Low-cost Road Surface Temperature sensing enabled by the Internet of Things

L. Chapman, S.J. Bell
The University of Birmingham, UK

Corresponding author's E-mail: l.chapman@bham.ac.uk

ABSTRACT

Winter maintenance engineers base their nightly decision making by consulting a Road Weather Information System (RWIS) which combines weather forecast data with road temperature and condition data. The current generation of RWIS is based on route based forecasts which take into account how the local geography interacts with the regional climate to produce a model of road surface temperatures for every 50m section of road. By knowing which road sections are likely to fall below the 0°C threshold on a night-by-night basis, highway engineers can selectively treat just the affected routes and make significant savings in salt usage. However, this is not happening. In an environment of increasing litigation, practitioners are nervous about making decisions based on model output as opposed to ground truth and means that the verification of route based road weather forecasts is urgently needed. Research at the University of Birmingham seeks to solve this problem by producing and deploying a new generation of low-cost, internet-enabled, road surface temperature sensors embedded within an Internet of Things ecosystem. This approach will not only provide a monitoring and verification solution, but also has the potential to form the basis of a new generation of 'nowcasting' models for winter road maintenance. This paper investigates the feasibility of harnessing the Internet of Things to develop a high resolution, but low cost, road surface temperature monitoring network. It focuses on the description and evaluation of a low-cost (<$200), self-contained sensor, which has now been tested in both a lab and field setting.

1 Introduction

The first generation of Road Weather Information Systems (RWIS), still frequently in use, relies on methods and tools developed in the 1980s. Early 'Ice Detection Systems' simply comprised of a small number of outstations which detected when ice formed upon the integral road sensor. This approach was limited for two reasons. Firstly, as the ice had already formed, it was too late to pre-salt the network. Secondly, the sensor only provided 'spot' measurements of road surface condition (RSC) and road surface temperature (RST) at a single location - it was then generally unknown how representative this location was in comparison to the remainder of the network. This lack of spatial information was the motivation behind the development of thermal mapping techniques (e.g. Gustavsson, 1999). Thermal mapping utilised an infrared thermometer to conduct a thermal survey from a moving platform driven around a road network. RSTs were typically taken at a 20m resolution to provide the highway engineer with the knowledge of local thermal variations. Throughout the 1980s, these technologies matured into an 'Ice Prediction System' which via an energy balance model, enabled RST and RSC to be forecast for each outstation before being interpolated with a thermal map. This system quickly spread across the globe and has been considerably refined as computer processing power improved. However, as technology moved on, the existing system began to look very dated and started to be superseded by route-based forecasting techniques e.g. XRWIS (Chapman & Thornes, 2006). XRWIS takes into account how the local geography interacts with the regional climate to produce a model of RST for every 50m section of road. By knowing which sections of road are likely to fall below the 0°C threshold on a night-by-night basis, highway engineers can selectively treat just the affected routes and thus make significant savings in salt usage. The forecast is displayed in a GIS environment and disseminated direct to the highway engineer via the Internet (Fig. 1).

Savings from using this approach are realised by adopting selective salting practices (i.e. treat only the coldest sections of road) and dynamic routing (thermally optimised salting routes that change daily in relation to the weather forecast: Handa et al. 2006). *However, it is here where the economic benefits of the technology have stalled*. To adopt selective salting and dynamic routing requires full confidence in the available technology. In an environment of increasing litigation, practitioners are nervous about making decisions based on model output as opposed to ground truth. This means that the verification of road weather forecasts is now sorely needed and it is only with this step that end users will be responsive to selective salting ultimately unlocking the potential for generating financial savings from route-based forecasts. However, existing verification techniques are simply not up to the job (Hammond et al, 2010). Point measurements using embedded sensors are expensive to install and lack the spatial resolution, where as mobile measurements lack the temporal resolution to provide the full picture. Fortunately, it has been proposed that a new generation of 'low cost' sensors embedded within an IoT ecosystem can solve the verification problem (Chapman et al, 2014)

