data from each sensor was saved to a file with a timestamp. The main functions of the DAU software were therefore to:

- Ensure valid GPS is received before collecting data from other sensors
- Ensure each sensor is connected and functioning
- Create a folder and set of files to store data
- Store the data from each sensor in real time
- Close all files and create a new set every x minutes
- Allow the user to press a physical button to safely shut down the system

5.2.2 Challenges

Trials were carried out to support the development, refinement and testing of the control system for the TRL prototype. These addressed a number of challenges.

5.2.2.1 Data storage

The operating system for the Raspberry Pi was on an SD card which claimed to have write speeds of around 10 MB/s. Data was also to be stored directly onto this card. However, the write speed was not achievable, and it was necessary to save data to an external USB drive. This came with its own issues as the Pi port assignment was inconsistent after powering the system on or off. A software solution was developed, but this meant that the external USB had always to be connected.

5.2.2.2 Power

Although the Pi should function adequately with a 5V supply, which can be provided from a car cigarette lighter, due to the current draw the 5V supply would dip sufficiently to cause system instability. The supply input to the Raspberry Pi has some protection on the input which reduces the incoming voltage. It was found that extra current pull on the Raspberry Pi caused voltage drops, triggering the Brown-out protection and leading to a shutdown. When the Raspberry Pi is powering connected sensors this happens more frequently. A power adaptor was developed that could robustly supply the required voltage.

5.2.2.3 Operating system

The Raspbian (the Raspberry Pi version of Linux) operating system is continuously being updated by an online community. This meant that some of the code we had written became deprecated even during the short time of this work, requiring some code to be re-written.

5.2.2.4 Control ports

The Pi comes with several ports on the pin header for the connection of sensors. Sensors are often designed to use I²C or UART, but there are few connections available on the DAU. It was therefore necessary to look at ways in which to collect data via non-preferred ports (i.e. connecting the GPS receiver directly into a USB port).

5.3 Camera

As discussed in Section 5.1.1, trials of cameras involved taking images of pavements to understand the clarity of images and also the size defects that could be seen. Several cameras were investigated - including hi-specification ethernet cameras (similar to those used on HARRIS), and a Raspberry Pi camera. Ethernet cameras were found to deliver high quality images, but these were of a size that could not be practically handled by the Pi, so were not appropriate for this application. Therefore, the most recently available iteration

(V2) of the Raspberry Pi camera (PiCam) was tested², and a number of setting and lenses trialled to understand the effects of these on the quality of the images provided by this camera. To benchmark its performance the Raspberry Pi and PiCam was installed on a high-end survey vehicle (HARRIS3) and collected images at the same time as HARRIS3 collected images using its hi-cost system.

The PiCam was setup to free run, which at the highest resolution gave a frame rate of around 1 frame per second (which is ~1 every 20m at 80km/h, whilst HARRIS3 delivers a frame every 5m regardless of speed). The images were saved sequentially onto a USB plugged into the Raspberry Pi for later inspection. There are a number of different options when capturing images (auto exposure, frame rate, image quality etc.) and as these options were trialled it was possible to produce different outputs depending on what may be required for different tasks (i.e. up to 5 images a second at lower resolution to get an understanding of the surroundings, or 2-3 images a second to allow the identification of major pavement defects). However, it was found that the different lenses made little difference to image quality - there was no notable difference between the Raspberry Pi V2 camera sensor pinhole lens and an upgraded lens in normal daylight conditions.

Figure 3 compares the HARRIS3 and Raspberry Pi images collected at a similar location under similar conditions. The resolution of these images is

- HARRIS3 - 3084 pixels by 1376 pixels (x,y)
- Raspberry Pi - 1920 pixels by 1080 pixels (x,y)

It can be seen that the higher resolution HARRIS images provide greater clarity than the PiCam (for example, within the leaves of the trees) and a wider field of view (highlighted by the red boxes). Although it is possible to increase the field of view of the PiCam by using a wide-angle lens there is a trade-off - to keep the image resolution the same as the smaller field of view requires increasing the image resolution and reducing the frame rate, conversely if the frame rate is kept the constant, the image resolution decreases. Although the resolution is lower, the PiCam image is surprisingly clear, considering the difference in technology and cost between the two image collection systems. However, it is also clear when zooming in to an image that the quality of the HARRIS 3 image is far greater (Figure 4). Further trials were carried out, which focussed on the identification of defects on the pavement surface (which is an expected application of the device). Figure 5 shows images collected using the PiCam in a standard vehicle traveling at up to 70 mph on both local and strategic roads, with defects highlighted. The image quality was found to be adequate to spot major defects within images, while also providing a visualisation of the surroundings. Therefore, given the strong results from the PiCam it was decided that this would be taken forward as the camera used for the system.

² Although the Raspberry Pi V2 camera is considered a vast improvement over the previous V1, since this project started a new Sony image sensor has become available which has higher resolution and larger per-pixel sensor area allowing it to be more sensitive to light and thus producing better images.



Figure 3: HARRIS3 image (top) and PiCam image (bottom) taken at approximately the same location (the PiCam mount was temporary and therefore the angle is not perfectly horizontal). HARRIS 3 has border shaded to match field of view of the PiCam



Figure 4: HARRIS 3 image, top, (from Figure 3) when zoomed in, and PiCam image at maximum zoom



Figure 5: Images from PiCam taken on the A3095 and A329(m) with defects highlighted

5.4 IMU

The selected IMU board was designed to work with the Raspberry Pi. Therefore, installation was straightforward and, once soldered into place, the remaining task was to produce the software required to store the data and test the capabilities of the IMU.

Basic tests were undertaken to ensure that the IMU measured gravity consistently in all 3-axes, and that general movements gave expected responses). The project did not have access to a test platform to check the fundamental accuracy of the device. In addition, for practical purposes the key concern regarding the accuracy of the system is its application in the measurement of road shape (profile), which is discussed below. Therefore, the key initial focus on the IMU data was to ensure that the data was collected and stored by the DAU.

The selected IMU had an upper limit data rate of 400 Hz and collects measurements in 6 “axes” (x, y, z, roll, pitch, heading). However, to successfully and consistently collect and store data from the GPS and camera, it was found that the rate of IMU capture needed to be reduced to around 100-200 Hz. Following a review of the expected requirements for the later calculation of IRI from the inertial data, and to ensure the accurate and consistent collection of gyro (pitch, roll and heading) and image data, a fixed sampling interval was required and therefore IMU data was set to be collected at 100 Hz in 6-axes.

5.5 GPS

5.5.1 *To measure coordinates*

As noted above, initial studies suggest that the horizontal GPS error of tablets and smart phones is ~20 metres on local roads. Tests were carried out using the GPS receiver integrated with the Pi to understand whether similar levels of accuracy would be provided by this system.

There are currently two proposed uses of the GPS:

- i) To spatially tag the location of images
- ii) To calculate chainage based on the distance between GPS points. This will be used as the basis for calculating IRI

To assess the accuracy of the GPS for the tagging of images the TRL prototype was driven along two sites (one shown in Figure 6, approximately 6 km long) three times at 50mph, and another, A329(M), approximately 25 km long, once. Each GPS point (once a second for the prototype) was converted into OSGR grid coordinates, and the horizontal error between each point and the ‘reference’ trace of the route itself was calculated. The reference was obtained by carrying out a survey with the HARRIS3 survey vehicle, equipped with a “reference”³ level GPS, from which reference OSGR coordinates of the routes were obtained.

³ For the purposes of these tests the term “reference level” is used because HARRIS3 is the reference device used for the measurement of coordinates for surveys on the SRN. It has an accuracy of ~1-2m.

Figure 7 shows the cumulative distribution of horizontal errors between the prototype coordinates and the reference. For this dataset ~50% of points are within 2 m of the reference, and 80% of points are within 5 m of the reference. To place this performance into the context of a higher end survey vehicle, it would typically be expected that up to 95% of the points would fall within 2 m of the reference on a site such as the A3095. However, the 5 m achieved by the prototype suggests that the location data, if used to label images collected at the location at which the coordinate was collected, should be adequate for the location of pavement defects. This is in the light of the how the device would be expected to be used - to understand the overall condition of a route or section for initial assessment, with any decision on maintenance subject to a site visit (locating defects to within 5 m and also providing an image would adequately guide an on-site assessor to their location).



Figure 6: GPS trace from low-cost GPS overlaid on a map

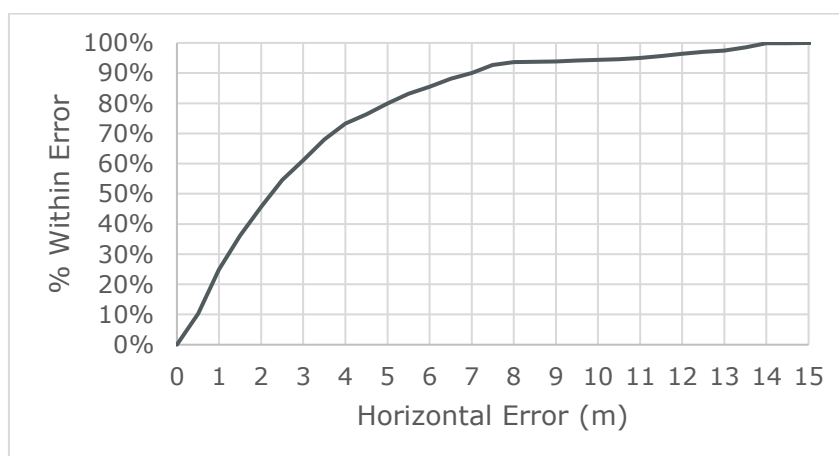


Figure 7: Cumulative horizontal error between prototype GPS and reference on A3095 and A329(M)

5.5.2 To measure distance

As the system has no DMI, and GPS is therefore to be used for distance measurement (in particular to support the measurement of roughness), an assessment of the effect of GPS error on the calculated distance (chainage) was carried out. The method to quantify the errors in the calculation of chainage from GPS was to:

- Align the prototype GPS to a point on the reference trace and reset the chainage on both datasets to 0m at this point
- Calculate the cumulative chainage in the prototype dataset by calculating the distance between successive GPS points. This created a dataset where each point had a chainage and OSGR value
- Calculate the OSGR error between the prototype OSGR and the reference OSGR at each chainage

Figure 8 shows that calculating chainage on the basis of GPS coordinates on the westbound length of the A329(M) has resulted in a cumulative chainage (distance) error that varies along the survey, with an error generally around 5-10 m. However, although there would be localised differences between the true and measured chainage of up to 10 m, this error appears to be “corrected” towards the end of the route. The behaviour seen in Figure 8 was actually somewhat unexpected. We might expect the distance error resulting from the calculation of distance from coordinates to increase during a survey as errors accumulate. The behaviour in Figure 8 suggests a “correction” is somehow achieved in this survey, after 12,000 m.

