dMap: A Low-Cost Distributed Mapping Concept for Future Road Asset Management

Dana Pordel, Lars Petersson, Philip Kilby

National ICT Australia^{*}, Australian National University
Locked Bag 8001, Canberra, Australia
{dana.pordel, lars.petersson, philip.kilby@nicta.com.au}

ABSTRACT

Road assets including road signs and utility poles are among the objects that are closely monitored by government bodies and transportation authorities in many countries. This paper presents a new approach to collect and analyse the information needed for road asset management and other digital cartography applications. Instead of using dedicated surveying vehicles, skilled drivers and high quality sensors that are common for traditional data collection, it envisions an infra-structure in which the data collection is done through cheap sensors that mainstream consumer cars can host in the near future. In order to investigate the new design, a prototype sensor named dSensor has been developed, each to be mounted on public transport vehicles, to collect information from the road scene. Each dSensor has three cameras, wireless modem, Global Positioning Systems (GPS), accelerometer, gyroscope, magnetometer and a laptop. A collection of vehicles equipped with dSensors and a server constitute the distributed mapping system called dMap for road asset management. The dMap design supports many features of an ideal road asset find and replace management system and in comparison to similar solutions or concepts for road asset management, the dMap is cost effective with high quality information particularly on matters such as update rate.

Keywords

Distributed Cameras, Road Sign Detection, Road Asset Management, Object Detection

1. INTRODUCTION

Government bodies need to monitor roads assets, e.g., road signs, to maintain safety standards and quality of roads,

^{*}National ICT Australia is funded by the Australian Department of Communications, Information Technology and the Arts and the Australian Research Council through Backing Australia's Ability and the ICT Centre of Excellence Programs.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

ICDSC '14 November 04 - 07 2014, Venezia Mestre, Italy

Copyright 2014 ACM. 978-1-4503-2925-5/14/11 \$15.00.

<http://dx.doi.org/10.1145/2659021.2659053>

Figure 1: Result of running stop sign detector over an RGB image. Currently, it is possible to detect 50 road signs within RGB images by dSensor.

minimise accident fatalities, and improve the driving experience. According to US Department of Transportation, in 2008, about 2.7 % of the all driving accidents is directly caused by issues related to road signs [1]. Currently, road asset monitoring is often done manually by experts who patrol the road network to observe their last status and possible damages made over time, and identify actions needed to maintain their quality. However, automating the detection and localisation of the objects of interest along the roads as part of a general road asset monitoring system is becoming more and more feasible thanks to the recent achievements in pattern recognition and low cost sensor technologies.

There are new researches that partially automate road asset management by identifying and localising signs in video frames captured using high quality sensors mounted on dedicated vehicles. However, these systems often have costly sensors, e.g., cars equipped with ladars [21], costly to use which limits the achievable update rate, and require dedicated manpower i.e. skilled drivers, to help collecting the information usually using cars that are very costly and hardly applicable at a large scale. Also it is often the case that the collected survey data is stored on some physical storage, i.e. hard drives, which need to be transferred to powerful servers, via a time-consuming process, for analysis such as

object detection and localisation. One example of such data collection for a different task than road asset management, is the data collection for Google Street View. The main downfall of such systems is their data refresh rate which is very low due to the high costs of its data collection process. For Google, it might not be highly desired to refresh its street view frequently, however, for road asset management systems and GPS mapping systems, it is important to update the data frequently, as it may cause serious casualties including fatalities [1]. One way to assist the current road management process is to utilise images from satellites. It is possible to detect roads and monitor traffic using very high resolution satellite imagery [14],[12],[7]. However, for road signs the satellite images are not very helpful mostly because of the viewing angle and occlusion.

In order to produce high resolution quality information concerning the roads and in such a way that changes can be discovered soon after they happen, we propose to use a mobile network of many sensors that provide information from most of the roads everyday. Implementing such a system using a large number of conventional surveying vehicles would become prohibitively expensive. We believe that a practical solution is to instead use many, cheap sensors in a way that does not require much manpower for the data collection process or the following analysis. The use of many low-cost sensors rather than a few high-cost high-resolution sensors has been discussed and attested in some other applications [13],[18]. However, for the purpose of collecting data useful for GPS-mapping as well as road asset management, to the best of the authors' knowledge, no such method has been developed for data collection yet. The lack of a system capable of providing road asset data at a high frequency and a low latency is the main motivation behind the development of our distributed mapping solution called dMap. The system resulting from this research can provide large scale inventories of road assets for use in asset management as well as a broader group of mapping applications, such as GPS navigation systems, that provide road information to cars and drivers. Advanced Driver Assistance Systems (ADAS), that are becoming more common, often can make use of more rich data to improve the driving experience.