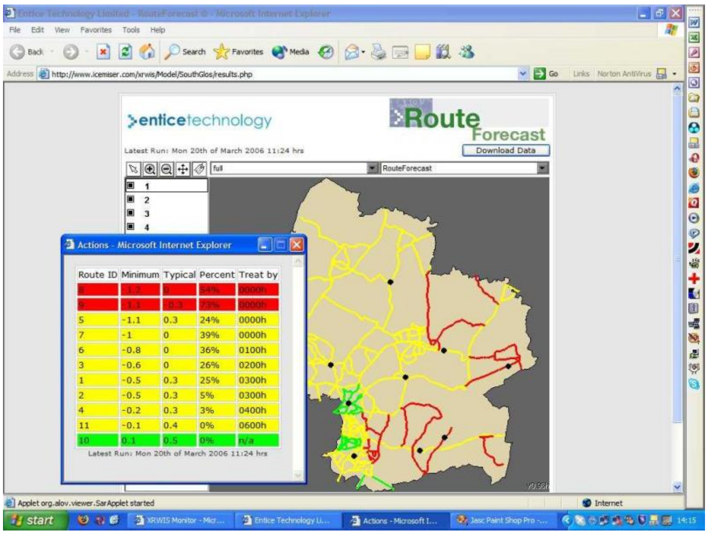

Figure 1: Screenshot of XRWIS. The fundamental difference between RWIS and XRWIS was the change from a reliance on thermal mapping to an increasingly high resolution modelling approach. Instead of running a model for a handful of outstations, the model was now run for thousands of sites across a network just metres apart. This is how these maps are produced – these consist of point data and not the line data which is how the visualisation appears.

The Internet of Things (IoT) quite literally means 'things' (e.g. sensors and other smart devices) which are connected to the internet. Since 2008, the number of 'things' has outnumbered users online. Chapman et al (2014) presented a viewpoint on the future role of the IoT in winter road maintenance. Using low cost (<$100) air temperature sensors developed on a previous project, they showed how dense network of sensors can easily be deployed on road corridors utilising existing WiFi communication networks. This paper documents how a road surface temperature sensor has now been produced and deployed based on the principles of the IoT.

2. A Prototype Sensor

The University of Birmingham have developed a road surface temperature sensor based upon the Melexis infrared thermopile MLX90614. This 5° field of view sensor has a small size (~4mm diameter), and has a range of -40 to 125 °C for sensor temperature and -70 to 380 °C for object temperature with a resolution of 0.1°C or better and a reported accuracy between 0.5°C to 1°C. It consumes very little current which is appropriate for the proposed application with power provided by a specialist single cell AA 3.6 V Lithium-Thionyl Chloride battery capable of providing the peak current (150 mA) required for radio transmission and under good wireless network conditions is reported to last up to 2 years. Data communication uses a low-powered wireless communications card using standard IEEE 802.11 b/g/n 2.4 GHz WiFi at bit rates of up to 11 Mbps. As such, the sensor can be deployed anywhere within range of a wireless network from which data packets can be transmitted periodically through the internet (or a local adhoc network) to software installed on a server. Overall, the sensor is self-contained within a weather proof enclosure (Figure 2) and requires no additional power or communications. This, coupled with its low cost (approx. $200 per unit) enables a large number of the devices to be rapidly deployed in a network. Using traditional sensors, communication costs can become significant when deployed in large numbers. However, given the increasing proliferation of WiFi networks in municipal areas, then the advantages of using IoT based sensors becomes apparent (Chapman et al, 2014)

Figure 2: Prototype road surface temperature mounted, as operationally, on street furniture adjacent to a road. The small black hole is the MLX90614 aperture, which is aimed at the road surface.

3. Laboratory Testing

Before deployment in the field, the sensor was tested using a new state of the art Weiss WKL 34 climatic chamber at the University of Birmingham. The chamber comprises a heat exchanger, electric heater and recirculating air fan along with a sophisticated control system capable of generating reproducible and stable temperature conditions from −40 °C to +180 °C within its 34 litre chamber. Temperature constancy in time is ±0.3 °C with cooling and heating rates of 6 °C/min and 4 °C/min respectively. Humidity can also be controlled in the range 10 – 98% with a built-in humidification/dehumidification system.

Figure 3: The Weiss WKL 34 climate chamber. Inside are the four test sensors mounted above the slab of asphalt into which the reference thermistor is embedded.

To simulate operational conditions, test sensors were aimed at a slab of asphalt (22 cm width, 3 cm depth) contained within a domestic baking tray. As the sensor corrects for the estimated emissivity of the target surface it is crucial the asphalt closely resembled that of public roads. As is customary for infrared road temperature sensors, sensors were configured to assume the asphalt surface has an emissivity of 0.95. To accurately measure the temperature of the asphalt slab, a reference thermometer was embedded in the centre at a depth of 2 mm. This thermometer is a flexible thermistor probe from Tinytag, which was calibrated 5 weeks prior to the chamber test using a reference meter in a UKAS approved laboratory. Four prototype sensors were tested (henceforth referred to as S1, S2, S3 & S4). Each mounted vertically with a 3 cm gap between *MLX90614* aperture and the asphalt. By positioning the sensors vertically side by side they each aim at different spot on the surface, with each spot ~5 cm from the reference thermistor in the centre. The test therefore relies on a constant temperature across the asphalt slab.