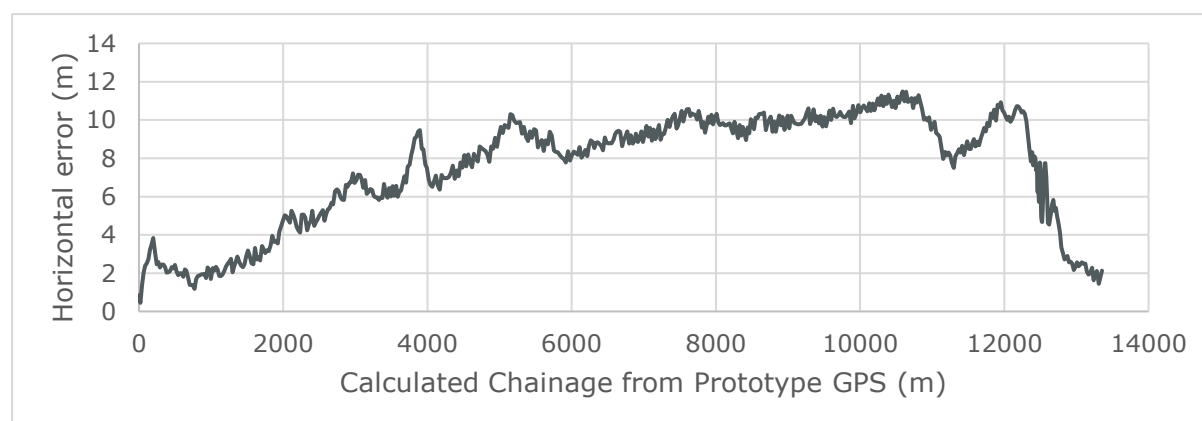


Figure 8: Horizontal error between prototype GPS and reference GPS points when comparing points at the same calculated chainage (Westbound)

We can see in Figure 9 that the distribution of errors peaks at around 10 m, and also around 2 m on this route, showing a bimodal distribution. To understand this behaviour we have plotted the reference and prototype OSGR points at the chainages corresponding to the start and end of the survey in Figure 10. The horizontal errors are only a couple of metres, in line with the errors seen at the start and end of Figure 9. However, we also show the points at 10,761 m, where the error is around 11 metres.

The decrease in error at around 12,300 m seen in Figure 8 corresponds to the vehicle entering and exiting a roundabout. Figure 11 shows the trace of the prototype at the roundabout. The first point plotted in both the reference and prototype occur at the same calculated chainage - the prototype is 'ahead' of the reference (the error is approximately 8 metres). The prototype records a GPS value every 1 second. As the survey vehicle traverses the roundabout, the gaps between measured coordinates results in the prototype measuring a shorter distance than the reference (which is accurately tracking the full roundabout). The result is a reduction in inter-coordinate chainages in this section and highlights an inherent weakness in the simple application of GPS for distance calculation on this route.

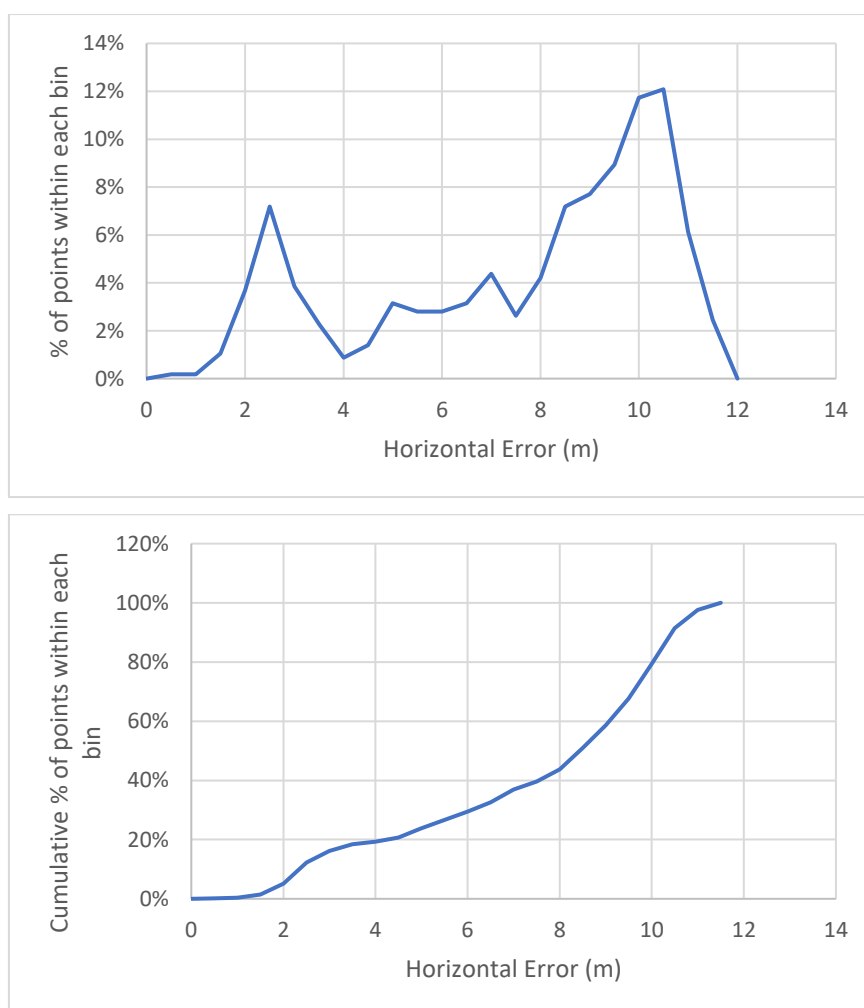


Figure 9: Distribution of horizontal errors in OSGR when comparing the same chainages (top) and cumulative distribution (bottom) on A329(M) Westbound

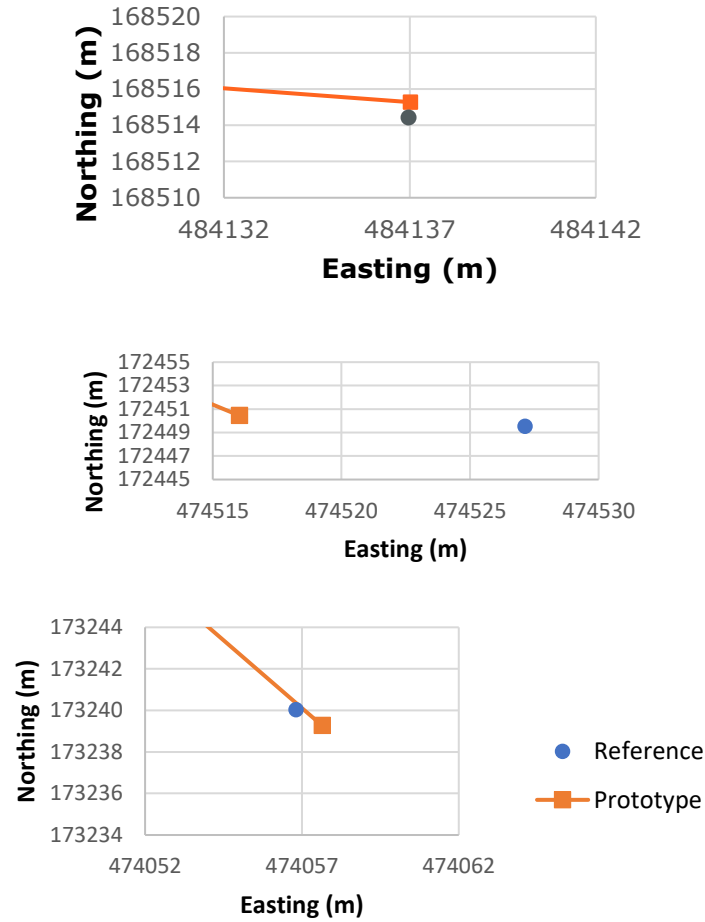


Figure 10: Visual comparison of the reported reference and prototype coordinates at the start (chainage 0 m, top), chainage of 10,761 m (middle, 11 m error), and end (chainage 13,354 m, bottom) of the survey route

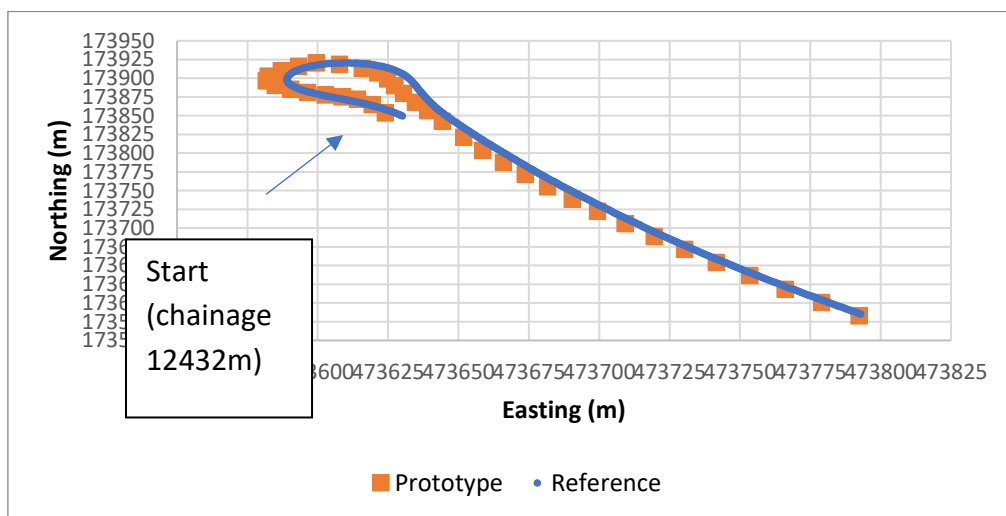


Figure 11: Reference and prototype OSGR showing the vehicle entering and leaving the roundabout

Figure 12 shows the horizontal error calculated (using the same methodology as above), between matching chainage points by the prototype and reference on the eastbound section of the A329(M). Whilst the range of error is similar to the above (~10 m) the anticipated increase in error along the site is demonstrated, with no fall off, suggesting that on such a site (no roundabouts etc.) the error would continue to grow. The frequency distribution (Figure 13) is also smooth.