The current implementation of dMap uses a prototype sensor called dSensor which includes three RGB cameras, GPS, gyroscope, accelerometer, magnetometer, laptop and a wireless modem. The dMap concept consists of many but cheap dSensors mounted on the dashboard of public transport vehicles such as taxis and buses. These fleets of public transport vehicles continue to work on their own schedule while, simultaneously, collecting road asset information actively and frequently. Using public transport vehicles, not only is it possible to continuously monitor most of all the roads that carries a reasonable amount of traffic, the costs of the surveying vehicles and their drivers are trimmed entirely too.

The dMap data collection system is following a probabilistic model in which a uniform distribution of dSensors over a set of public vehicles is expected to get higher update rates of road assets on roads that have more traffic and more likely to have a higher priority/need for maintenance. By changing the distribution of sensors, the expected update rate on different parts of the road network can be controlled. Achieving

a competitive performance with such a design, requires new technologies to address many new challenges of the system, e.g., how to utilise low-cost sensors, limited communication network bandwidth, lack of adopted distributed light-weight image processing algorithms etc. While we envision the general concept of the dMap system in this paper, its current implementation includes the first prototype of dSensor, and the development of software drivers of different components of dSensor that facilitates interactions inside nodes. Many features of the dMap concept are left for future development.

The remainder of the paper is organised as follows. Recent works on distributed sensor platforms are discussed in Section 2. The requirements of a close to ideal mapping system for road asset management is presented in Section 3. The new dMap concept with specification of its first developed sensor as well as the current implementation of its software architecture and algorithms are described in Section 4. Conclusion of this paper is discussed in Section 5 and finally, the future works are presented in Section 6.

2. RECENT WORKS

There are numerous works on image processing and communication protocols over distributed ad-hoc sensor platforms. Most notably, Soro et al. [17] provide a survey about challenges and research trends in visual sensor networks. The survey includes communication protocols that are useful to utilise for image processing, light-weight algorithms able to run on resource constrained nodes, node management system that utilises the resources of each node based on different hardware architecture for each node, middle-ware support to facilitate interactions between nodes and provide system level functionalities, and finally distributed image processing algorithms that utilise such infrastructure using node level processes, collaboration between nodes and fusion centres. Another survey concerning mobile phone sensing [8] describes several opportunities that are emerging when mobile phones are increasingly equipped with sensors, e.g., gyroscope, accelerometer, compass and cameras. Among them are vision based applications that require some local processing before submitting a more abstract level of information to a server for further processing.

Lane et al. [8] suggest the need to integrate crowd-sourcing, social network and active learning for mobile platforms to enable developers to use the huge available infrastructure of mobile phones for applications in healthcare, environmental monitoring etc. A survey about the problem of coverage in video-based wireless sensor networks [6] discusses energy usage of sensor nodes and presents three modes of each visual sensor node; active, sleeping and off. Modes change in order to maximise the battery life of each node. In addition, it discusses some problems with graphical 2D reconstruction from multiple field of views of different cameras.

Costa et al. [6] show that a large number of low cost cameras, in general, provide better information from the environment compared to a few high-resolution sensors. Soro et al. [16] provide an algorithm to reduce data redundancy and energy consumption of a visual sensor network by selecting particular image frames from each camera of the network to construct a 3D field of view at a server. Wang et al. [19] describe how to use a Kalman filter-based tracking

algorithm to prevent detecting the same sign over several different frames. Their method skips repeated objects, and decreases the searching time. Similarly, Amirsadri et al. [4] simulate a mapping system for road sign detection that includes a fusion centre and couple of vehicles for data collection. They present an optimiser for communication between vehicles and the central server such that the vehicles send only newly detected signs to a server and avoid sending unnecessary information.

