

Sensor Networks For Smart Roads

Marcin Karpiński, Aline Senart and Vinny Cahill
Distributed Systems Group
Department of Computer Science
Trinity College, Dublin, Ireland
{karpinsm, aline.senart, vinny.cahill}@cs.tcd.ie

Abstract

This paper proposes the use of a wireless sensor network of “cat’s eye” augmented with embedded processing, communication, and sensing capabilities to monitor vehicle behaviour on augmented roads. The primary goal of the system is to provide drivers with early warning of potentially dangerous situations that may arise. The focus of the paper is on the software architecture needed and the technical challenges to be overcome in order to support this and related applications.

1. Introduction

While in recent few years a number of projects [13, 8, 2, 3] have proposed the use of sensor networks for vehicle tracking in military applications we propose that a distributed and autonomous system of sensor network nodes be deployed with the goal of improving driving safety on public roads. According to the Irish National Roads Authority [1], 41% of all two-vehicle only accidents were caused by travelling on the wrong side of the road and by improper overtaking. Because of that the purpose of our system is to provide drivers with a consistent view of the road situation a few hundred meters ahead of them, so that they can react to potential dangers early enough. Also, because 72% of all road accidents occurred outside of urban environments, and 42% were on national roads we consider those roads as the primary target deployment environment for our system.

2. The Concept

The basic idea is as follows: sensor nodes are placed along both sides of a road every few meters. We focus on two-way single carriageways only, because almost 80% of all road accidents in Ireland (see [1]) occur on this type of road. The nodes are to be placed inside cat’s eyes and

are equipped with magnetic sensors. After deployment they form an ad hoc radio network to exchange information about passing cars, as determined by their magnetometers, between each other. The nodes need to cooperate by means of communication to maintain consistent real-time local road-state information, which consists of the relative positions and speeds of all of the vehicles travelling along the road at a particular time and place. This information is then communicated to the vehicles, so that on-board computers can use it to infer dangerous situations. The amount of reasoning and computation that needs to be done by sensor nodes and vehicle on-board computers is to be investigated. To make our ideas more clear we consider a typical improper overtaking scenario that could possibly be avoided by use of our system: a vehicle tries to overtake another vehicle on a road curve, where it cannot be seen by other vehicles travelling from the opposite direction. This situation is depicted in Figure 1. The sensor nodes forward the information about the vehicles they have detected along the road and drivers can be informed early enough to react safely.

The system should be easily deployable and robust. The nodes should be replaceable and failure of single nodes should not affect operation of the system as a whole, which means that node failure should have impact only on local system operation. It should also be possible to replace broken devices without any additional system reconfiguration. The newly installed ones should detect their neighbours, and quickly find their place within the whole system.

3. Other Applications

Although the focus of our system is to improve driving safety, we can also consider additional applications. In further research we envision a system with enhanced functionality that comprises of a configurable number of *services*. The particular set of services comprising a particular system instance would be chosen at the design level, with respect to such factors as node hardware or power constraints




Figure 1. Possible application of sensor networks to safe driving monitoring.

or simply the need for particular functionality. We expect that this system, besides the basic road-state information service, would also provide:

1. Vehicle tracking over longer distances for travel-time estimation. By correlating spatio-temporal detection data, vehicles can be tracked along the road. This requires further investigation.
2. Vehicle counting at intersections. This service is of major importance in cities, where traffic volume information is used by traffic management systems.
3. Pedestrian on-road presence notification. This service can be provided assuming the use of more sophisticated sensors or pedestrians equipped with some device to announce their presence on the road.
4. Erratic driving detection and driver warning. Assuming a bigger computational capability, the sensor nodes can make use of the road-state information they generate (the vehicle tracking service) to infer dangerous

driving patterns (e.g., speeding or driving in the wrong direction) and warn drivers by using some actuators, e.g., LEDs in cat's eyes. In the case of this service the reasoning would have to be done on the sensor network side.

