

6th Transport Research Arena April 18-21, 2016



Sensors on vehicles (SENSOVO) – proof-of-concept for road surface distress detection with wheel accelerations and ToF camera data collected by a fleet of ordinary vehicles

Carl Van Geem ^{a,*}, Marleen Bellen ^b, Boris Bogaerts ^d, Bart Beusen ^e,
Bruno Berlémont ^a, Tobias Denys ^e, Paul De Meulenaere ^f, Luc Mertens ^d, Peter Hellinckx ^c.

^a Belgian Road Research Centre, Bld. de la Woluwe 42, B-1200 Brussels, Belgium

^b Flanders Institute for Mobility, Wetenschapspark 13, B-3590 Diepenbeek, Belgium

^c MOSAIC - University of Antwerp, Rodestraat 4, B-2000 Antwerp, Belgium

^d Op3Mech - University of Antwerp, Salasianenlaan 90, B-2660 Hoboken, Belgium

^e Flemish Institute for Technological Research, Boeretang 200, B-2400 Mol, Belgium

^f CoSys lab - University of Antwerp, Rodestraat 4, B-2000 Antwerp, Belgium

Abstract

This contribution presents the results of the “SENSOVO” project initiated by the Flanders Institute for Mobility (VIM), executed by the University of Antwerp (UAntwerp), the Flemish Institute for Technological Research (VITO) and the Belgian Road Research Centre (BRRC), and supported by several other parties.

Both road users and road managers could benefit from massively, continuously, automatized collecting of information on road surface distress (potholes, cracking, subsidence,...) by a fleet of vehicles equipped with low-cost sensors. Road users could have immediate information on road conditions while road managers could get year-round insight on the general performance of the road network in addition to the data they obtain from annual inspections with specialized monitoring devices.

The project’s objective was to investigate possibilities of road surface distress detection using data collected by a fleet of vehicles. Two scenarios were considered: a large fleet of ordinary cars and trucks transmitting collected relevant sensor data already available on the CAN-bus in such vehicles; a vehicle, equipped with one or more Time-of-Flight cameras (ToF).

Data on the CAN-bus were collected using either a CAN-logger that sends its data to a central server using GPRS where it is processed for road surface distress detection, or using an OBD scan-tool that sends its data using Bluetooth to a smartphone that already processes the data into detection events, forwarding only these to a central server. The performances were tested by

* Corresponding author. Tel.: +32-10-236522; fax: +32-10-236505.

E-mail address: c.vangeem@brcc.be

UAntwerp and VITO. A simple computation developed by UAntwerp on the speed of all four wheels and on the vertical accelerations often allows indicating road distress. Either the CAN data-logger or the smartphone delivers the GPS location of the event. Several cars equipped with this technology sent their observations to a central database.

Several ToF cameras were benchmarked by UAntwerp. Road data were collected at 40km/h and 40frames/s with ToF-cameras of brands Mesa and Fotonic. Several image processing algorithms dedicated to the identification of road distress on a flow of images were developed by UAntwerp. It has been demonstrated that several types of “unevenness” of the road surface (including potholes) can be detected. Some ToF-camera observations were added to the central database.

The BRRC developed and implemented an algorithm for treatment and interpretation of the collected events in the central database using the ArcGIS software, simulating both sending out alerts to road users and providing daily “quality scores” for each road section in the network to the road manager. For this, the network was defined from a geographical map provided by the Flanders Geographical Information Agency (FGIA/AGIV). As soon as “enough” observations are made taking into account the frequency of observations in time, the defect is considered as real and will be reported. When the defect is no longer observed, it is considered as being repaired. A “score” for each road section was computed daily, from the amount of defects present that day.

The SENSOVO project along other recent research projects worldwide delivers a proof-of-concept for fleet probing of road surface distress.

© 2016 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Peer-review under responsibility of Road and Bridge Research Institute (IBDiM)

Keywords: Probe fleet monitoring; CAN-bus data; ToF camera; V2I/I2V communication; pavement management

1. Introduction

Road surface distress can be an immediate hazard for the road user or an indication of road deterioration. Road users would benefit of the availability of adequate real-time information on the presence of potholes and other distresses that have an impact on the safety or the comfort for the user. Pavement management by the road authority can also be more efficient when the network is regularly monitored. Monitoring of the road network by fleet probing has been suggested in many previous works (e.g. Deuss et al. (2008), Saarikivi et al. (2014)). It would allow almost continuous data collection on the whole road network, which means that the information can potentially be used for deployment of rapid urgent interventions and temporary repairs with a direct benefit for the road users. Such data collection techniques include extraction of CAN data as for example in Deuss et al. (2008), Varanka et al. (2008), and Nitsche et al. (2014), and the use of smartphone data as for example in Mohan et al. (2008), Yagi (2010), Byrne et al. (2013), Yagi (2014), and Van Geem (2014).