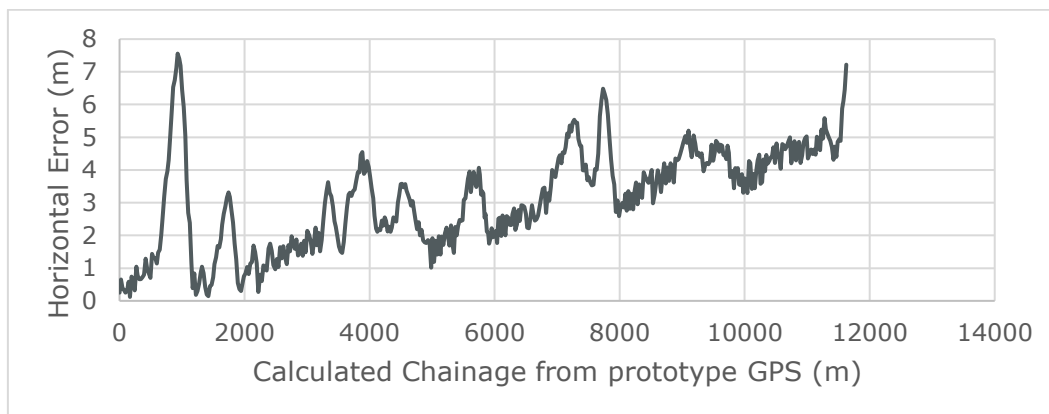


Figure 12: Horizontal error between prototype GPS and reference GPS points when comparing points at the same calculated chainage (Eastbound)

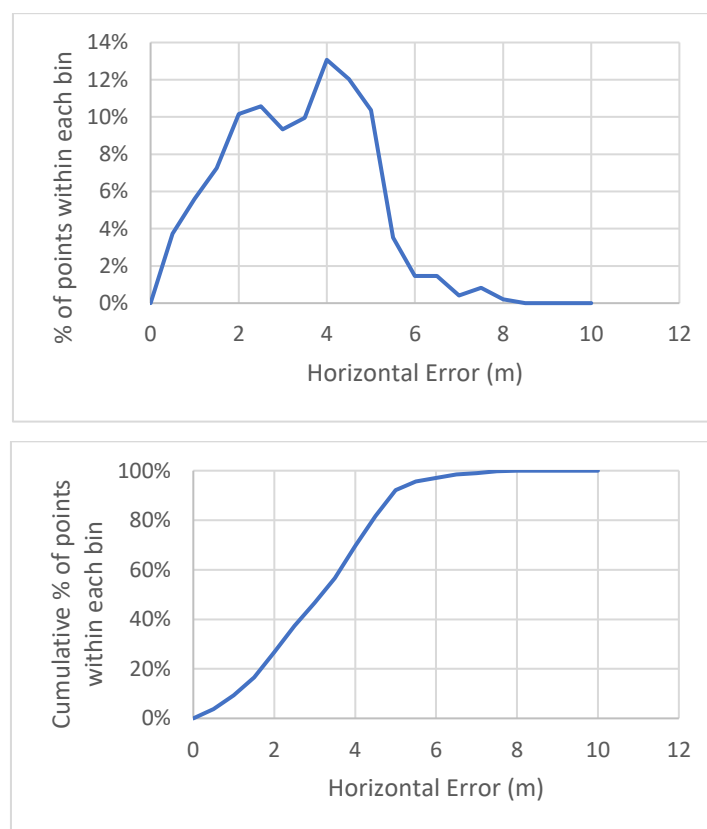


Figure 13: Distribution of horizontal errors in OSGR when comparing the same chainages (top) and cumulative distribution (bottom) on A329(M) Westbound

Two GPS systems were tested in this way, and both gave similar performance. The unit selected for the device was that with the form factor that allowed the easiest integration with the Pi. However, the above observations highlight the limitations of using this simple approach to location measurement (basic GPS) and a simple calculation of chainage from GPS data. Therefore, although not implemented within this project, there are enhancements that could potentially improve aspects of the location measurement, such as:

- Using a Bluetooth GPS receiver, which could be WAAS enabled and also sit outside the vehicle for a better view of satellites
- Integrating the IMU data with GPS data to provide inertially corrected GPS
- Applying post processing rules to the data based on the expected distance travelled between GPS points (i.e. based on expected vehicle speeds some spurious GPS data could be removed from the analysis)

Each of these improvements could not only improve the absolute accuracy of the reporting of coordinates, but also result in a more accurate measurement of chainage (which forms the basis of algorithms such as IRI). Finally, in addition to the standard GPGLL string output from the GPS, the receiver purchased included a 1 PPS (Pulse Per Second) signal. This was used to give a greater accuracy in relation to the time stamping of GPS data, which in turn gives a great level of synchronisation to any other sensors.

5.5.3 *LiDAR*

As noted above, by collecting measurements of the surface shape, it should be practical to better identify operational defects such as potholes, which have a depth of around 10-50 mm. Testing of the LiDAR system focussed on its likely capability to achieve this. One LiDAR system was purchased and was trialled (LeiShen LS02). Software was written to collect and interpret the LiDAR data. This was developed in two stages. The first stage developed code in C on a windows machine for initial tests and to understand the basics of the device. Further code was then written in Python on the DAU to trial the device in representative conditions

A set of indoor tests were carried out to understand the capabilities of the device by measuring the distance to objects. Figure 14 shows the output from the LiDAR in graphical form, the top chart represents a scan over a width of 1.2 metres centred around the LiDAR, the middle graph zooms into the data collected over a width of 0.2 metres (also centred around the LiDAR), and the bottom graph indicates the returned signal intensity (the units are arbitrary for the device). To collect the data in Figure 14 a flat white, but diffuse, object was placed in front of the device at a distance of 0.18 m. The system measures the object as expected. However, the central value is reported as 0 – it appears that the returned intensity of an object directly facing the sensor causes problems. In the middle chart there is a slight step at around +0.1 m. Finally, we can see the higher returned intensities towards the centre of the scan.




Figure 14: Lidar data of a flat object. Zoomed out point cloud (top) zoomed in point cloud (middle) and intensity data (bottom)

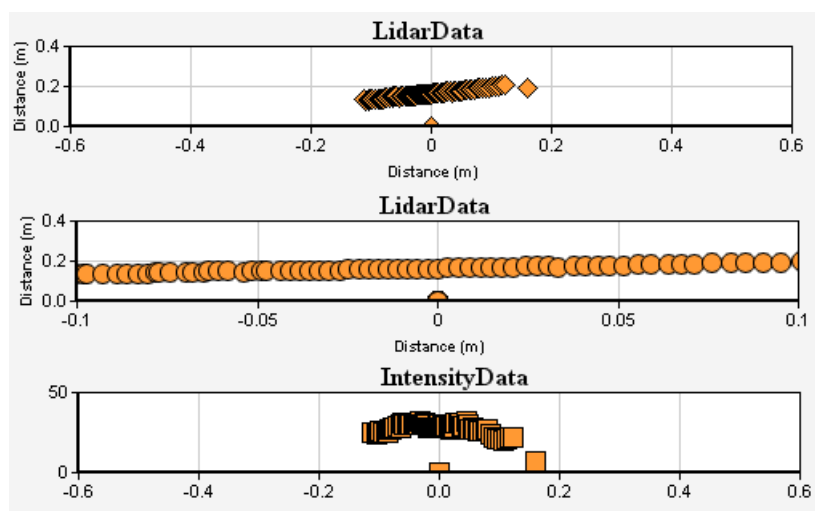


Figure 15: Lidar data of a flat object tilted. Zoomed out point cloud (top) zoomed in point cloud (middle) and intensity data (bottom)

To collect the data in Figure 15 the object measured in Figure 14 was tilted around its centre point. We start to see noise in the data, as the flat object seems to now have a slightly non-flat profile, and its “intensity profile” differs from when the surfaces were aligned.

A convex object has been scanned in Figure 16, and a concave object scanned in Figure 17. For the convex object it can be seen that as the intensity increases (due to the centre of the object being closer to the LiDAR than the edges of the object) dropouts (lost data) occur. In both Figure 16 and Figure 17 we see the profile isn’t smooth, with errors of around 1 cm.

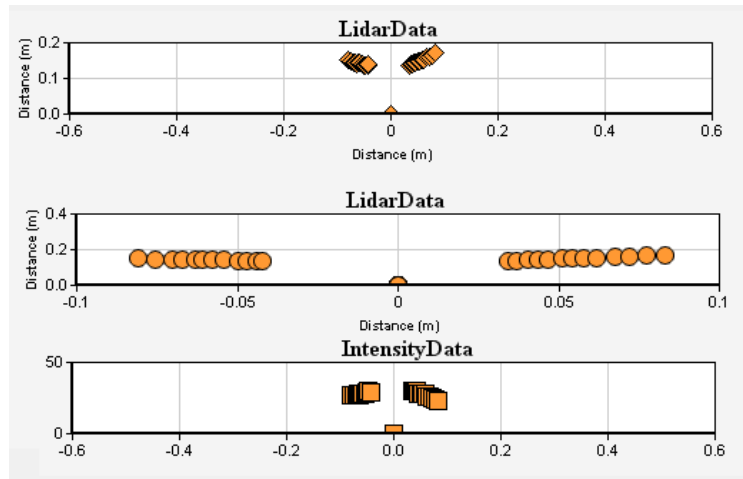


Figure 16: Lidar data of a Convex object. Zoomed out point cloud (top) zoomed in point cloud (middle) and intensity data (bottom)

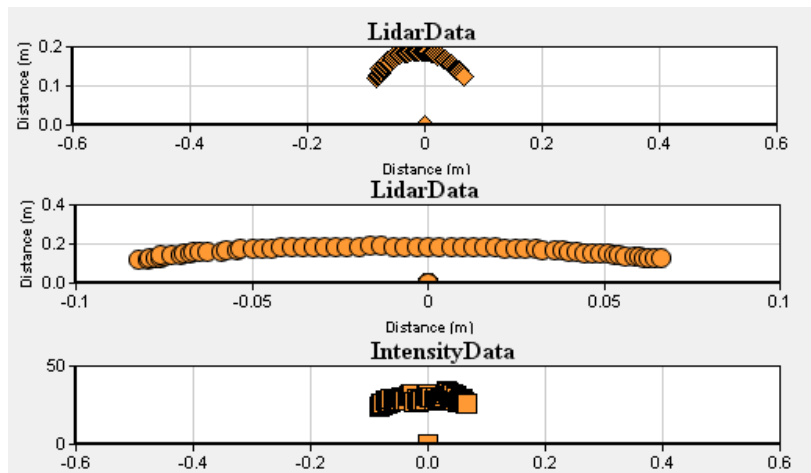


Figure 17: Lidar data of a Concave object. Zoomed out point cloud (top) zoomed in point cloud (middle) and intensity data (bottom)

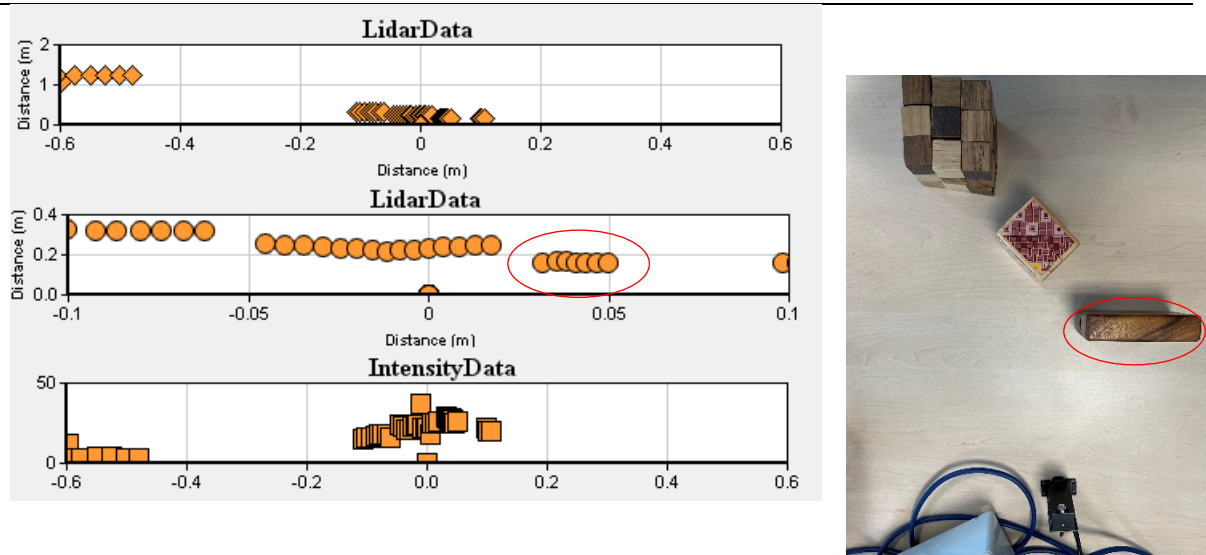


Figure 18: Lidar data of 3 objects. Point clouds (left), and image of the setup (right)

In Figure 18 a test has been setup using a set of blocks as shown to the right (LiDAR sensor at the bottom), with the resulting LiDAR data shown in the graphs. We can also see some of the issues with the LiDAR measurements - the flat object towards the right of the image appears to be slightly curved in the chart (both circled in red).