5. Incident detection. A service similar to erratic driving detection, with the difference that the nodes try to infer an incident (e.g., a crash) and forward that information along the road to warn other drivers or to send it to a traffic management centre, by communicating it to a more powerful device capable of long-range communication.
6. Periodic traffic volume snapshot. This service counts or estimates the number of vehicles along the whole road at some particular time and sends it to a traffic management centre to make it available to other users (e.g., through the Internet). Periodic service execution would allow for real-time traffic information.
7. Road-surface condition reporting service. Equipping sensor nodes with additional sensors like temperature and humidity sensors would allow to periodically gather information about the road-surface condition. The information could be sent to a traffic management centre, displayed on a variable message sign or communicated to passing vehicles.
8. Message broadcast service. Nodes can also forward messages along the road and broadcast them to passing cars. Their purpose would depend on the particular application, e.g., they could be warning messages generated by a traffic management centre.

The above listed services are meant as an extension to the basic functionality, which is the vehicle tracking service. Their development will be undertaken in the second phase of our research.

4. Software Architecture

Since the sensor nodes are to provide various services, the software architecture of the nodes must be modular to allow for service set composition and further service enhancement. The software architecture of a single node has to consist of two main layers: a lower layer containing a set of low-level software components that provide functionality common to the set of implemented services, and a higher layer defining a subset of the services listed above. We assume that computing constraints of the nodes require the architecture of the lower layer to be component-based, so only the minimum set of components can be chosen, with respect to the functionality required by the services of the higher layer.

All software components developed during our research will constitute a highly configurable component-based middleware, whose architecture will be an instantiation of the architecture developed by the Aithne project¹, a project investigating novel approaches to building distributed applications composed of a large number of autonomous devices physically dispersed in space and cooperating to achieve common goals. Central to the Aithne middleware is a concept of the *sentient object* [5] which models a single system entity (a device). Sentient objects are typically equipped with a range of sensors and wireless communication. They gather knowledge about the surrounding environment through sensor fusion and they collaborate with each other by means of communication. We will model sensor nodes as sentient objects and apply Aithne middleware architectural guidelines during system design and development.

At this early stage of the project we envision the following groups of lower layer software components being implemented:

- System deployment: components responsible for self-initialisation of a node including neighbour discovery, initial time synchronisation and localisation relative to node's neighbours.
- Sensor algorithms: components containing sensor data processing algorithms, for instance, signal processing, sensor calibration, vehicle detection, categorisation and tracking. All based on the output of a magnetometer and possibly a PIR sensor.
- Actuator drivers: if nodes were to be equipped with some actuators (e.g., LEDs) those components would provide routines for using them.
- Communication: components defining the communication sublayer. They will include a MAC protocol, a (possible) routing protocol, common node intercommunication module, node-to-vehicle communication module.
- System operation: components providing common services used during system operation. This might include localisation and network topology maintenance, neighbour monitoring for node failure and new node discovery.

It is very important that the whole architecture has an event based structure, because the operation of the nodes is driven by environmental events such as vehicles driving on the road. Also, the energy saving requirement forces us to put the nodes in a stand-by mode as often as possible, so their operation would be triggered in many cases by timer events.

¹<http://www.dsg.cs.tcd.ie>

Additionally, the organisation of the higher layer can be hierarchical, because we could build some services upon others. A good example is the road-state maintenance service, which would be the basis of many more sophisticated services such as erratic-driving detection.

5. Implementation Issues

The most fundamental problem in sensor networks is energy consumption. Due to constrained battery capacities, sensor nodes need to save energy to prolong their life time. It is usually done by duty cycling and minimising the amount of communication needed to achieve system's goals, as well as, reducing sensor reading time. Although very important, it always comes at a price - system response time or accuracy are the common trade-offs (e.g., see [15, 16, 7, 12]). In our case responsiveness and accuracy are of most importance. Hence we expect the smart road system to be energy-efficient enough to be powered by solar cells rather than batteries which would require further reductions in its functionality or performance.

The first optimisation we want to perform is to use a fixed node localisation scheme. To accurately estimate vehicle speeds all nodes have to know their relative order and precise distances to their neighbours. Nevertheless, once deployed they are not expected to move in any way, so it makes sense to assign location estimates to each of them once during system deployment allowing them to then infer their relative order. To ease the actual deployment we will evaluate the use of photo sensors and a GPS receiver to quickly assign nodes their positions. The idea has recently been elaborated in [14] and we reduce it to the following setting: each sensor node has an on-board photo detector, a user installing sensors in the road surface carries a PDA device equipped with a GPS receiver, the user touches the photo sensor with the PDA's screen which, in turn, by blinking black and white in a certain way transmits its position to the sensor node, as determined by the GPS receiver. The sensor node transforms photo sensor's readings to a bit sequence which encodes its position. The whole situation has to be initiated by the PDA by sending a broadcast message to all nearby nodes which then start sampling their photo sensors expecting the location data.