Usually road managers of major road networks do monitor by (bi-) annual inspection rounds with dedicated probe vehicles. Local road networks usually are less well monitored and may only be inspected visually once in a while. Data collected by a probe fleet could be used as additional information for road managers. On major road networks the dedicated probe vehicles could be sent for an additional check-up to road sections that seem to have deteriorated more rapidly than expected according to the fleet data. On local road networks the fleet data could allow a more frequent follow-up of the state of the network and provide additional data for road management.

This paper presents a proof-of-concept for probe vehicle monitoring and for the exploitation of the collected data. Two technologies are considered: CAN data collection and interpretation using a fleet of ordinary cars or trucks, and ToF cameras to be installed on a small fleet of vehicles that already visit the whole road network. Collected data were put on a map and transformed into quality scores for each of the road sections in the network. Both rapid communication to road users and data analysis over a longer period of time for road management were simulated.

Nomenclature

CAN	Controller Area Network: serial bus inside a vehicle on which data are transported between different microcontrollers and other electronic devices
OBD	On Board Diagnostics
ToF camera	Time-of-Flight camera: a range imaging camera system
Quality score	An indicator of the severity of the probable road surface distress, as reported by fleet data
GIS	Geographic Information System

2. CAN-bus data

All kinds of data are continuously collected by sensors on board of ordinary cars and trucks. These data are communicated over a CAN-bus. Depending on the kind of vehicle the data can be read through the OBD plug or by an immediate connection with the CAN-bus itself using a CAN-bus clamp.

Several tools were tested for the continuous monitoring of data on the CAN-bus through the OBD plug. Several communication techniques and set-ups were tested for retrieving the data. A selection of relevant available CAN data was made and an algorithm was devised in order to report a possible presence of a road defect. A condition for the detection of such a defect is that a wheel of the vehicle traverses the defect.

The set-up was implemented in six cars and tested during more than six months on the Flemish road network. Collected data on the potential presence of defects was stored in an event data base in order to make it available for further interpretation. The algorithm was also successfully tested for a truck using simulation techniques.

2.1. Data retrieval and communication techniques

During the different stages of this project three different data retrieval methodologies were used. This was done in order to anticipate on the needs in the different stages.

In a first stage it was important to directly analyse the CAN data while driving in order to retrieve the required vehicle specific information (CAN identifiers, algorithm thresholds,...) to fine-tune the algorithm and in order to get immediate feedback of the algorithm while driving. To tackle these needs the Kvaser Leaf SemiPro HS was used in combination with an in house made data analysis program written in .NET. This in-vehicle analysis tool ran on a laptop connected to the CAN-bus using the Kvaser Leaf SemiPro USB device. This allowed for real-time monitoring in the car and retrieving relevant CAN data in combination with visual observations (Figure 1).

In a second stage, data on wheel speeds was gathered via an Android smartphone that was connected to the CAN-bus via an OBD-scan tool on the OBD connector. OBD-scan tools are mostly used to read specific OBD parameters, but the tools also support the option to passively listen to all information traffic that is available on the CAN-bus. By filtering out all unwanted messages, only the CAN-bus messages that the algorithm needs are processed by the scan tool and transmitted to the smartphone via a Bluetooth link. As shown in Figure 2 the Kvaser module was replaced by a ELM 327 and the laptop by a common smartphone. The ELM module was connected to the smartphone using Bluetooth. The smartphone itself did the pothole detection and submitted the defects to a central server using a 3G connection.



Fig. 1. in vehicle data analysis tool.



Fig. 2. Different types of data connections.

The wheel speed information from the CAN-bus is then combined with acceleration info from the built-in accelerometer, providing us with all necessary parameters to implement the pothole detection algorithm on the smartphone. When a pothole is detected, an event is created, including the position information that is read from the built-in GPS receiver. Events are immediately uploaded to the project data server using a web service interface.