The sensor node localisation in many researches use GPS, accelerometers and gyroscopes. GPS receivers may not work properly in regions where the GPS signals are weak, e.g., tunnels and valleys. Accelerometer and gyroscope readings can help localise the sensor node using reference points provided by GPS in areas where the GPS signals are reliable. These techniques are widely used to estimate a more accurate location in different applications, e.g., Accelerometer Augmented Mobile Phone Localisation (AAMPL) for mobile phones [10], and cooperative GPS and internal systems [9] for vehicles.

3. REQUIREMENTS

Road asset management systems need to change in many ways from its current implementation to keep in pace with the recent changes in automotive industry and safety standards required for roads all over the world [15], [2]. Amid all the improvements needed, they should be more cost efficient such that government agencies, as their main customer, will be able to use them on scale.

The new system should be able to identify and geo-locate all road scene objects of interest including signs, rail guards, road potholes and puddles, traffic lights, road surface marks, poles etc. The latest state of the detected asset inventory should be represented in a database as well as lend itself to be visualised in a graphical map, or be exported in a format usable with other tool suites. For road asset management use, it is important to calculate an accurate location of each asset as compliance with regulations of asset placement is often required. However, for GPS mapping purposes, there is usually more flexibility in terms of geo-locating the object.

Different applications are concerned with different road side objects that usually differ between countries. Therefore developing region specific versions and supporting new objects should be a feature of the new system. New algorithms for image processing over a distributed platform may need to be developed. Such algorithms may concern themselves with maximising the utility of the available resources in the sensor network, simultaneously considering issues such as communication bandwidth, location of data relative to computing resources etc. This challenge should be distinguished from similar trend on computer vision architecture that focuses on multi-core and multi-threaded platforms [20],[5] where instead of the limitation being caused by wireless network bandwidth, other challenges such as communication time between processors and data storage devices emerge.

The new system should provide a platform for new distributed algorithms that are able to use the resources of both the sensor nodes and the central server. It should consist of a software framework that facilitates interaction between

Table 1: System Level Requirements of the dMap.

No	Requirements
1	Able to detect all classes of assets of interest.
2	Should be affordable and cost efficient compared to traditional methods.
3	Should be easy to use and deploy.
4	Support central management of the system.
5	Provide high update rates of road information.
6	Support region specific versions.
7	Support remote software updates.
8	Must be non-intrusive as to not distract the driver.
9	Should provide accurate geo-location estimate of individual assets.

Table 2: The Secondary Requirements of the dMap

No	Requirements
1	Provide a high level of interoperability between server and sensors.
2	Optimise the use of communication bandwidth.
3	Enable remote debugging facilities.
4	Provide easy to use sensor UI for interaction with the drivers.
5	Enable efficient distribution of tasks between sensors and the server.
6	Configurable for different modes of operation.

processes running on distributed hardware. It might embed Service Oriented Programming (SOP) or other technologies to increase interoperability and support different programming languages that are required for its development. In order to be maintainable, the system must support the ability to be remotely managed, allow remote software updates and be remotely debugged. It is not feasible to physically visit every installed unit so care must be taken such that the vast majority of issues can be dealt with remotely.

The new road asset management system should provide a friendly User Interface (UI) at both the sensor and the server location. The sensor UI should provide basic information of its running status. The drivers should be able to check if the device is working properly or if it requires some actions, e.g., the power cable is disconnected and need to connect again. At the same time, the sensor unit should be non-intrusive such that it does not cause any distraction for the driver. The server UI can be made arbitrarily advanced depending on which algorithms are being deployed. Being able to visualise the history of changes within the monitored area, performing software updates on one or more sensor nodes and track individual sensor node performance are just a few examples of what a server UI will need. A purposeful road asset management system need to provide updates to its users at a meaningful rate and without too much delay. In order to achieve this, the system should be able to estimate the performance in various parts of the monitored region and as a result either change the parameters of the deployed system, or suggest incorporating additional sensor nodes in certain areas. For example, certain sensor nodes may be remotely changed to perform tasks more or less often or possibly perform a completely different task. Table 1 lists the system requirements that should be fulfilled by an ideal road asset management system, and Table 2 lists the secondary requirements that can be derived from the system level requirements.

4. NEW MAPPING CONCEPT

Based on the many requirements discussed in Section 3, we are proposing a new road asset management system called dMap that consists of three main components; a set of dis-

tributed sensors called dSensor, a central server, and a software framework that facilitates implementing distributed algorithms.