It therefore appears that even in the well-controlled indoor office environment this LiDAR sensor struggles on specular surfaces (seen in the intensity of the returned signal). There was also a relatively large amount error of in the distance measured to the objects (up to 1 cm when the object was placed at around 30 cm from the LiDAR).

After the initial tests, the LiDAR was mounted onto the HARRIS3 survey vehicle and driven at low speeds along lengths pavement (the LiDAR collects data at 10Hz and therefore driving at slow speed was necessary to try to ensure sufficient pavement coverage to potentially identify defects). The LiDAR data was connected to the Raspberry Pi via USB and was synchronised with other data using timestamps.

Due to the stand-off required to collect valid LiDAR data, it was not possible to mount the sensor underneath the vehicle (where sunlight would be less of an issue), so the sensor was attached to the front bumper and generally within the vehicle's shadow. It was found that, in general, the LiDAR could not resolve distances when subjected to sunlight⁴. As for the indoor tests above, the data was noisy. When attached to the vehicle there were a large number of dropouts and the data itself did not appear to be very repeatable. This, coupled with the fact that we had to ensure that the surface being scanned wasn't exposed to sunlight, meant that it was very challenging to get any data to analyse.

⁴ Note: susceptibility to sunlight was a known issue for all of the low-cost LiDAR systems researched




Figure 19: Two plots of the same LiDAR data, taken over a 20 metre length pavement surface. All values reported in mm

The data from the outdoor tests was extracted from the data files for visual and numerical analysis. Figure 19 shows the noisy data, with dropouts that were exacerbated by the presence of sunlight (many points are reported far above the road surface). The measurements on the road surface are very noisy, and there is banding (on the left chart we see bands of colour). There are a number of points on the right chart which are registered at up to 70 mm above the pavement surface (circled in red) and more than one profile is registered at an entirely different height to the rest of the surface (circled in blue).

Therefore, although it was possible to collect data with the Raspberry Pi and LiDAR, the quality of the data itself was disappointing. Suitable levels of accuracy could only be achieved in perfect conditions. As with many of those researched, the system did not appear to have the capability to overcome sunlight and provide robust data when used outdoors. With the high level of noise present across the data it was not considered feasible to be able to remove this noise reliably while maintaining the 'real' profiles to a level such that defects such as potholes would be reliably detected. However, the trials found that the simple DAU could collect and store the data. If a LiDAR (or suitable distance measurement sensor) can be found with a similar technical specification, but without the performance limitations, there would be a real possibility to detect operational defects.

5.6 Wireless Communications

By incorporating 4G wireless communication there is potential to collect, upload to the cloud, and hence view survey data in almost real time. As noted above, as the specifications and prices for commercial 4G dongles were similar, selection was based on those that had been proven to be usable with the Raspberry Pi. Before trying to upload data to the cloud, a cloud-based storage area needed to be set up. Azure was selected as TRL already had capacity within an existing Azure area.

The objective was to automatically upload each 'run' (as discussed, a run was a dataset collected every x minutes) as soon as it had ended. While the run for the next data was

being collected, the software on the DAU would try to upload data (generally images, GPS and IMU data) to the cloud storage area. A number of 'runs' were carried out to understand how the system performed, each collecting images at 3 Hz, IMU data at 100-200 Hz, and GPS data at 1 Hz. Each run was around 2 minutes long. The software was modified to inform the user when a 'run' had been fully uploaded to the cloud.

It was found that the limitation of the upload was constrained by 4G speeds and the amount of data to upload. Even though the PiCam compresses jpegs to generally under 1 MB per image, when collecting images at around 3 Hz this still meant that a relatively large amount of data needed to be uploaded. As more runs were completed the backlog of data to upload to the cloud became prohibitively large. In some tests it would take up to 20 minutes to upload 2 minutes' worth of data even with a stable 4G connection.

However, although the data could not all be sent, the basic methodology worked. With faster communications (5G) almost real-time data uploading to the cloud would be feasible using this low-cost system. With the addition of battery power backup, it would be feasible to return the system to an office and upload data over Wi-Fi each day.

6 Integration and testing

In the above sections we have presented the background to the low-cost survey device, the selection and testing of the individual sensors. It has shown the following:

- A low-cost Data Acquisition Unit (Raspberry Pi) can be bought, and software written to collect and store data from multiple sensors
 - the system can be powered via the cigarette lighter in a vehicle. However, the voltage needs to be very carefully managed otherwise the system may power down
 - The system can be designed so that no user input is required
- Low-cost cameras have the capability to collect images with high enough resolution, in which significant pavement defects can be seen
- Low-cost GPS receivers can locate position within enough accuracy to generate distance data and to allow mapping (although algorithms would be required to check whether the accuracy of distance could be used for calculating roughness)
- Inertial Measurement Units can provide data at high frequency, which can be stored
- Low-cost LiDAR, although becoming more advanced in its capability and reduced in cost, is probably not at the stage where road profile can be measured within our low-cost target (a few £100)
- A GSM module can be used to automatically upload data to the cloud for storage and processing. However, 4G is not unlikely to be able to sustainably upload data at a fast enough rate to halt a data backlog

Figure 20 shows the prototype at this stage of development.



Figure 20: Original Prototype, comprising DAU, GPS receiver, IMU and camera

The next stage of this project aimed to further develop the prototype into a useable system, overcoming some of the practical challenges that were expected for such a system, including:

- Power consumption and management

- Data storage and management
- Physical robustness (housing and mounting)
- Robust software

Following this, the work aimed to carry out trials of the system and understand if it could be applied to provide data of use to the management of highway assets – in this case the delivery of road roughness data (Section 7.2).

6.1 Power management

As noted above, the initial tests had found that the prototype required a stable continuous voltage. Powering of the system through the cigarette lighter of a vehicle was felt to be too high risk, as when the voltage dropped the system would shut down in an unsafe manner. An unsafe system shut down can damage the Operating System, causing a failure to re-boot on system power up.

A power management system was designed. This utilised a 3.7V Lithium cell, lithium charger circuit, 3.7V to 5V booster circuit, power relay and bespoke software running on the Raspberry Pi (RPI). The only external requirement was a 5V USB power source either from a mains power supply or a 12V to 5V vehicle cigarette lighter socket. The external 5V supply then charges the battery *and* provides power to the Raspberry Pi. If the 5V is removed the system would still operate from the battery via the 3.7V to 5V booster for several hours. In between this battery backed power supply and the Raspberry Pi is an electronic switch that physically switches the power to the Raspberry Pi. Power on is controlled via a physical ON button that controls the electronic switch. Power off is controlled via a physical OFF button connected to the RPi GPIO pins that, via software, allows the Raspberry Pi to execute a safe shutdown service and indicate to the electronic switch to remove power once complete (Figure 21). As a result the prototype can operate for a full day with or without being connected to the cigarette lighter as long as the battery has been allowed to fully charge previously.

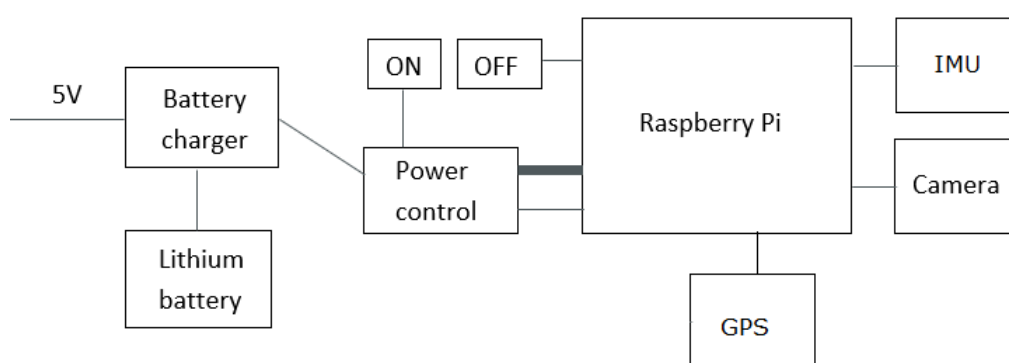


Figure 21: Basic schematic of powering and charging system

6.2 Data storage

In the initial tests all data was stored to the SD card which also contained the DAU Operating System. However, as noted above, it was found that the write speed of the SD card was not

sufficient to save images at an acceptable rate. It was therefore decided that small USB drives should be used to store the data in real time, which also had the added advantage that after collecting data for a number of hours, the system could be powered down, and the USB drive could be swapped, which was not possible with the SD card (as it contained the Operating System and was housed inside the DAU).

6.3 Housing and mounting the system

An off-the-shelf enclosure was selected to house the Raspberry Pi. This has an internal section to mount the Raspberry Pi camera and space underneath to mount other components such as the battery and battery management system. As with all commercially available Raspberry Pi enclosures, they are not designed for mounting outside. The enclosure was modified to facilitate mounting inside the vehicle via windscreen attachments (Figure 22). A benefit of an internally mounted system is that there is no need for waterproofing. However, the angle of mounting is difficult to control, and there is the possibility the GPS antenna will not have a clear view of the sky. In addition, image quality can be affected by reflections from the windscreen.



Figure 22: Raspberry Pi enclosure, and mounting on a windscreen

6.4 Control and processing software

6.4.1 Further development of the control system

Following completion of the initial development, additional software functions needed to be written to allow any sensors to work robustly and to allow the DAU to interface with them. In addition to storing data from the sensors, other required functionality included simple uploading of data to the cloud (4G or Wi-Fi) and a method to provide confidence to the user on the system status (which was achieved using a set of LEDs mounted on the case – visible in Figure 22).

6.4.2 Post processing system

Although sensors had been integrated into the system, and the ability to store raw data demonstrated, the value provided by the system (and the data) is unclear until it has been shown to be capable of providing data of use in highways Asset management (which is the current application of the high-end survey systems discussed in Section 0). This requires the development of post processing tools to convert the data into information. Two areas were considered

- **Visual condition:** To assess visual condition using images the system should provide image data that can be associated with its location on the road network (coordinate, section and distance). An algorithm was developed to create a “route” representing the path driven by the vehicle, by converting the GPS data into Ordnance Survey Grid Reference (OSGR) points and using trigonometry to calculate the distance between points. By using the time stamp the vehicle velocity could also be calculated. As the GPS error can be up to 20 metres, there will clearly be distance errors in comparison to a reference. However, the dataset containing OSGR, time and distance could be aligned with the images provided by the camera to assign each image a real-world position.
- **Roughness:** Nearly all network condition survey vehicles provide a measure of surface roughness. The roughness parameter International Road Roughness (IRI) is an international measure applied worldwide. Within high-end systems this parameter is obtained by fusing laser and IMU data parameters to obtain Longitudinal Profile, to which an algorithm is applied to obtain IRI. Research has shown that International Road Roughness (IRI) can be estimated from IMU data alone (using accelerometer-based systems). Therefore, this was seen as a suitable first test for the Raspberry Pi system. When calculating IRI using accelerometers alone, it is known that the characteristics of the vehicle suspension must be considered, or, that the algorithm must be “tuned” to a reference. TRL has existing processes/code that enable estimates of IRI to be obtained from accelerometer-based data. This was adopted to read the raw IMU data (after a distance had been assigned to each value using the GPS based distance and time stamps) and generate IRI.