In a last stage the whole concept had to be deployed on a large fleet of vehicles. To make this integration possible and reliable on a large fleet we had to move towards off the shelf solutions. For this we opted for a technology of Altran (prior TASS). As shown in Figure 2 the uCan module connects to a server of Altran which will on his turn connect to the central server described in the second solution.

2.2. Algorithm for pothole detection

As part of the huge amount of CAN-bus data, the real-time sensor data represent the actual physical state of some of the passenger car’s devices. The first step in the development of the algorithm is therefore to decide which sensor data may provide valuable input for the defect detection.

Aiming for the observation of vibrations generated by a defect, the sensors listed in Table 1 are deemed good candidates. The table also shows in which vehicle systems these sensors can be found. Obviously, not all cars dispose of all of the mentioned vehicle systems, so neither of the associated sensors. As our goal is to develop a widely usable pothole detection algorithm, it is necessary to rely only on those sensors that are present in a broad range of passenger cars and trucks if possible.

Table 1: Overview of relevant vehicle systems.

		Vehicle systems				
		ABS	TCS	ESP	Active steering	(semi-) active suspension
Sensors	Wheel speed	x	x	x	x	x
	Steering wheel angle			x	x	x
	Longitudinal and lateral acceleration			x		x
	Vertical acceleration					x
	Yaw rate			x		x
	Altitude					x

From a mechanical perspective, the vertical acceleration sensor would yield direct evidence of a car hitting a road defect. This sensor is however not broadly present in cars. As an alternative, Nilsson et al. (2012) uses the smartphone’s acceleration sensor.

The wheel speed sensors are broadly available as almost every car is equipped with an ABS (Anti-lock Braking System). Pothole detection based on these sensors is however less intuitive. Prior to building a pothole detection based on the wheel speed sensors, we first ran simulations to gain more insights in vehicle behaviour while hitting a road defect. Afterwards we used the same simulation methodology to test the proposed detection algorithm.

A detailed mechanical model of the car has been realized using the LMS Amesim simulation package. The model allows for a virtual car to drive on a predefined road surface at various speeds. Next to a complete suspension system, the model also contains a gear, to investigate whether the effect of a gear shift can be discerned from hitting a pothole. The model for the road surface contains a pothole of 4 cm depth and 40 cm diameter, and can be hit by the front right wheel only.




Fig. 3. Simulated wheel speeds (ordinates) at 50 km/h with a pothole at time step 32 s for a straight track (top) and in a curve (bottom).




Fig. 4. Deviation between front and back wheel speed differences (ordinate) in function of time laps (abscissa), with threshold on 1.3 km/h.

Figure 3 shows a simulation result for all four wheel speeds both on a straight road part and in a curve. The wheel speed of the front right wheel clearly deviates from the others, which is a clear indication of the possible presence of a pothole. For obvious reasons, wheel speeds also differ in road curves. The proposed pothole detection algorithm therefore takes all four wheel speeds into account and decides on the presence of a pothole if:

$$|(v_{FL} - v_{FR}) - (v_{RL} - v_{RR})| > T \quad (1)$$

where v stands for the individual wheel speed and the index FL for Front Left, FR for Front Right, RL for Rear Left and RR for Rear Right. The threshold T needs to be chosen such that false positives will be excluded as far as possible. Gear shifts, passing pedestrian refuges or the impact of different types of pavement (e.g. cobble stones) can be excluded this way.

2.3. Tests on the Flemish road network

The algorithm has been implemented on a computer connected to the CAN-bus as described in section 2.1. The computer collects the wheel-speed data from the CAN-bus and shows the results in real-time. The events of the individual pothole detections are sent instantaneously over the mobile communication network to a central server.

Figure 4 shows the wheel speed difference, as defined in expression (1), for a real measurement. These observations indicate that, for the actual car used, it is meaningful to put the threshold on 1.3 km/h. The peaks in the centre of the figure, surpassing the threshold, correspond to one pothole. It is hence feasible to sense all potholes on a certain trajectory, provided that they are hit by one of the wheels and that they are sufficiently large to produce peaks that surpass the threshold. To gain more data, a fleet of 6 cars has been equipped with the monitoring devices.