The key feature of dMap is to replace relatively few, expensive, high-quality dedicated surveying vehicles with a much larger number of low-cost sensors. Today's smartphones include nearly every piece of hardware needed for the task; one or more cameras, wireless modem, GPS, accelerometer, gyro and magnetometer. This means that it is feasible to either integrate in a similar form factor, or in a few years time use off-the-shelf smartphones when they have reached enough computational power to perform the tasks. Each dSensor should be installed in public transport vehicles or other fleet vehicles that run on their own schedules. It should require only a power connection from the vehicles to operate, and after installation it should automatically perform a calibration cycle to calculate its position and orientation relative to the vehicle. The bulk of the data captured by the dSensor are the video inputs from the cameras that require a lot of storage capacity if saved and kept in raw format. The preferred solution is to temporarily buffer data on a local medium and only process it in order of priority. Some data may even be automatically discarded if it is considered unlikely to contribute any new information.

The extracted information should be queued and sent to a central server, prioritised in order of value, through wireless modem. Using public transport vehicles enables dMap to collect the information at a small fraction of the cost of traditional systems. However, since dSensors are not very powerful and at the same time, vehicles are not aiming to collect the information for the road asset management system but to carry out their own tasks, it is expected that the overall detection rate of road assets will be less than for conventional ones, on a per sensor node basis. This is overcome by dMap by using a large number of sensors for data collection that operate all the time and fuse their results in the central server to achieve a desired overall performance. The performance of dMap can further be improved by letting the central server ask individual sensor nodes to change the priority of their current tasks. For example, through analysis of the information that the central server has received, it may be plausible that there are assets that may have been missed or are poorly geo located. The central server can then change the use of relevant sensor nodes to prioritise processing that would help resolve such issues.

Apart from building a database of the current status of the monitored region, the server should provide the necessary interfaces to control the overall operation. That includes visualising the current status, request additional information from individual sensor nodes like raw sensory data if needed etc. Data transmission over the wireless modems is a bottleneck and the sensory nodes must at all times prioritise which information to buffer and which to send right away.

4.1 System Architecture

We believe a service oriented architecture is suitable for developing dMap providing high interoperability between dSensor and the central server, and following that we decided on using ROS (Robotic Operating System). The dMap design using ROS is shown in Figure 2. The dSensor will be

installed in public vehicles and should support a simple UI for the basic functionalities such as reset option so the driver can check if the sensor is running properly. Node manager is a key package in the ROS architecture which represents both the server and dSensor. The dSensor node manager sends the detected objects or other information to the server through another package that is called Communication.

The Error handler package at dSensor captures all the errors that may happen for dSensor at run-time and default actions are activated in the node manager to handle the errors. The node manager of each dSensor may receive information from different sensors, e.g., cameras, with different timestamps. Hence, it is needed to synchronise all the processes together, i.e., sync the location information with the video streams from the cameras through another package that we call Synchroniser as shown in Figure 2.

Although each dSensor detects and localises road scene objects from the video frames it captures, it might not be able to process all the sensory information because of its constrained resources. Therefore, the dSensor should utilise fast algorithms that detect most of the objects in a reasonable timeframe and also have an understanding of where in the map that information is of poor quality. It should be able to investigate the situation further by providing abstract information from its local understanding of the problem and request for further analysis that are possible with the powerful processing units of the server. We believe that in an efficient solution not only the server should be able to identify regions of poor quality, but the dSensor may contribute too. A package named Resource manager handles these requests at dSensor and optimally save and queue the information as well as requests for transmission to the server as shown in Figure 2.

4.2 The Current Implementation

dMap has three main components; dSensor, Communication and the Server. We have built a hardware prototype for dSensor, a software framework based on ROS for handling all the drivers of dSensor, utilised fast road sign detectors using road sign detection algorithms [11] for dSense, and considered techniques presented by Amirsadri et al. [4] for optimised communication.

The implemented sign detector is able to detect 50 types of Australian road signs of Australia using 50 sign detectors. It runs each object detector over an input video frame, over many scales, fuses their results and extracts the detected objects by calculating the size of the bounding box surrounding each cluster of detections. Figure 1 shows a stop sign being detected within an RGB image by the stop sign detector. The dSensor uses ROS as the underlying software framework for all the packages. Python is used to develop the UI package of dSensor whereas C++ is used for the other packages. The dSensor hardware consists of three RGB cameras, GPS, gyroscope, accelerometer, magnetometer, laptop and a wireless modem. It is designed to be installed in large numbers on dashboards of public transport vehicles including buses and taxis. Figure 3 shows the first prototype of the dSensor and Figure 4 shows an inside view of it with the arrangement of its components. The presented design of the dSensor is not intended for large scale deployment, but

Figure 2: Block diagram of dMap architecture.