7 Network trials

7.1 UK local roads

7.1.1 Test site

Trials were carried out on the network near to TRL to investigate the application of the system to estimate pavement condition. Surveys were carried out on around 50 km of single and dual carriageways around TRL for which reference data was provided by HARRIS3. The sites included the A329(m) (30km), A3095 (6km) and B3430 (14km) as shown in Figure 23.

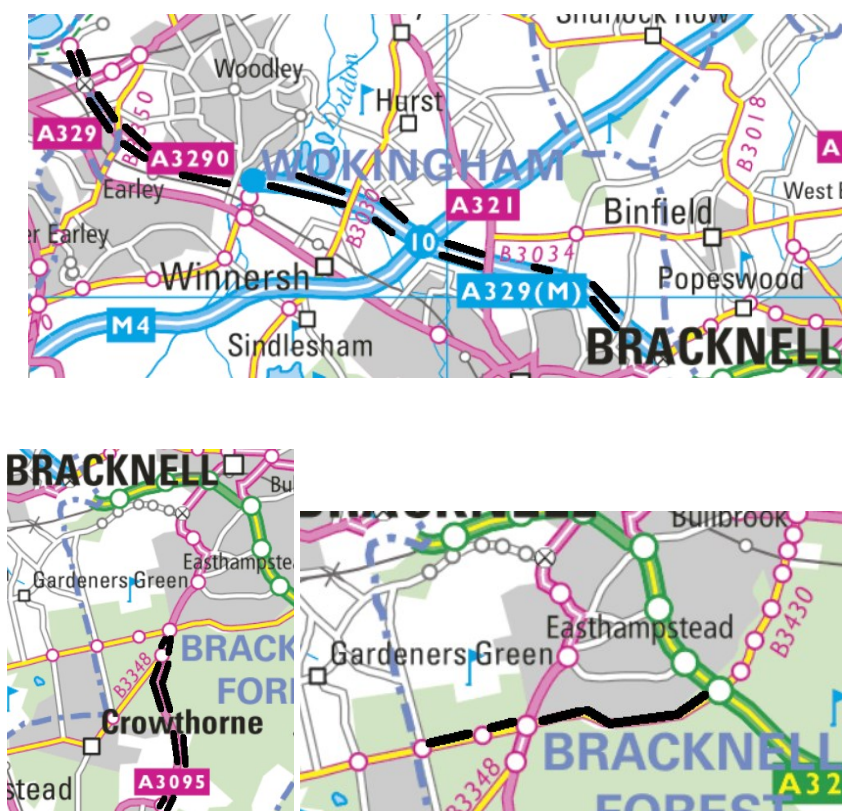


Figure 23: Maps of 3 sites used for UK network testing of system. A329(m) top, A3095 left and B3430 right

7.1.2 Communication / delivery and visualisation of the data

As noted above, although the prototype contains on-board data storage, the ability to deliver the data in real time through wireless (4G) communications was a desirable capability for the system. Therefore for the network trials, work was initiated to develop software that would upload data from the prototype into the Cloud during the survey, which would then be automatically processed through algorithms (to calculate distance, IRI) and then made available for viewing within asset management tools that have access to the cloud storage area to which the data had been uploaded. Although not fully developed, key

components of the software were completed: the data could automatically be uploaded from the device to a cloud storage area, and automatic processing of data files to convert raw data into a usable XML file format (which contained the location of images and a simplified IRI referenced to geographical location and survey chainage) was achieved. The final, uncompleted stages were associated with the asset management tool (comprising the automatic upload of the XML files into the visualisation tool, and the completion of the visualisation tool itself). However, the software did reach a level such that data could be delivered and broadly visualised using prototype asset tools provided by TRL’s software group, as discussed below.

7.1.3 Measurement of location

We can see from Figure 24 that the OSGR from the prototype GPS on the A329(m) match the reference well when viewed at this scale (as in the previous GPS tests, the reference used was the HARRIS3). An analysis on the cumulative chainage errors as described in Section 5.5, shows a distribution of errors as seen in Figure 25 where 50% of the data is within 5m of the reference, and 100% within 12m. We see that there appears to be a bimodal distribution in Figure 25, which is a similar shape to that seen in Figure 9, suggesting similar issues are present on this site associated with coordinates reported close-together on bends, leading to shortened point-to-point distances.

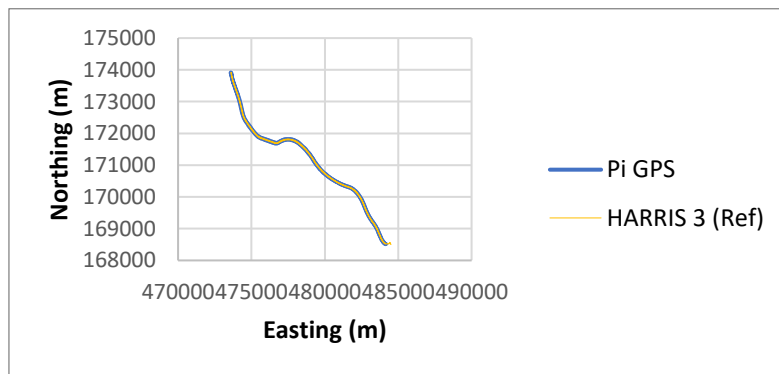


Figure 24: OSGR on A329(M) by the reference device (HARRIS3) and prototype system




Figure 25: Distribution of horizontal errors in OSGR when comparing the same chainages (left) and cumulative distribution (right)

7.1.4 Visual condition

A trial was carried out to visualise the data in an asset management tool. This applied the software discussed in Section 7.1.2, with visualisation provided by TRL's software group. Figure 26 shows an example of an image provided by the prototype from the A3095 site. In this example the data was uploaded to the cloud from the device, and this was accessed by the software for download and display⁵. The collected images could be scrolled through on a map background showing a red dot on the map corresponding to the location of the displayed image.



Figure 26: Visualisation of geo-located image data in TRL software

Figure 27 shows a selection of images from the different trial routes, highlighting that both relatively small and major defects could be successfully identified within the images. This network test confirmed the initial observations that the prototype could provide the facility to collect, deliver and display visual condition data. However, as it was not feasible within the constraints of this project to collect reference visual condition information, the tests did not quantify the “accuracy” of the measurement of visual condition. Large scale trials would be required to obtain a robust understanding of this.

⁵ As discussed in previous sections of this report, although the device has wireless communications, the 4G communications available on the site were not sufficient to enable “real time” delivery of the data during the survey. Hence the data was uploaded from the prototype following completion of the survey.



Figure 27: Images showing low level surface loss (top), sunken patches (middle), short longitudinal crack (bottom)

7.1.5 *Measurement of IRI*

Figure 28 compares the IRI provided by the reference (HARRIS3) and the prototype system on the A329M. Both the reference and the prototype reported IRI values over lengths of 10m, but a 100m moving average has been applied to both datasets to reduce noise⁶. The following observations can be made:

- Although we have seen above that the chainage derived from GPS can introduce errors of ~10m, this does not appear to have significantly affected the alignment of the reference and prototype IRI datasets. No manipulation of the datasets has been

⁶ 100m lengths are typically used to report condition in routine network surveys

carried out, as visual assessment suggested good alignment, with the 100m moving average further reducing the effect of any misalignment.

- In relation to the typical levels used to assess IRI this site has a level of ride quality that ranges from Good to Fair (Excellent would be <math><1.5</math>, Poor is >4 to 5 on paved roads)
- There are a few locations where there are larger differences between the datasets but, visually, a good level of agreement can be seen throughout the site. Note that lengths covering roundabouts were removed from the analysis, as network surveys undertaken with systems such as HARRIS3 do not currently cover roundabouts.
- The survey velocity of the prototype is shown because it is known that this can affect the measurement of IRI by these types of devices. This does not appear to have had a strong influence on the results apart from at the roundabout and perhaps at the start and end of the survey. We would typically expect lengths collected at slow speed (or high acceleration/deceleration) to be removed from the survey data.
- Figure 29 shows that 62% of points fell within $\pm 0.5\text{mm/m}$, and 90% within $\pm 1.5\text{mm/m}$ of the HARRIS3 reference. It is of note that the differences are centred on 0mm/m , which suggests no systematic bias in the prototype.

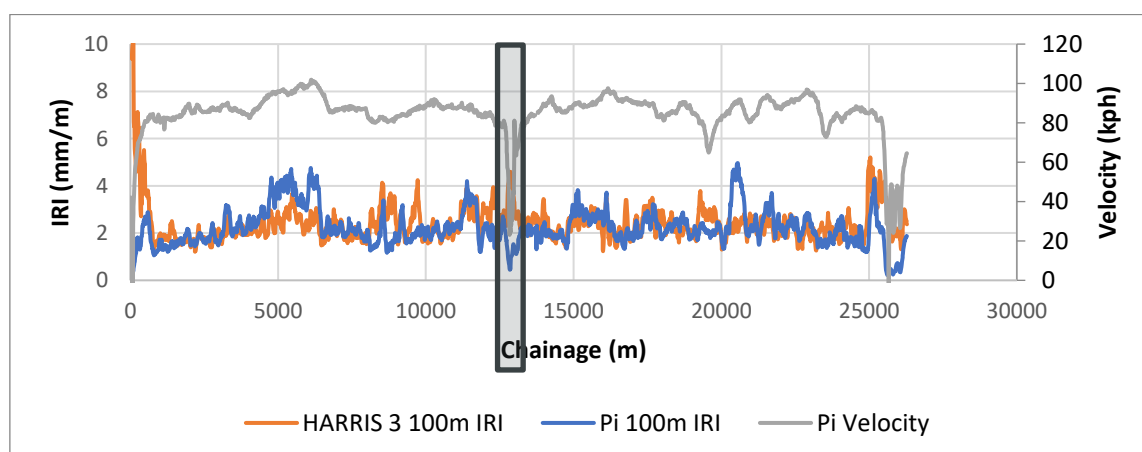


Figure 28:A329M 100m moving average IRI (in mm/m) from reference (HARRIS3) and Raspberry Pi (shaded area is at a roundabout)




Figure 29: Histogram of differences between HARRIS3 IRI and prototype IRI (left), and cumulative (right)

To place the performance in context, longitudinal profile data was obtained for this site from two current network survey devices (HARRIS3 and a commercial survey system) and IRI values calculated. The IRI agree closely (Figure 30) with 90% of differences between the reported IRI values falling within ± 0.5 (Figure 31). As might be expected, the higher-resolution devices give better performance in the measurement of IRI (80-90% within ± 0.5 of the reference compared to 62% for the prototype). This suggests the potential applications for the device would be in the course assessment of roughness at the network level, especially on rougher networks (such as the developing world, forestry roads etc).