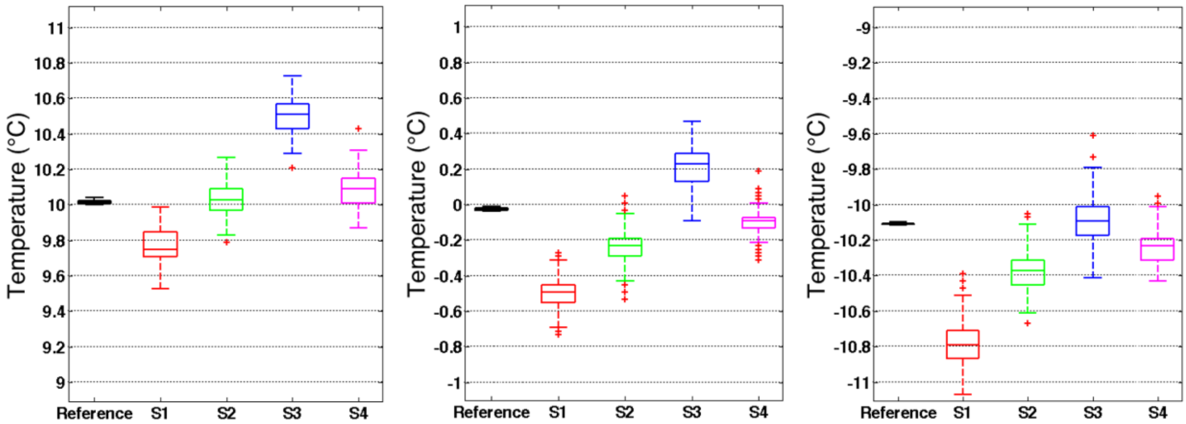

Figure 4: Box plot summary of temperature readings at three different chamber temperatures (+10 °C, 0°C and −10°C) by the reference thermistor and four prototype sensors. Each box represents data collected over 1 hour.

Each climate chamber test took a total of 12 hours to complete (comprising of three 4-hourly periods). The temperature of the chamber was set differently for each period. +10 °C, 0°C and −10°C were selected as winter maintenance engineers rarely deal with temperatures outside of this range. Only data from the last hour of each 4-hour period was used: the first 3 hours are solely to allow the asphalt slab to reach equilibrium with air temperature within the chamber (it takes < 10 mins for air temperature to reach the target temperature). With the Tinytag sampling every 5 seconds, and our sensors every 15 seconds, 720 and ~240 data points were recorded respectively over each hour. The results of the test are summarised in Figure 4 and Table 1.

Table 1: Mean temperature bias for each of the four sensors (S1, S2, S3, S4) at each of the three test temperatures relative to observations made by the reference thermistor (only reference observations within 3 seconds of infrared sensor observations are used). The sample size for each cell is ~240, with the value in brackets denoting the standard deviation of the bias.

REFERENCE (°C):	10.02	-0.02	-10.11
S1	-0.24 (0.09)	-0.47 (0.10)	-0.67 (0.12)
S2	0.01 (0.09)	-0.21 (0.10)	-0.27 (0.11)
S3	0.49 (0.09)	0.24 (0.11)	0.02 (0.13)
S4	0.07 (0.09)	-0.07 (0.08)	-0.13 (0.09)

The variance of the reference thermistor (black box plot in Figure 4) is small, confirming that the asphalt slab reaches a stable state of equilibrium. Whilst the climate chamber uses its own thermometer to regulate its internal temperature, this internal thermometer actually has a slight calibration bias explaining why the more reliable Tinytag reference thermistor deviates slightly from the target temperatures of +10 °C, 0°C and −10°C. Whilst the median and mean of the prototype sensors rarely agree exactly with that of the reference thermistor, they do however fall within acceptable limits of ±0.7 °C. The sensor's firmware includes the option to remotely apply a constant bias correction if required. As 0 °C is a critical threshold for winter road maintenance the bias measured during a climate chamber test at this temperature would be most appropriate to use for this correction. Observations from the test sensors are much noisier than that of the reference thermistor, but still acceptable. Interestingly, at all three test temperatures, S1 always reads coldest on average, followed by S2, then S4, with S3 the warmest. However, relative to the reference thermistor they all drift to a cooler bias as the chamber becomes colder. This experiment only tests the performance of the sensors when air, asphalt, and sensor temperatures are all in equilibrium. In further experiments the performance of the sensor will be tested when air and sensor temperature are significantly different from that of the asphalt surface.