Figure 3: First prototype of the dSensor.

Table 3: List of the components of the dSensor.

Name	Company	Description
MT9V023	Imaging Source Company	Wide VGA CMOS camera
ThinIMU	Future Technology Devices Int.	Measurement unit
3 Modem	Vodafone Australia Pty.	3G wireless modem
GM-158-232	SANJOSE Navigation	USB GPS

rather caters for flexibility while carrying out the research. The final product, when all the requirements are known, can easily be made significantly smaller and less intrusive. It also uses a laptop as processing unit that later will be replaced by efficient mobile CPUs and GPUs of similar type present in smartphones. The dSensor requires electrical power from the vehicle during its operation which includes the period it captures video, analyses the video and transfer its information to the server. Since locally buffered video is being processed, it is even possible for the dSensor to continue its analysis after the vehicle is turned off if necessary. The components of dSensor are low-cost, low-spec sensors, e.g., 744*480 pixel RGB cameras.

4.3 Evaluation Scenario

The dMap is not designed to collect high quality sensory data from roads as the other high-end surveying platforms do, but is instead relying on a probabilistic approach to achieving good performance resulting from multiple passes

Figure 4: Inside view of the dSensor. It consists of three cameras, wireless modem, accelerometer, gyroscope, magnetometer that all connect to a laptop.

of dSensors. Hence, in comparing performance, we should not compare a single pass of a dSensor with a single pass of a high-end surveying vehicle. Attributes that are interesting to monitor for comparison are; detection rate as a function of number of passes, update rate as a function of number of participating taxis, buses and other fleet vehicles and geo-location accuracy of detected objects as a function of number of passes. Given these attributes and costs of individual installations, it can be estimated what the total costs would be to deploy dMap at a very large scale to produce road asset information. For system development purposes, other meta data associated with each detection will also be monitored, e.g., the path and speed of each dSensor. These system level comparisons can only be made once the system is fully de-

Figure 5: Taxi trajectories are being studied within about 4130 km of roads in Beijing.

ployed, however, a precursor to these is an estimate of the road coverage as a function of number of vehicles and for how long they have been on the road. The following section uses real world taxi trajectory data to make such estimates.

4.4 Road Coverage

Road coverage is here defined as the proportion of the road network that is traversed by at least one vehicle during the sample period. Related statistics on the number of times a certain road segment is visited is also considered. Such measures are important in order to estimate the number of vehicles to equip with sensors to reach a particular level of minimum coverage and what average update rates are to be expected. The expected total detection rate is also directly dependent on road traversal frequency as the number of opportunities to detect an asset increases. A fair comparison against traditional systems requires an in depth investigation of the probabilistic behaviour of fleet vehicles, and at this stage we present the result of our studies that considers a limited number of taxi trajectories and road coverage scenarios.

We use the T-Drive taxi trajectory data sample from Microsoft [22] and [23] to illustrate the correlation between road coverage and the number of taxis as well as their period of operation. The original dataset includes data from 10357 taxis over a period of one-week. However, in this paper, we visualise the data from 24 randomly selected taxis that have been driven for 7 days, and that have generated about 1 million GPS points. We have mapped the location information using Nearest Neighbour (NN) to road segments after some noise filtering to remove erroneous data points. Figure 5 shows the urban area of Beijing that contains data from the taxi trajectories considered in this paper.

Figure 6 shows 24 taxi trajectories traversed over one week within Beijing. Each taxi trajectory is shown using a unique colour while Beijing roads are shown in pink. Both the number of taxis, and the number of days each taxi operates influence the road coverage percentage. Figure 7 shows how the road coverage increases over time for one taxi that op-

Figure 6: Trajectories of 24 taxis in Beijing.

erates for a week, and Figure 8 shows how road coverage increases if we fix the operation time of each taxi to 7 days but increase the number of taxis incrementally from 1 to 24.