Figure 30: A329M 100m moving average IRI (mm/m) from two typical network survey systems

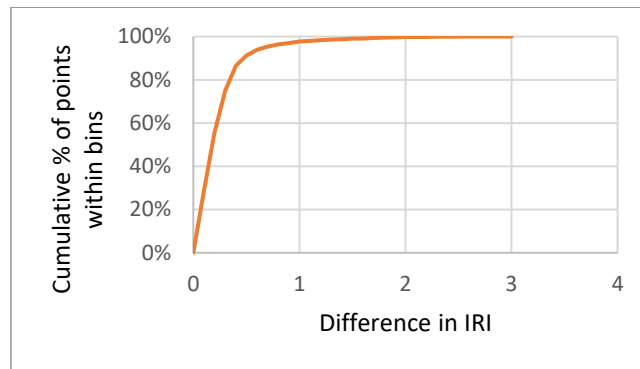


Figure 31: Cumulative histogram of differences between commercial survey devices

7.2 Network trial in the Global South

7.2.1 Test routes

A further network trial of the prototype was carried out as part of an evaluation of road condition on a 292 km length of road, being undertaken by the Millennium Challenge Corporation (MCC), in El Salvador. The road was the Northern Transnational Highway (NTH): an upgraded two-lane transnational highway across northern El Salvador serving the Northern Zone running parallel to the international border with Honduras, see Figure 32. The test network is located in a remote part of El Salvador that is mountainous and difficult to reach.

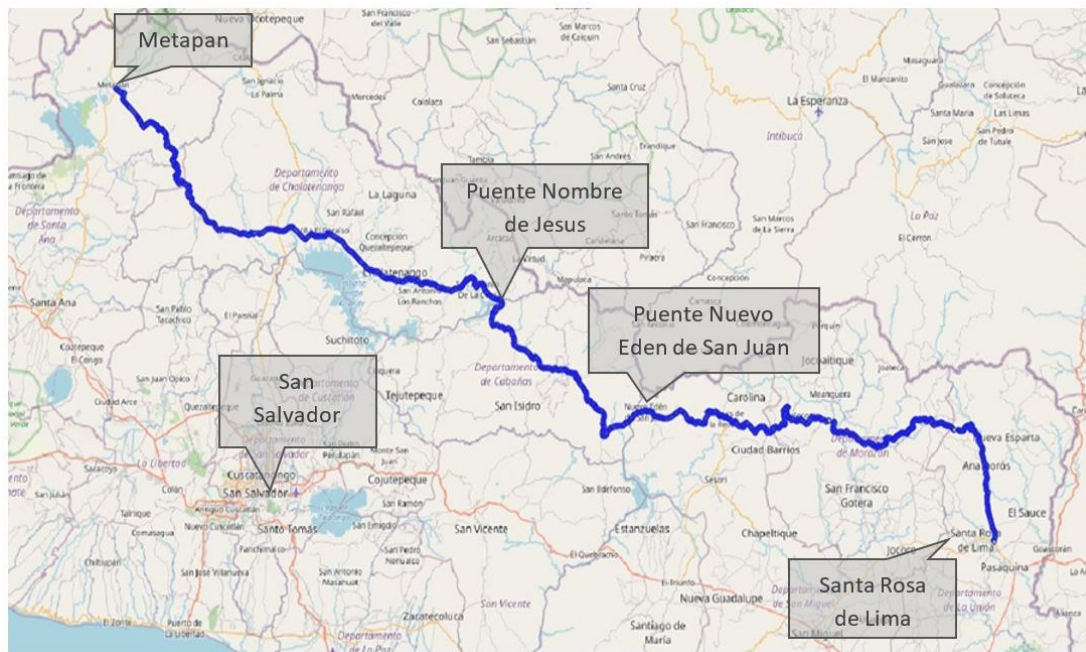


Figure 32: Map of route surveyed in El Salvador

Data collection activities were originally planned for late spring/early summer 2020, but in March/April 2020, all international travel effectively stopped with the outbreak of the COVID-19 pandemic. Many countries have been badly affected by the pandemic, with lockdowns and shelter-in-place orders, in many cases restricting employees to working from home. Therefore, data collection activities could not proceed as planned.

In discussion with MCC, an Evaluation Design Report (EDR) Addendum was approved in October 2020 setting out proposed data collection processes that would allow the work to proceed within the constraints of COVID-19 guidance issued by MCC.

7.2.2 Devices

The updated processes agreed with the MCC provided the opportunity to investigate innovative remote data collection systems in order to remove the need for international travel by the project team. This created an opportunity for the prototype to be applied alongside alternative data collection devices. These alternatives included two smart phone apps and a dashcam (Figure 33). The objective of the data collection exercise was to collect IRI and imagery, from which defects could be identified. All devices were operated by non-TRL staff, who were provided with training remotely. A training course was devised for the TRL prototype and was delivered remotely via communications media (Microsoft Teams).



Figure 33: Setup of the different systems used as part of the trial

7.2.3 Calibration

Although it is possible to report a general measure of IRI based on assumptions regarding the vehicle in which the equipment is mounted, a more robust measure of IRI can be obtained by calibrating the system. For calibration a reference device is used to measure IRI

over specific lengths over which IRI data is also collected using the test device (installed on the vehicle to be used in the full survey). By comparing the IRI measurements, certain parameters used in the IRI calculation in the test device can be adjusted to optimise the accuracy of the IRI measurements.

Hence calibration of each device was carried out locally in El Salvador on an 800 metre long calibration site that was also surveyed by a walking-speed Class 1 instrument called SurPRO (<https://www.internationalcybernetics.com/surpro/>). The test devices were mounted to a vehicle and driven along the calibration site at a range of speeds and data used to tune the IRI algorithms for each device before undertaking the network trial.

7.2.4 Surveys

All devices were installed on a single survey vehicle, which was driven over the 292km network route over a 4-day period. Where possible the driver tried to maintain a velocity of around 50kph. Each device had its own GPS sensor and saved data to its own internal storage, which was then processed using its own software to provide IRI. For the TRL prototype, the data was sent back to TRL in the UK to process and provide IRI and geo-located images.

7.2.5 Results

7.2.5.1 Initial observations / lessons learnt

One of the purposes of carrying out this larger trial was to better understand the usability of the TRL prototype in “real world” conditions. A number of observations were made:

General

- The system was operated and mounted without any issues and the training provided allowed the local team to operate the TRL prototype effectively.
- The two-point mounting connection was relatively stable using suction mounts, but some movement was experienced. The mounting in the vehicle could be improved by refining the connections points.
- In operation, there was some reflection from the windscreen which affected the image quality, despite using dark coverings on the dashboard to minimise this. A polarising filter could be added to improve this. However, the light colour of the prototype case also caused some windscreen reflections to other equipment in the vehicle. If the prototype is to be used inside the vehicle the case should be of a dark or black, non-reflective, material. This also applies to the leads. Alternatively, both problems could be overcome by mounting outside the vehicle. This would also allow closer proximity to the road. There are examples of this type of equipment being attached in this way, and it should be straightforward to implement a waterproof cover.

System operation

- As noted above, the prototype does not produce IRI directly. In the trial the data had to be sent to UK for processing. In the future it would be beneficial if the IRI could be produced as part of the data output, to reduce data handling.
- The TRL prototype provides no visual feedback to the user (it's a "black box"). Mobile phone devices offer more feedback as they have screens. However, providing this additional component would add to the complexity of the TRL prototype and deviate from the original objective for a self-contained automated system. It is possible that functionality could be added so that a third-party device could be linked to the prototype via Bluetooth to check on the data collected.
- The data requirements were high – ~3 images were collected per second and the surveys took 4 days (around 8 hours of driving per day), this totalled around 160GB of data. Because the internet in this remote area of El Salvador was not good it was not possible to use the cloud, so the data had to be removed from the prototype and uploaded every evening. Because of the large number of files, this took some time.

7.2.6 Performance – measurement of IRI

The focus of the performance assessment has been the measurement of IRI, which was the key parameter being assessed for all the devices. Figure 34 compares the IRI reported by the prototype and by the Roadlab device, over the whole route. The survey vehicle velocity is also shown. The IRI has been calculated over 100m lengths and a 10 point (1km) moving average applied. There is a good visual agreement between the IRI reported by the two devices, although the prototype values appear to be lower in general. In a difference analysis we find that 80% of TRL prototype IRI values are within ± 2 of the Roadlab values (Figure 35). However, frequency distributions of the prototype and the Roadlab values (Figure 36) show that the Roadlab device reports a higher value of IRI, biased from the prototype by about 1mm/m. This bias will have reduced the agreement between the devices reported in the difference analysis - removal of the bias could improve the agreement to 80% within ± 1 , which is close to the agreement seen between the prototype and HARRIS3 in the UK trials.

In addition to the bias, in Figure 34 it can be seen that there are localised differences between the prototype and the Roadlab device. While the underlying shape is similar (where there are long lengths of generally higher IRI the devices both report this) each device also reports a number of sharp peaks. These are not reported at the same locations. The reasons for this are unknown. However, it is possible that changes in vehicle velocity could be the cause of some of the differences, if these changes are being handled differently by the IRI algorithms contained within the different devices. The prototype removes values where the vehicle velocity drops below a lower threshold (20km/h). However, it does not filter out lengths where there were high levels of acceleration or deceleration, which can also affect the inertial measurements. The devices may also have different sensitivity to larger bumps. For example, there were speed bumps were present in some locations (Figure 37 and Figure 38), different responses in each system to the presence of such features could lead to differences in the reported IRI. There is also the possibility that the (correct) high

levels of roughness at the speedhumps have been removed from the Roadlab data, potentially under-reporting the roughness at this location.