2.4. Truck simulation

In a final test stage it was important to know whether the proposed technology was also applicable to busses and trucks. In order to prove this we fell back on the simulation technology described in 2.2. Instead of working with a normal passenger car we replaced the simulation models with those of a Volvo FH12 420. We did the tests with a large set of different potholes (40 cm x 4 cm, 40x2, 30x4, 30x2, 20x4, 20x2, 10x4, 10x2, 5x4, 5x2) and at different speeds (10, 20, 30, 40, 50, 90 km/h). We notice that every pothole except for the 5x4 and 5x2 at 90 km/h is detected. In analogy to the passenger car we believe that we can extrapolate these results from simulation level towards fleet level. And for this we conclude that the technology can also be used on trucks and even busses.

3. ToF camera data

The challenge for the use of a ToF camera on a probe vehicle mainly resided in the speed of the movement of the camera on a moving vehicle and the object in the image: the road surface. Several cameras of different brands were benchmarked with data collected on the road at different speeds. Several image processing algorithms were developed and applied to the collected images and criteria were established for the detection of several types of unevenness of the road surface including potholes. The ToF camera technology is very recent and still in full evolution. During the project it could be observed that performances improved just by technological progress and it can be expected that the

more than promising results of the SENSOVO project will already be out-performed by newer versions of ToF cameras.

In order to show that the information collected by a ToF camera could be exploited in an analogous way as the CAN-bus data, a small amount of data from a ToF camera inspection of a road was added to the data base of “events” for further processing.

3.1. Benchmarked ToF cameras

Recently, novel solid state 3D technologies have emerged, leading to 3D vision systems with radically improved characteristics. In the SENSOVO project the possibilities of ToF cameras are evaluated in real outdoor experiments. The problems with ToF image blur in moving scenes were studied profoundly. Basic experiments in which a tennis ball is presented at high speed to a ToF camera will show the ball changing into a snake alike subject. The ball smears out and seems to oscillate around the floor it is rolling on. With some phantasy it could be associated with the famous ‘Loch Ness monster’. For bumps or potholes in the road the same anomaly will happen. Therefore we have to distinguish two different types of ToF cameras. One group acts like the MESA Swiss Ranger cameras (SR400 & SR4500) that calculates the ‘phase shift’ between a send NIR sine wave (30 MHz) relative to the received reflection. In static applications this principle works well. For moving objects or moving cameras the situation is quite different. The received signal is a mix over time from different reflection conditions of a continuous changing world point leading to more or less chaotic measurements. Discontinuities like e.g. road mark signs are given back as unrealistic spikes. A second group of ToF cameras behave better in dynamic situations. Cameras from Fotonic or Melexis, for instance, use ‘block wave light integration’ in a shorter amount of time. With such cameras we got acceptable results up to camera speeds around 40 km/h, making about 40 images a second. Progress will come since higher frame rates are announced and better resistance against sun light might be expected in the near future. Details around recent time of flight cameras can be found in documents like Hussmann et al. (2013) and Mertens et al. (2013).

3.2. Data collection and image processing

A general overview of the SENSOVO ToF-experiments will be given by means of a typical example. The image analysis was carried out in different steps and the results are presented in an image format. It concerns a stroke of about 43 m at the following GPS coordinates: Latitude = 51.1033 ; Altitude = 3.7461. First of all the raw ToF image data (depth values) are median filtered (Figure 5) in order to suppress local out of controls (OoC). A lot of OoC’s remain in the corners of the image. This is due to the combination of the incident angle and the reflection of the NIR light on the road. A minimum reflection must be found before a pixel measurement can be accepted. Most of these OoC’s will disappear after the rectification step (Figure 6).

NCC is a quick and accurate method to derive the camera speed between every two consecutive images. 167 random pixels (white spots) from the second half ‘previous image’ are correlated to row shifted points in the ‘next image’. The calculation is repeated 15 times and the best correlation delivers the best fit shift. Here a row shift of 24 lines occurs (Figure 7).




Fig. 5. Depth image, curved and with ‘out of controls’ at the border of the image.




Fig. 6. Rectified relative values: 'out of control' pixels are overwritten by a mean distance value (1661 mm).




Fig. 7. Camera speed evaluation based on normalized cross-correlation (NCC) between images.