Looking at figures 7 and 8, it might seem that an increase in the number of vehicles does not significantly contribute to a higher degree of road coverage. There are primarily two reasons that may have led to such a result. Firstly, although we chose the fleets randomly, there might still be a strong correlation between selected taxis as they might operate in specific areas. Since we use a small number of taxis in this preliminary study, there is a chance of strong correlation between taxi trajectories. Increasing the number of taxis from the current 24 and also including taxis from different companies that operate in different suburbs may help reducing the risk of correlation between taxis trajectories, and therefore increase the relation between road coverage and the number of taxis. Secondly, the road map information of Beijing that we use [3] includes only two main types of road from the five different categories of roads. Highways and main streets are covered but alleys and other narrow paths do not exist in the publically available version of the road map. Therefore taxis road trajectories could not be mapped to proper roads from the recorded GPS points and were removed from our study as noise. We have reason to believe that supporting all road types and mapping taxi trajectories to the actual roads that they traverse increases the estimated road coverage substantially.

Although using taxis and other non-dedicated fleets for road coverage may fail to cover 100% of roads, they provide very high frequency of coverage for main and busy roads including highways and city centres. Figure 9 shows how it is possible to traverse roads multiple times within one week using 24 taxis.

5. CONCLUSION AND DISCUSSION

Managing road assets is a very costly, but necessary, undertaking. The traditional road asset management systems are very expensive mainly because of their expensive data collection methods that require dedicated vehicles, expensive sensors and skilled drivers. Also, due to the high cost of data collection, it is not done very often and asset inventories become old and out of date. This paper presents a new

Figure 9: Road coverage heat map from 24 taxi trajectories in one week.

Figure 7: Road coverage changes over time using one taxi.

Figure 8: Road coverage correlation with number of taxis.

concept that differs from the traditional road asset management systems widely by its data collection process. A few high-end, high-spec surveying vehicles are exchanged for using many low-cost, low-spec sensors on public transport vehicles within a distributed mapping concept named dMap. The dMap can also be coupled with conventional mapping systems that can provide both an initial inventory as well as additional information contributing to dMap's probabilistic model of asset detection. There may be regions that are rarely passed by taxis, buses or other fleet vehicles and those regions may need to be specifically surveyed the conventional way.

While the dMap data capture model offers many advantages, it also introduces many challenges in computer vision and

optimisation in order to maximise the utility. With solutions to those challenges in place, it is expected that its deployment on large scale will cost many times less than the cost of conventional systems coupled with better performance along certain dimensions such as, in particular, update rate.

Given that installations of dSensors in host vehicles will experience a rather large variance of circumstances, it is of particular interest to investigate robust methods of self-calibration that will require no input from the driver. Also, in order to maximise the utility of the sensor nodes and associated compute capacity, in various parts of the road network as well as at different times a day, different parameters of the system might be needed depending on situation. It would therefore be interesting to investigate methods for automatic

adaptation of these parameters. Examples of such parameters are frame rate, number of scales in the image that is scanned, which set of object detectors to run at which times etc.

Further, methods are needed for the central server, with knowledge of the complete current state of the asset database, how to best advise individual dSensors to perform optimally. Given rules about how certain assets are supposed to be placed, it might be possible to infer the existence of a sign given the other observations. This can be used to request every dSensor that passes that region to spend additional resources (higher frame rate, more scales, smaller step size in image) in order to confirm its existence. Inconsistent observations in a region made at different times, by different dSensors, may also be used as a trigger to spend additional resources there. In particularly challenging cases, that need to be automatically identified, it makes sense to flag it as a problem that need intervention from the dMap operator which can then request to have raw images sent to the central server for manual inspection. Incorporating information in a sensible way from other modes of capture might also be challenging. Crowd-sourcing in various forms can likely add considerable value to the overall dMap system.

6. FUTURE WORK

There are many desired requirements/properties for road asset management systems as presented in Section 3 and while the dMap concept that is presented in Section 4 does not fall short of the requirements by its definition, its current implementation is limited to twelve prototype units of dSensor and implementation of its low-level services. Future work consists of the completion of the implementation of the dMap architecture, and the verification of the fulfilment of the requirements.