4. Field Deployments

The sensor has also been operationally tested at a number of sites across the UK, but the primary testbed remains the Hagley Road winter road maintenance testbed (see Chapman et al, 2014). A total of 24 sensors have been deployed along this arterial road running from Birmingham City Centre to the rural countryside (Figures 5 and 6). Sensor locations were chosen to be close to the access points of a recently installed proprietary roadside telecommunications network (WiFi based Wireless Mesh) that connects roadside equipment across the city (e.g. traffic signals, message signs, car park counters, CCTV cameras) via the internet.

5. Discussion & Conclusions

Despite the low cost nature of the sensor, sensor performance in both lab and field trials is very positive. One frequent concern with low-cost sensors is sensor drift and given the age of the sensor, this has been impossible to fully ascertain in this study. However, during the trials over winter 2015/16, no significant issue has been discovered. Overall, this work has confirmed that it is certainly possible to produce a low-cost sensor which is fit for purpose, however further research and development is needed in the IoT space, as well as consideration as to how the sensors will be maintained.

As highlighted by Chapman et al, 2014, there are a number of challenges with the IoT approach. Reliable communciations are particularly important, as intermittent WiFi signals can cause the internal sensor batteries to quickly drain. Wireless access continues to improve but also remains a limiting factor. However, this is expected to change rapidly in a short space of time given the advent of a new generation of Sub-GHz communication networks. Once these problems are overcome, it is evident that the IoT will deliver the data long needed to verify and improve RWIS. In the meantime, WiFi Hotspots can deliver localised connectivity and sensors can still be used in a logging mode where data can be later retrieved by simply pressing a button on the device (data is routed to the cloud via the button pressers mobile phone). Indeed, this approach alone is receiving interest in the UK by end-users who wish to use the data for forecast verification but this passive use of the sensors does not do justice to the potential of the approach. For example, it is a relatively straightforward step to enable sensors to directly feed their data into the models and decision support systems and therefore continually modify the forecast (and subsequently the actions of the winter maintenance team) based on the newly available high resolution data.

However, there is potential to change RWIS even further. The availability of detailed, and more importantly 'open', observation data along with the improved confidence which that brings, may empower decision makers to once again take a reactive approach to winter road maintenance. Currently, large sums of money are spent each year on road weather forecasts. If the IoT approach can be further extended to include a nowcasting solution (commercial products already exist which do this, e.g. the Vaisala nowcast controller), then there is the potential to re-adopt an ice detection strategy, thus making considerable savings on forecasts. As most nights in winter are clear-cut, an early decision on whether to salt or not can be made with a cheap and simple text forecast delivered at midday to the highway engineer. In the case of marginal nights, the information from the sensors could be closely watched, deploying the gritters only when absolutely necessary and thus making considerable savings. Effectively, this is what was done in the early years of RWIS. The difference now, is the resolution of observations in an IoT framework.

In summary, open data and the IoT are going to completely change winter maintenance decision making and operations. The traditional lines between instrumentation manufacturers and forecast providers are going to become increasingly blurred. There will be commercial winners and losers as a result of this, but we will be another step closer to making our roads a safer place on which to travel.