To achieve system accuracy that is sufficient to support all applications listed in Section 4 we need to take advantage of the structure of the smart road network. Vehicle or, in general, object tracking has been present in the literature for a long time. Inspired by military applications, the usual approach assumes placing sensor nodes randomly in the field and sensing for objects moving in arbitrary directions. The nodes are meant to collaboratively compute object positions and send it to a gateway node that further reports it to a control centre. Such a random setting in-

creases the computational complexity of the detection algorithms and reduces their efficiency. Limitations include bounding maximum object speed and requiring different objects to be located far enough apart so they can be disambiguated (see, for example, [8, 9]). Our scenario assumes sensor nodes equidistantly placed along roads. This constrains possible network configurations to a particular family of graphs. Also, vehicles move only in one of two directions which simplifies the detection algorithms - there is no need to run computationally expensive algorithms for bearing estimation and prediction of most probable target presence area in the near future (e.g., [17]). By taking advantage of these simplifications we want to achieve lower system response times and be able to track vehicles (or platoons of vehicles in situations in which distinguishing between them is too expensive) moving with ordinary road speeds (i.e., around 90 km/h).

6. Related work

Object detection classification and tracking are some of the most basic sensor network applications. They originate from military needs to detect intruders and track enemy activity in the battlefield. Thus, most of the scientific projects investigating these issues were inspired by such applications. As a consequence, such systems have very specific requirements, as sketched in the previous section, and the developed algorithms cannot be universally applied to all object tracking applications.

Before sensor networks came into existence, there had been a large body of research in the area of algorithms for object tracking and the algorithms were mostly designed for centralised systems (a review of the most popular techniques can be found in [11]). With the introduction of the sensor network idea a number of theoretical investigations have been carried out ([9] provides a good overview) and number of real system implementations and experiments have been performed.

So far the most advanced field demonstration of a vehicle tracking system has been done by PARC's CoSense² project. The researchers used seismic and acoustic sensors and have shown that it is possible to track a single object with a reasonable accuracy. Theoretical underpinnings and experiment results can be found in [10, 17]. Another approach to object tracking which also aimed, apart from vehicles, at detecting soldiers and civilians was taken by the *Line in the Sand* [3] project. The researchers used CrossBow's³ Mica2 motes equipped with custom-built mini-radar sensor boards and magnetic sensors from Honeywell⁴. They showed successful tracking of soldiers and vehicles moving with speeds

²<http://www2.parc.com/spl/projects/cosense/>

³<http://www.xbow.com>

⁴<http://www.ssec.honeywell.com/>

not exceeding 3-25mph. A different approach was taken by scientists from Berkeley, who demonstrated a scenario in which a number of small sensor nodes (the René mote, a predecessor of the Mica platform) were dropped from an unmanned aerial vehicle and formed an ad hoc radio network to detect a number of passenger and military vehicles using magnetic sensors [4] only. However, not many details were given about the accuracy of the system and algorithms used. Another vehicle detection system of 70 Mica2 motes with magnetic sensors has been implemented and demonstrated in [8]. The sensor nodes were placed on both sides of a field road and were able to track military vehicles. The main drawback of the system was the assumption of targets being separated so no sensor node sensed two different objects at the same time.

To our knowledge, the only non-military application of sensor networks to detect vehicles was presented in [6]. The authors proposed embedding a number of MicaDot nodes (also from Crossbow) with magnetic sensors on board to count vehicles approaching intersections. The idea was to replace more traditional traffic sensors, such as inductive loop detectors. This approach is closest to the ours, but is much simpler since the authors were interested only in the total number of vehicles waiting at a junction. Also, a similar idea of placing sensors in the road surface has been employed by Austucia⁵ which manufactures devices called *road studs*. They are equipped with light, humidity, and temperature sensors and also with a number of bright LEDs that change colour depending on the weather conditions sensed by the sensors. However, there is no communication between the devices so they are not considered to be a sensor network. Road studs are powered by solar cells.