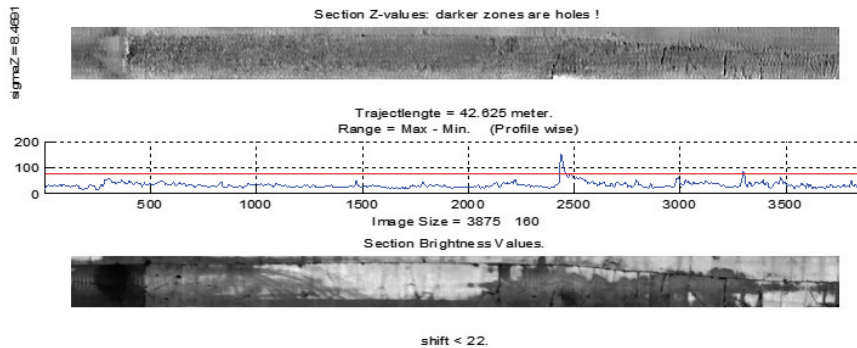


Fig. 8. sectional image.

If speed is known precisely, coinciding parts of consecutive images can be averaged giving nice sectional images. Depending on the speed more columns are found (here 3875 columns, Figure 8).

3.3. Successful road distress detection

Time of flight cameras are able to evaluate 'Road Damage'. If the cameras spatial resolution increases the camera can be placed higher (e.g. at a height of 2.5 m.). At that moment each pixel can cover e.g. 16x16 mm² and higher camera speeds are allowed. Higher frame rates will make the job easier and more accurate. Such cameras recently became available (Melexis, Panasonic,...). Important is that the overlap of the images is high enough in order to reduce noise. As an overall quality one might expect results as given in Figure 9.

From the connected holes/bumps locations the overall area is calculated knowing that 1 pixel covers 11x11 mm². If an area of 1 dm² is taken as a threshold, damage locations can be found. If depth is combined the local damage volume can be derived (e.g. in dm³). This gives an impression of the seriousness of the local damage. The world positions of the holes/bumps can be found relative to the measured GPS-coordinates at the image borders.

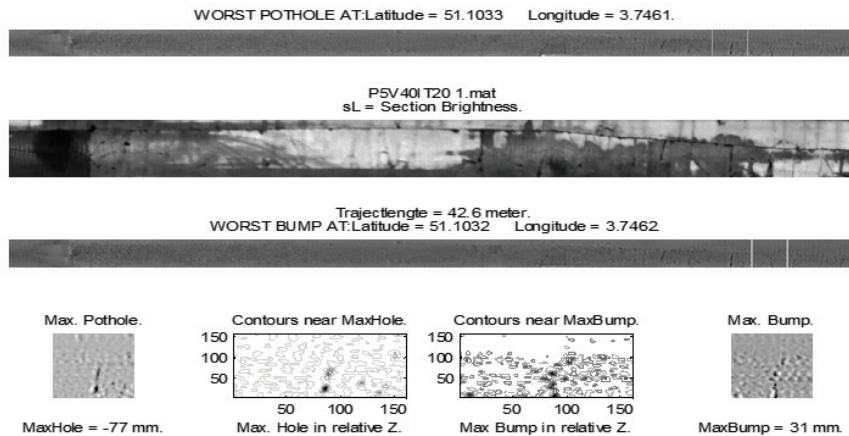


Fig. 9. Illustration of final result.

4. Maps and scores

Both fleet monitoring techniques delivered entries to a data base of “events”. Each event represented an observation of a potentially present defect of the road surface. Each event also came with a GPS-position with rather feeble precision. In order to further exploit these events, we first grouped those events that occurred on GPS-locations that were not “too far away” from each other and assigned them to the same “potential defect”. The location of each “potential defect” was also assigned to a road section represented by a polygon in the map of the road network.

An algorithm was developed that simulated the incremental arrival of “events” and the sending out of an alert to road users as soon as a “sufficient number” of events were assigned to the same potential defect. Frequency of observations was taken account for. When a potential defect was observed no longer, it was considered repaired. Also, on a daily basis, the algorithm computed a “daily score” for each of the road sections in the network. The daily scores were stored in a data base and this data base could be examined further from a road management perspective: e.g. by studying the evolution in time of the daily score of a particular road section, by computing an aggregate score for the whole network, etc.