7. REFERENCES

- [1] National motor vehicle crash causation survey, report to congress. July 2008. Available at: <http://www-nrd.nhtsa.dot.gov/Pubs/811059.PDF>.
- [2] Moving ahead for progress in the 21st century act (map-21). *The Federal Highway Administration (FHWA)*, 2012.
- [3] Openstreetmap community of mappers. 2014. Available at: <http://www.openstreetmap.org/>.
- [4] A. Amirsadri, A. Bishop, J. Kim, J. Trumpf, and L. Petersson. A computationally efficient low-bandwidth method for very-large-scale mapping of road signs with multiple vehicles. pages 1351–1358, july 2012.
- [5] N.-M. Cheung, X. Fan, O. Au, and M.-C. Kung. Video coding on multicore graphics processors. *IEEE Signal Processing Magazine*, 27(2):79–89, 2010.
- [6] D. Costa and L. Guedes. The coverage problem in video-based wireless sensor networks: A survey. *Sensors*, 10(9):8215–8247, 2010.
- [7] J. Gilles and Y. Meyer. Properties of bv - g structures + textures decomposition models. application to road detection in satellite images. *IEEE Transactions on Image Processing*, 19(11), 2010.
- [8] N. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. Campbell. A survey of mobile phone sensing. *IEEE Communications Magazine*, 48(9):140–150, 2010.
- [9] V. Milans, J. Naranjo, C. Gonzlez, J. Alonso, and T. De Pedro. Autonomous vehicle based in cooperative gps and inertial systems. *Robotica*, 26(5):627–633, 2008.
- [10] A. Ofstad, E. Nicholas, R. Szcudronski, and R. Choudhury. Aampl: Accelerometer augmented mobile phone localization. *Proceedings of the Annual International Conference on Mobile Computing and Networking, MOBICOM*, pages 13–18, 2008.
- [11] G. Overett and L. Petersson. Large scale sign detection using hog feature variants. *IEEE Intelligent Vehicles Symposium, Proceedings*, pages 326–331, 2011.
- [12] F. Qu, F.-Y. Wang, and L. Yang. Intelligent transportation spaces: Vehicles, traffic, communications, and beyond. *IEEE Communications Magazine*, 48(11):136–142, 2010.
- [13] M. Rahimi, S. Ahmadian, D. Zats, R. LAufer, and D. Estrin. Magic numbers in networks of wireless image sensors. pages 71–81, October 2006.
- [14] Y. Rianto, T. Kim, and S. Kondo. Detection of roads from satellite image using optimal search. *IEEE Asia-Pacific Conference on Circuits and Systems - Proceedings*, pages 587–590, 1998.
- [15] A. Schanz and R. Hamann. Towards "vision zero. *SAE Technical Papers*, 2012.
- [16] S. Soro and W. Heinzelman. Camera selection in visual sensor networks. *2007 IEEE Conference on Advanced Video and Signal Based Surveillance, AVSS 2007 Proceedings*, pages 81–86, 2007.
- [17] S. Soro and W. Heinzelman. A survey of visual sensor networks. *Advances in Multimedia*, 2009, 2009.
- [18] N. Tenzcan and W. Wang. Self-orienting wireless sensor networks for occlusion-free viewpoints. pages 2558–2567, 2008.
- [19] K. Wang, Z. Hou, and W. Gong. Automated road sign inventory system based on stereo vision and tracking. *Computer-Aided Civil and Infrastructure Engineering*, 25(6):468–477, 2010.
- [20] P. Wang, J. Collins, G. Chinya, H. Jiang, X. Tian, M. Girkar, N. Yang, G.-Y. Lueh, and H. Wang. Exochi: Architecture and programming environment for a heterogeneous multi-core multithreaded system. *ACM SIGPLAN Notices*, 42(6):156–166, 2007.
- [21] W. Wijesoma, K. Kodagoda, and A. Balasuriya. Road-boundary detection and tracking using lidar sensing. *IEEE Transactions on Robotics and Automation*, 20, 2004.
- [22] J. Yuan, Y. Zheng, X. Xie, and G. Sun. Driving with knowledge from the physical world. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '11*, pages 316–324, New York, NY, USA, 2011. ACM.
- [23] J. Yuan, Y. Zheng, C. Zhang, W. Xie, X. Xie, G. Sun, and Y. Huang. T-drive: Driving directions based on taxi trajectories. In *Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems, GIS '10*, pages 99–108, New York, NY, USA, 2010. ACM.