7. Summary

This paper presents the concept of a smart road system, that is able to accurately track vehicles moving on national roads with ordinary road speeds. The system is to be deployed over long distances and actively respond to drivers' activity. It should be fully scalable and its architecture allow for extension of its basic functionality, so that it can be tailored to the needs of local authorities that decide to deploy it. The idea falls into the area of sensor networks, but the purpose of the system relates it to the fields of ubiquitous computing and ambient intelligence. We believe that in the not distant future such systems will become an everyday experience and our research is one step towards that vision.

⁵<http://www.astucia.co.uk/>

References

- [1] *Road Collision Facts*. National Roads Authority, Ireland, 2003.
- [2] A. Arora. Exscal: Extreme scale wireless sensor networking - <http://www.cast.cse.ohio-state.edu/exscal/>, 2005.
- [3] A. Arora, P. Dutta, S. Bapat, V. Kulathumani, H. Zhang, V. Naik, V. Mittal, H. Cao, M. Gouda, Y. Choi, T. Herman, S. Kulkarni, U. Arumugam, M. Nesterenko, A. Vora, and M. Miyashita. A line in the sand: A wireless sensor network for target detection, classification, and tracking. In *Computer Networks*, 2004.
- [4] U. Berkeley. 29 palms fixed/mobile experiment, website on vehicle position monitoring: Tracking vehicles with a uav-delivered sensor network. <http://www.eecs.berkeley.edu/pister/29palms0103/>, 2001.
- [5] G. Biegel and V. Cahill. A framework for developing mobile, context-aware applications. In *2nd IEEE Conference on Pervasive Computing and Communications, Percom 2004*, Orlando, Florida, March 2004.
- [6] S. Coleri, S. Y. Cheung, and P. Varaiya. Sensor networks for monitoring traffic. January 2005.
- [7] C. Gui and P. Mohapatra. Power conservation and quality of surveillance in target tracking sensor networks. In *MobiCom*, May 2004.
- [8] T. He, S. Krishnamurthy, J. A. Stankovic, T. Abdelzaher, L. Luo, R. Stoleru, T. Yan, and L. Gu. Energy-efficient surveillance system using wireless sensor networks. In *ACM MobiSys*, June 2004.
- [9] D. Li, K. Wong, Y. Hu, and A. Sayeed. Detection, classification, and tracking of targets. In *IEEE Signal Processing Magazine*, March 2002.
- [10] J. Liu, J. Reich, and F. Zhao. Collaborative in-network processing for target tracking. In *EURASIP Journal on Applied Signal Processing*, 2003.
- [11] E. Mazor, A. Averbuch, Y. Bar-Shalom, and J. Dayan. Interacting multiple model methods in target tracking: A survey. In *IEEE Transactions on Aerospace and Electronic Systems*, 1998.
- [12] S. Patten, S. Poduri, and B. Krishnamachari. Energy-quality trade-offs for target tracking in wireless sensor networks. In *2nd Workshop on Information Processing in Sensor Networks*, April 2003.
- [13] G. Simon, M. Maróti, A. Lédeczi, G. Balogh, B. Kusy, A. Nádas, G. Pap, J. Sallai, and K. Frampton. Sensor network-based counter-sniper system. In *2nd ACM Conf. Embedded Networked Sensor Systems*, November 2004.
- [14] R. Stoleru, T. He, J. A. Stankovic, and D. Luebke. A high-accuracy, low-cost localization system for wireless sensor networks. In *ACM SenSys*, November 2005.
- [15] R. Szewczyk, A. Mainwaring, J. Polastre, and D. Culler. An analysis of a large scale habitat monitoring application. In *ACM SenSys*, November 2004.
- [16] Y. Yu, B. Krishnamachari, and V. Prasanna. Energy-latency tradeoffs for data gathering in wireless sensor networks. In *INFOCOM*, March 2004.
- [17] F. Zhao, J. Liu, J. Liu, L. Guibas, and J. Reich. Collaborative signal and information processing: an information-directed approach. In *Proceedings of the IEEE*, August 2003.