4.1. Grouping events and road sections on a map

The Flanders Geographical Information Agency (FGIA/AGIV) provided the geographical maps of four cities in Flanders. The maps were used in the ArcGIS software and presented pieces of roads as polygons: each road section between road crossings as well as each of the road crossings is a polygon. Several techniques were tried out for determining the polygon within which a GPS location of a reported event lies. Due to inaccuracies of GPS collection, some events were reported outside any polygon in the map. It was observed that most of these points are less than 3 m away from the nearest polygon and they were attributed to that polygon. Also some events were reported on private terrains or parking areas not represented in the map and were ignored.

Due to inaccuracies of GPS collection, two events that represent the same real road surface defect will have different GPS positions. Therefore the events that were reported on GPS locations considered as “nearby” were grouped together as representatives of the same potential road surface defect. However, since the vehicles circulate on the same lane in the same direction, GPS positions of reported events typically form long stretches rather than circular areas around a central position. Therefore we used the Wald test (first presented in Wald (1943)) for the iterative formation of clusters. Nearby clusters can be merged to one. The position of the cluster is defined as the gravitational centre of the GPS locations of the events belonging to the cluster and considered as the location of the potential road surface defect.

The algorithms working on the data and the maps were developed in the script language Python. It was observed that the execution times of the computations with the different in ArcGIS available API-functions could vary a lot. In particular the choice between distance related operations for the computation of inclusion of a point in a polygon, distance from a point to a polygon, inclusion of a point in a polygon with offset, etc. had a very important impact on execution times. Note however that this is strongly related to the architecture of the simulation program we realized in the project and that this must be tackled differently in a future real application.

4.2. Reporting potential defects to the road users

More and more events will join an existing cluster. When the number of recent events forming a cluster reaches a threshold, it is considered that a local road surface defect really exists and an alert is created. This alert could then be sent to all road users.

Since the monitoring techniques are not controlled as is the case when inspections are done with dedicated road monitoring devices, the risk of false positives and unreliable events in the data base is real. This already is the motivation for sending out an alert only when several events are reported on almost the same location. False positives will not be confirmed by other events. So when an isolated event is not confirmed within a reasonable time frame, the event and its cluster can be deleted.

When a local road surface defect gets repaired, there will no longer be any reporting of events at that position. Therefore, when during a reasonable time no more events can be added to a cluster for which an alert was sent out, we consider the defect as being repaired and we delete the cluster.

However, since the traffic numbers on different roads can vary a lot, the frequency with which a road surface defect is reported will also vary a lot. Therefore we must define the time windows in function of the traffic numbers. For most of the local road sections there are no traffic data available so we use the frequency of the reporting of events belonging to a same cluster as the indicator of traffic density. The time windows used are defined by the frequency of previous observations multiplied by a constant factor. This constant factor is set to a higher value as soon as the cluster generated an alert: indeed, for instance, we can expect that a driver alerted on the existence of a pothole will try to avoid driving through it and although the pothole is still there it will less frequently be reported.

4.3. Daily scores for the road manager

At the end of each day, we evaluate the number of clusters within each road section. The definition of the daily score could also be made dependent on the type of reported defects by associating a different weight to each type. The definition of the daily score can follow the format suggested in the final report of COST action 354, Litzka et al. (2008). Over time, more clusters may appear in a same road section, which corresponds to a degradation of the state of the road surface. However, the local repairs also make clusters disappear again. The analysis over time of the evolution of the daily scores of all road sections in the network gives an indication on the necessity and regularity of urgent, local maintenance interventions and where these have taken place. The road manager can use this information for the planning on the somewhat longer term of more definitive maintenance works.

4.4. Simulation results on some real data

More than 35000 events were registered in the data base, of which more than 16000 were located in the areas covered by the available maps. The events were concentrated in more than 1800 different polygons. Making use of somewhat too relaxed threshold values, we generated 58 alerts. In order to test the potential for statistical exploitation of daily scores, we multiplied these events by copying them and assigning them to following days and weeks. We produced some graphs of the evolution of weekly averages of the daily scores of a few road sections.

These experiments show that the potential is there for both alerting road users in near real-time and supporting road authorities in decision making on road maintenance. However, the large amount of data collected in the frame of the SENSOVO project did not yet suffice in order to fine-tune the different parameters for the interpretation of the data.