Figure 34: IRI from Roadlab and TRL prototype, along with vehicle speed, on the Westbound route of the trial site

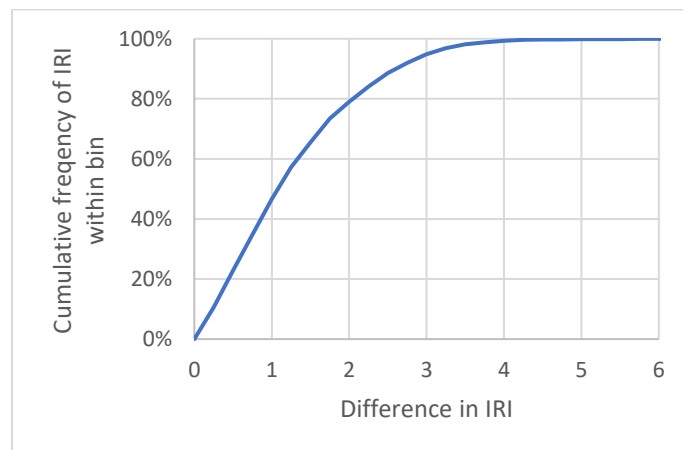


Figure 35: Cumulative distribution of differences between Prototype and Roadlab IRI values

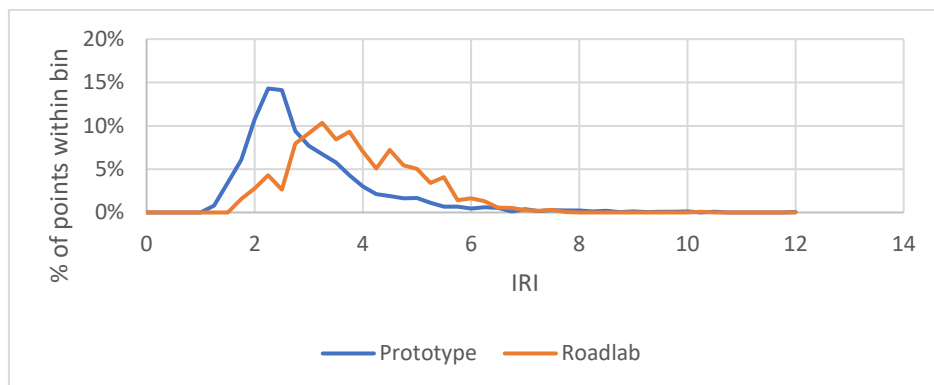


Figure 36: Frequency distribution of IRI from the prototype and Roadlab on the Westbound data




Figure 37: TRL prototype IRI and vehicle velocity highlighting variability in vehicle velocity (around 96000 m)



Figure 38: Speed bumps present where high levels of IRI also reported

7.2.7 Visual IRI comparison

In addition to the reporting of IRI by the devices directly using the inertial measurements, a further measurement of “IRI” was obtained by undertaking a visual assessment of the route against the World Bank IRI visual scale (Appendix A, Sayers et al. 1986). This assessment was carried out in real time by a passenger traveling in the vehicle who was asked to estimate IRI to the nearest unit (m/km (=mm/m)) every kilometre. This measurement is subjective and cannot be considered “reference” level data. However, it can be used as a further check that the inertial IRI measurements reported by a test device have been reported in the expected range.

Figure 39 and Figure 40 compare the “visual IRI” with the inertial IRI reported by the prototype and the Roadlab device for the westbound and eastbound surveys of the test

route. It is immediately clear that the detail and resolution of the visual IRI survey is limited. However, some detail can be discerned, such as the higher levels around 100000m (WB).

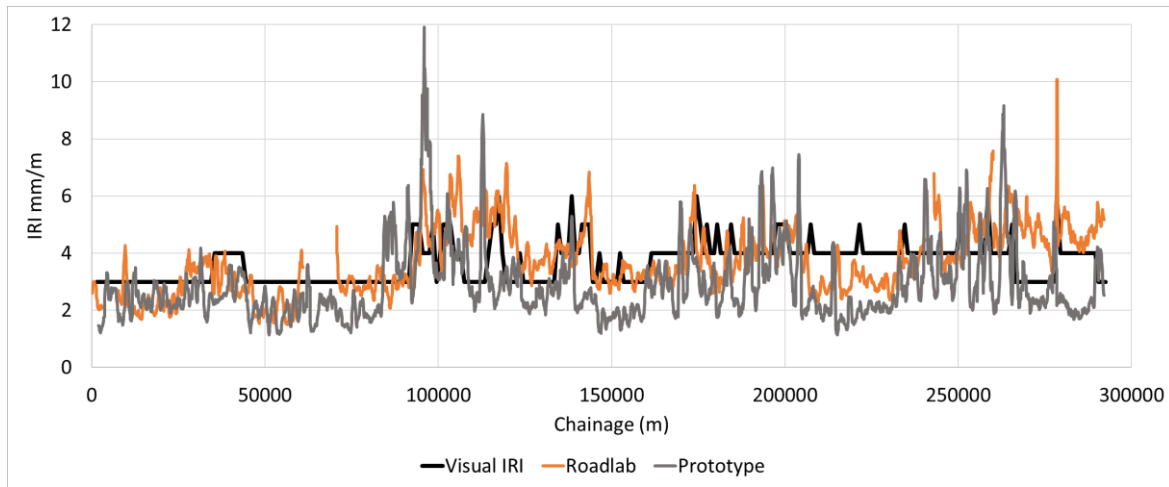


Figure 39: 1 km moving average IRI on 300 km site Westbound, showing visual IRI, Roadlab app and TRL prototype

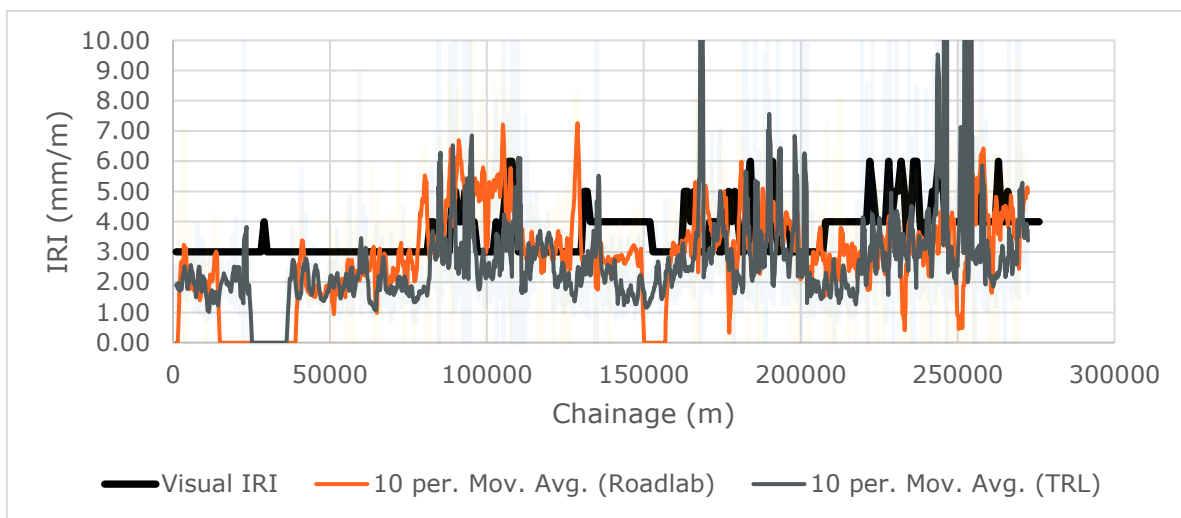


Figure 40: 1 km moving average IRI on 300 km site Eastbound, showing visual IRI, Roadlab app and TRL prototype

It can be seen in Figure 39 that the visual IRI inspection only reported 3 levels of IRI, (3, 4 and 5). To try and understand the performance of the RoadLab and TRL prototype in the context of this data, we have also thresholded the inertial IRI data reported by the services into three bands. The upper and lower inertial IRI levels (thresholds) for each band were obtained by determining the level of IRI that would give similar amounts of data within each band for each device and the visual IRI data. This results in the thresholds shown in Table 1 and the data is split as shown in Figure 41. We can see from the thresholds that the Roadlab

generally reports higher values of IRI than the TRL prototype, which agrees with the distributions seen in Figure 36.

Table 1: Thresholds for IRI data

Band/Device	Visual IRI	Roadlab	TRL prototype
Low	3	<3.5	< 2.5
Medium	4	>=3.5 and <5.5	>= 2.5 and < 4.5
High	5	>= 5.5	>= 4.5

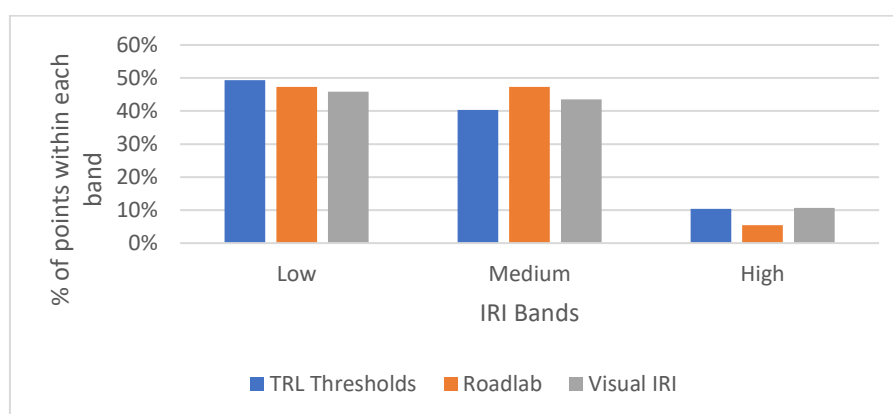


Figure 41: IRI split into bands with approximately the same amount of data in each band for each device

After applying this broad three level “normalisation” of the device IRI data we can then attempt to understand how well the Roadlab and TRL prototype match the visual IRI in terms of the locations at which each level of IRI is reported. Table 2 shows the confusion matrices for the two devices vs the visual assessment. For a perfect agreement all values in the leading diagonal would be 100%, which would say (e.g.) that all locations reported as “low” by the device were also reported as “low” in the visual assessment.

Although Table 2 shows that the performance falls short of the ideal 100% in the leading diagonal, it can be seen that both the TRL Prototype and the Roadlab devices give similar levels of performance. Also, both are generally within the correct range (i.e. when the visual reports high roughness the devices both report a large majority of the data to be high or medium roughness). Unfortunately, there was no further or more robust reference available in the El Salvador trials.

Table 2: Performance tables of thresholded data (Visual vs TRL Prototype, top, and Visual vs Roadlab, bottom)

		Visual IRI		
		low	medium	high
TRL Prototype	low	63%	40%	29%
	medium	31%	50%	42%
	high	6%	10%	29%

		Visual IRI		
		low	Medium	high
Roadlab	low	57%	39%	16%
	medium	31%	54%	68%
	high	2%	6%	16%

8 Discussion and Conclusions

This research aimed to determine whether a low-cost prototype system could be developed to provide a viable alternative to high-end equipment deployed for the assessment of condition of highway pavements. The work has considered the low-cost sensors that are available and how these can be compared with higher resolution devices, with the following observations and conclusions.

For system control and data collection/delivery the work has found that the Raspberry Pi provides a practical system that can be used as the Data Acquisition Unit (DAU). It is capable of integrating different sensors whilst also meeting the low-cost demands and has the ability to connect to both Wi-fi and Bluetooth for wireless communication. 4G communications can also be added. This work has found that it is not practical to use 4G to achieve real-time delivery of the data collected, due to the large volumes. However, the system does provide simple and rapid delivery of the data via the on-board storage, or wirelessly “in the office”. The data delivery is therefore suited to application in operational asset management, such as for day-to-day surveys or to rapidly identify emerging visual defects.

Accurate location referencing is important for all network surveys. The low-cost GPS of the prototype unit delivered coordinates such that ~80% fall within $\pm 5\text{m}$ of the reference. This would be sufficient for locating visual defects and for reporting condition over longer reporting lengths (e.g. the 100m lengths typically used in asset management systems). However, the performance in the use of GPS to calculate survey chainages does vary, depending on the geometry of the site (e.g. at tight bends on roundabouts). This would complicate application for assessing local defects or to detect change in condition. Further development could include the integration of a WAAS enabled GPS which, although costing up to £200, would provide benefits for both coordinate and distance-based reporting. More sophisticated processing of the data could also be used to improve the distance performance, drawing on the inertial measurements to provide basic fill-in of gaps in the GPS data.