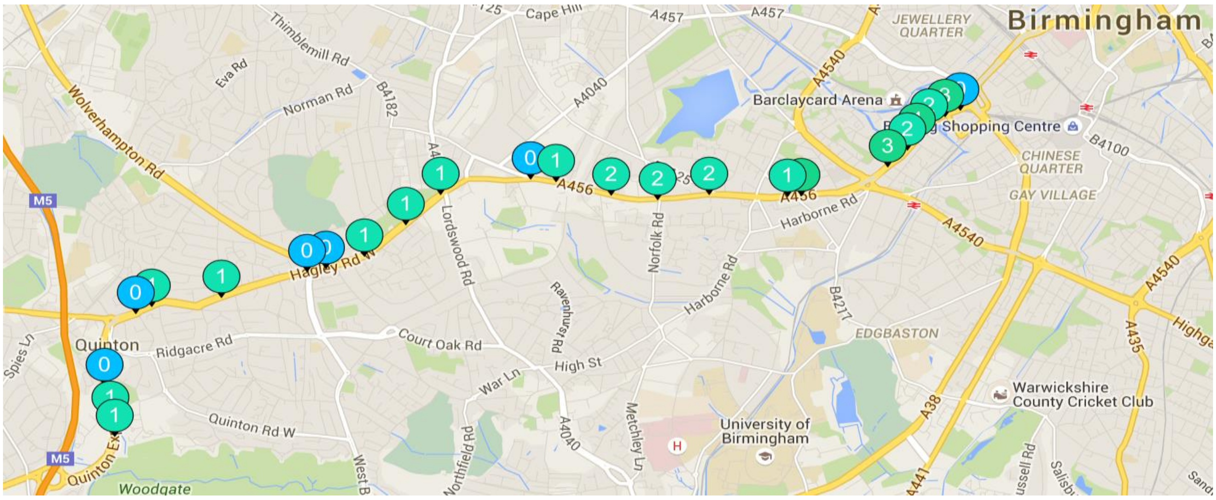

Figure 5: A screenshot of the web application used to visualise sensor data. Shown are real sensors mounted on lightning columns adjacent to Hagley Road, Birmingham. Map data ©2016 Google.

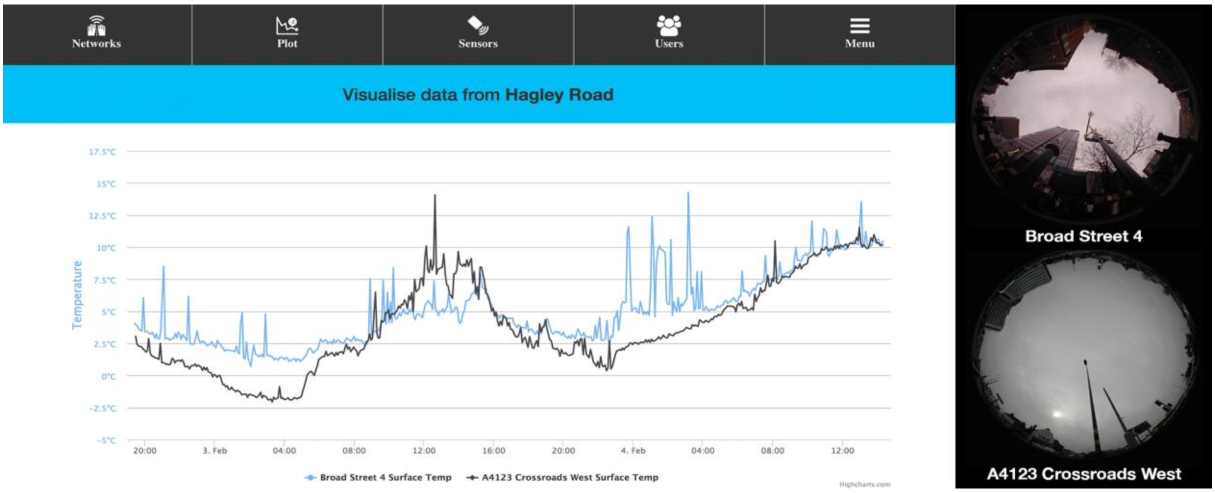

Figure 6: A screenshot of the web application being used to compare time series data from two of the sensors deployed along Hagley Road Birmingham. The sensors are ~3.5 km apart. The adjacent fisheye lens photographs illustrate the differences in sky view factor at the two locations.

References

Chapman, L. & Thornes, J.E. (2006) A geomatics based road surface temperature prediction model. *Science of the Total Environment* **360**:68-80

Chapman, L**.,** Young, D.T., Muller, C.L., Rose, P., Lucas, C., Walden, J. (2014) Winter Road Maintenance and the Internet of Things. *Proceedings of the 17th SIRWEC Conference, 28th-1st February 2014, La Massana, Andorra*

Gustavsson T (1999) Thermal mapping—a technique for road climatological studies. Meteorological *Applications* **6**:385–394

Hammond, D., Chapman, L., Thornes, J.E. & White, S.P. (2010) Verification of route-based winter road maintenance weather forecasts. *Theoretical and Applied Climatology* **100**:371-384

Handa, H., Chapman L. & Yao X. (2006) Robust route optimisation for gritting/salting trucks: A CERCIA experience. *IEEE Computational Intelligence Magazine* **1**:6-9