5. Conclusions

The SENSOVO project delivered a proof-of-concept of a fleet probing strategy for the collection of data on road surface distress. Collecting information with a selection of CAN-bus data available in modern cars and trucks allows indicating the presence of surface distress quite accurately. Vehicle manufacturers or distributors of GPS-route planners for instance could be interested in offering such information. The ToF camera (hardware cost of approximately 6000€ at the beginning of the SENSOVO project) is a promising tool that can be mounted on a small fleet of vehicles (owned by the road administrator for instance) and can detect several types of surface distress. The huge amount of data collected by a fleet using these rather inexpensive tools can be mapped and transformed in “scores” for road sections giving an indication on road surface quality. The project showed that this information can be used for rapid communication to road users who can then adapt their driving style. When treated statistically over a longer period, the same information can indicate the evolution of road surface distress and can give an overview of road surface quality on the whole road network. Hence, the data can also be an input for pavement management on network level by a road manager. However, several difficulties will still have to be overcome, such as the lack of standards for CAN-bus data formats, fine tuning of the ToF camera image interpretation techniques, data mining on large sets of fleet probing data and the implementation of an appropriate business model.

Acknowledgements

The SENSOVO project was initiated and coordinated by VIM and financially supported by the Agency for Innovation by Science and Technology (IWT) of the Flemish government, as well as by project partners Agiv, Aswebo, Athlon, AWW, Beijer Automotive, Caeleste, Coyote, Datavision, Port of Ghent, Melexis, Mobistar, Siemens Industry Software NV, Altran, VAB, and Voxdale.

References

- Byrne, M., Parry, T., Isola, R., Dawson, A., 2013, Identifying road defect information from smartphones, in *Road & Transport Research*, Vol. 22 No. 1.
- Deuss, M., et al., 2008, *Intelligent Roads*, Final Summary Report of the European FP6 “INTRO” project.
- Hussmann S., et al., 2013, A review on commercial solid state 3D cameras for machine vision applications, Chapter 11 in Buytaert, J. (editor), *Recent Advances in Topography Research*, ISBN: 979-1-62618-840-2, 2013 Nova Science Publishers, Inc.
- Litzka, J., et al., 2008, The way forward for pavement performance indicators across Europe, Final report of COST action 354, *Performance Indicators for Road Pavements*, ISBN 978-3-200-01238-7.
- Mertens L. et al., 2013, Time of Flight Cameras (3D Vision), Chapter 12 in Buytaert, J. (editor), *Recent Advances in Topography Research*, ISBN: 979-1-62618-840-2, 2013 Nova Science Publishers, Inc.
- Mohan P., Padmanabhan V.N., Ramjee R., 2008, Nericell: Rich Monitoring of Road and Traffic Conditions using Mobile Smartphones, in *Proceedings of SenSys'08*, Raleigh, NC, USA.
- Nilsson S., Sjögren L., Lundberg T., 2012, Ojämnhetsmätning med smartphone, Ett inledande testfall, VTI report 2012/06/21 (in Swedish).
- Nitsche, Ph., Van Geem, C., Stütz, R., Mocanu, I., Sjögren, L., 2014, Monitoring Ride Quality on Roads with Existing Sensors in Passenger Cars, at the 26th ARRB Conference - Research driving efficiency, Sydney, New South Wales.
- Saarikivi, P., Andersson, M.C., Ekström, P., Zachrisson, E., Gustavsson, T., Müller, S., 2014, MOBI-ROMA tool for mobile road monitoring, at the *Transport Research Arena 2014*, Paris.
- Van Geem, C., 2014, Fleet Monitoring of Road Comfort, at the ERPUG Forum and TRIMM Final Conference, Brussels.
- Varanka, M., Erkkilä, K., Noponen, K., Mäkelä, K., Seppänen, T., 2008, Automatic road slipperiness detection of heavy duty vehicles, in *Proceedings of the 11th International Symposium on Wireless Personal Multimedia Communications (WPMC 2008)*, Lapland, Finland.
- Wald, A., 1943, Tests of Statistical Hypotheses Concerning Several Parameters When the Number of Observations is Large, in *Transactions of the American Mathematical Society*, 54, pp.426-482.
- Yagi, K., 2010, Extensional Smartphone Probe For Road Bump Detection, in *Proceedings of the 17th ITS World Congress*, Busan.
- Yagi, K., 2014, Road Cracking Detection by Using Smartphone Accelerometer, at the *Pavement Evaluation 2014*, Blacksburg, VA, USA.