As a key component of the data provided by the higher cost systems is visual condition, the research has trialled a number of low-cost cameras. The PiCam offers a low-cost solution, delivering images in which defects can be identified. Although higher resolution cameras are available at a reasonable price, the PiCam is designed to offer high-speed data capture through a designated port on the Pi, and therefore offers the most practical implementation. The use of a higher specification camera may require changing the DAU (from the Pi) to a more powerful system, which would increase costs.

By integrating an IMU with the DAU, the prototype is able to measure vehicle motion. This provides core data for the measurement of ride quality, which is a further key measure provided by high-cost systems. However, the measurement of ride quality in high-cost systems typically combines inertial and height measurements (see below) to measure the true road profile, from which IRI is calculated. Therefore, the use of IMU data alone provides only an approximation of the IRI. The UK trials have shown that the performance is lower than high-end systems (as may be expected). However, the prototype does give a broad indication of the roughness of the road, with an accuracy of around 90% within ± 1.5 of the reference. Refinement of the processing (e.g. implementing functions to filter out data with high levels of acceleration) would improve the performance.

As noted above, high end survey systems provide higher levels of performance in the measurement of profile by combining the IMU and laser measurements. Therefore, we have explored the use of low-cost laser measurement devices as a proxy for the high-resolution lasers deployed on survey vehicles. The solid-state LiDAR showed potential in the broad measurement of shape, but suffers from both inconsistency, noise and adverse effects resulting from environmental conditions (sunlight). Current low-cost sensors appear unable to provide a direct proxy for current high-cost systems. Further LiDAR companies were approached to understand the emerging capabilities and it is apparent that systems costing several £k should be able to measure features such as potholes at traffic speed. There are constant developments being made in this area, with the costs falling so it should not be long until this vision is able to be achieved at low cost.

The international trials of the prototype were undertaken to explore the potential of the system in an example application to which it should be highly suited – the measurement of roughness on developing world roads. Roughness (IRI) is a common measure used to provide an indicator of both local and network level condition in developing world networks. The trials show that other systems have been developed that apply the concept of low cost and portability for this application. However, as in these trials, they are typically based on smartphones running an application that is set up by the user. Hence the user is tasked with inputting a number of settings regarding the vehicle suspension, the road type the vehicle will be driven on etc. to generate IRI. Furthermore, smartphone data can vary as a result of differences in the software, phone accelerometers etc. Hence their performance can be adversely affected by the way they are set-up. It is also difficult to get phones to collect images whilst surveying (in the trials, the smartphones did not collect images). The prototype delivered IRI to a similar range and performance as the smartphone approach, with the added ability to collect images, and the potential to add additional sensors and provide real time reporting via the cloud (where communications systems are available). The prototype could be set up in a controlled manner, deployed with no user input required and, as in the trials, easily operated. It was superior to the smartphones as a result of its comparable performance, combined with higher levels of practicality and capability, and the potential for high levels of consistency through a common low-cost measurement platform.

9 Recommendations and Future Use

The above work has identified a few areas where refinements could increase the practicality and performance of the current system:

- **Mounts:** The mounting and housing of the equipment could be improved to reduce movement during a survey, and to improve the ability to repeatedly mount the system in the same orientation within the vehicle (which would further improve the consistency of the IRI). A custom-made housing and mounting could provide the ability to more easily deploy the system both inside and outside the vehicle.
- **Measurement of Location.** Investing in a WAAS enabled GPS receiver would increase locational accuracy for coordinate-based location referencing and combined with processing improvements, would also increase the accuracy of distance-based referencing. This would provide more robust/practical section and chainage referencing.
- **Measurement of Roughness:** The IRI showed good performance in the UK trials, and also performed adequately in the El Salvador trials. However, development of the algorithms for calculating IRI could improve the filtering (e.g. of poor survey conditions), and there would be benefit in deploying an improved/formalised calibration procedure. These would enable more fully automated delivery of IRI with little need for user input.
- **Data delivery/Communications.** Further development of the data delivery (e.g. improving the upload data to the cloud, or use of wi-fi) could enhance the practicality and “autonomy” of system operation. It would also move the device closer to a real-time tool that supports operational network asset management.

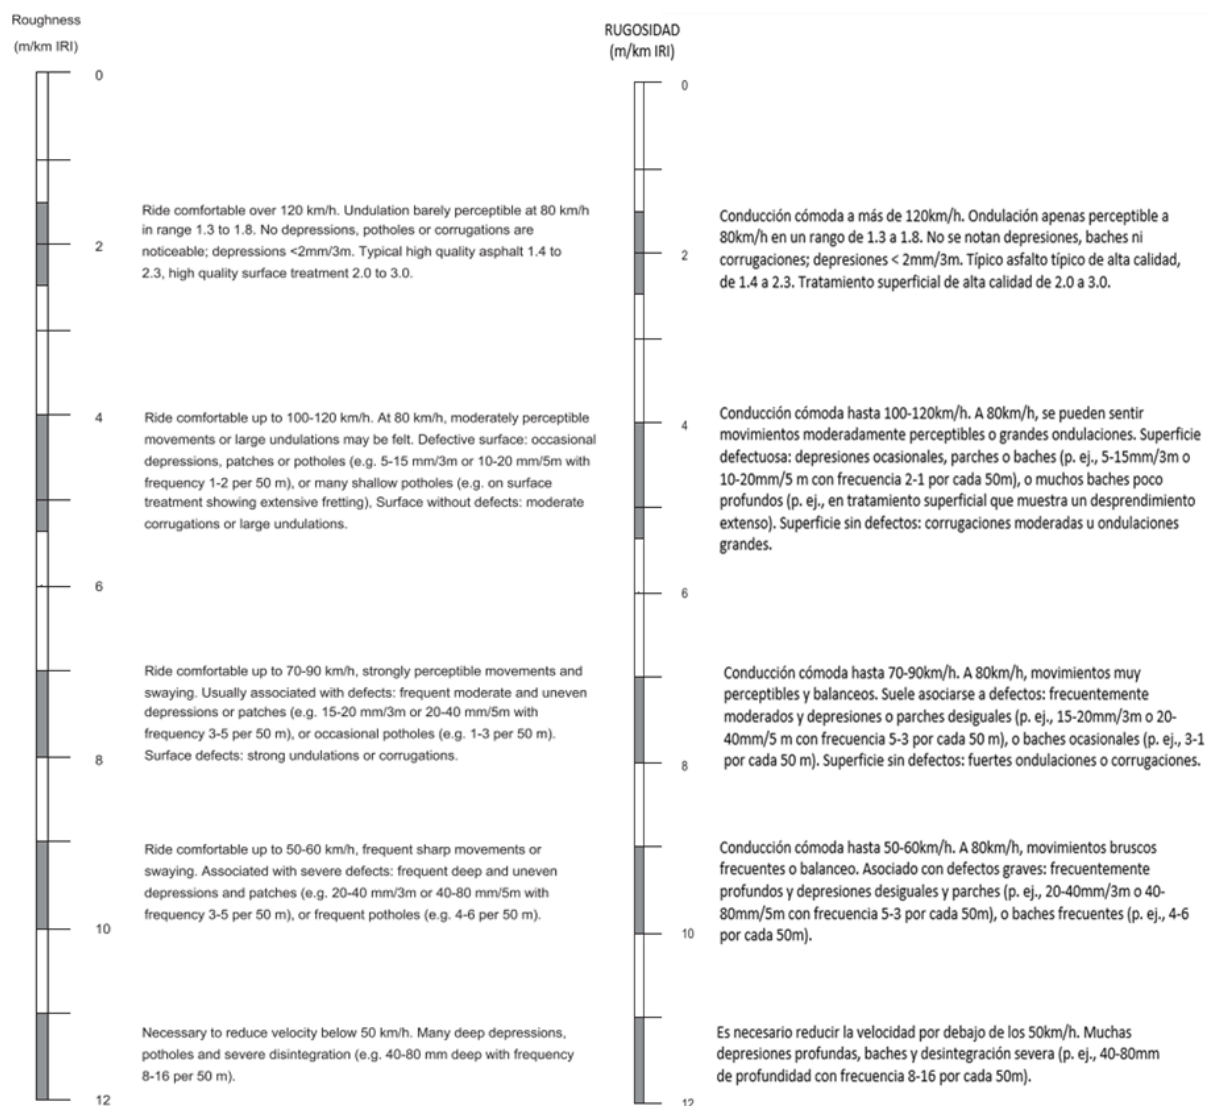
In the light of the research, development and trials of the device, it is felt that its initial application for condition surveys would lie in global south nations. Many of these nations are unable to deploy the complex equipment required to carry out network level condition assessments, and therefore rely on incomplete or inaccurate data when making decisions on maintenance. Implementing the above refinements would support the delivery of a robust network assessment tool that could be deployed to developing world nations as multiple low-cost surveying units providing semi-autonomous, real time, network wide measurements of roughness, supported by imagery for assessment of visual condition:

- As the quality of images increases there is potential to apply machine learning to automate the detection of operational defects within images.
- As communication systems develop, global south nations could access a fully working tool that allows all data to be visualised (e.g. as shown in Figure 26) and analysed in almost real-time.
- As the system is expandable, additional sensors could be added. For example, adding air quality sensors and collecting data on a large scale to support developing world nations tackling increasing pollution in urban areas, potentially linking the data to traffic levels, weather etc.

-
- The system could also support the iRAP initiative within developing nations. iRAP is a system to improve road safety of infrastructure by inspecting high-risk roads to develop Star Ratings, Risk Maps and Safer Roads Investment Plans (<https://irap.org/>). iRAP uses images of roads to understand the safety of pavements and their surroundings - inspecting the road alignment, inventory and the condition of pavement, structures and furniture. The prototype system offers great potential within iRAP as a low-cost transportable system with minimal setup and easy operation.

Appendix A World Bank IRI Visual Scale

World Bank IRI Visual Scale



Next Generation Monitoring Systems



Survey vehicles, operating at traffic-speed, are deployed across the road network to assess the condition of road pavements. These apply high-quality (and high cost) equipment to measure condition. However, significant progress has been made in the development of low-cost sensors and data collection units that may have potential for application in highways. This project has aimed to understand the capabilities of this emerging technology. The project explores the technologies and combines a Raspberry-Pi based Data Acquisition System, compact camera, GPS, inertial measurement system, Wifi and 4G GSM comms and a low-cost Solid State LiDAR into a prototype device. The total cost is a few hundred pounds.

Trials characterise the prototype system. Although the solid-state LiDAR sensors are not found to be robust in this application, the remaining sensors show strong potential for use in road condition assessment. A wider trial of the prototype system in a potential application – the measurement of roughness (IRI) on developing world road networks – was carried out in El Salvador. The prototype shows comparable performance with alternatives, combined with higher levels of practicality and capability, and the potential for higher levels of consistency through a common low-cost measurement platform. In the light of this research, it is felt that, following refinements to the prototype, the initial application for the device would be to condition surveys in developing world nations.

TRL

Crowthorne House, Nine Mile Ride,
Wokingham, Berkshire, RG40 3GA,
United Kingdom
T: +44 (0) 1344 773131
F: +44 (0) 1344 770356
E: enquiries@trl.co.uk
W: www.trl.co.uk

ISSN 2514-9695

ISBN 978-1-915227-22-5

ACA022