

# Wireless Sensor Networks for Aerospace Applications—Thermal Monitoring for a Gas Turbine Engine

D. Goldsmith\*, E. Gaura\*, J. Brusey\*, J. Shuttleworth\*, R. Hazelden\*\* and M. Langley\*\*\*

\* Cogent Computing Applied Research Centre, Coventry University CV1 5FB

\*\* TRW Conekt, Stratford Road, Solihull, B90 4GW

\*\*\* Vibro Meter UK, Viabes Industrial Estate, Basingstoke, Hampshire, RG22 4BS

## ABSTRACT

This paper reports on the development of a prototype wireless sensor network for thermal monitoring of aircraft gas turbine engines. The prototype acts as a concept demonstrator for the application at hand. Building upon the state of the art in the domain, the authors pursued a rapid prototyping approach, supported by a prototyping framework: FieldMAP. As a key property, the framework enables the conceptual shift from *data to real-time, user relevant information*, a feature that is less explored in the wireless sensor networks research. Consequently, an information extraction and visualisation component is put forward as an addition to traditional “sense and send” WSN systems. This component offers an intuitive approach to user understanding of the global evolution of the observed phenomena. Integrated within the prototype, the component makes use of the processing power available across the network coupled with interpolation algorithms borrowed from the geosciences domain. Reconstruction of field representations of the phenomena from the sparse sensed data allows identification of abnormalities and inference of their likely cause.

**Keywords:** wireless sensor network, engine monitoring

## 1 INTRODUCTION

Wireless Sensor Networks (WSNs) have developed in recent years from a small collection of simple experimental devices, to a diverse range of hardware platforms [1, 2]. Similarly the design space of WSN has grown from a set of dream applications to a large number of attempted real-life deployments over a diverse range of application areas [3]. Much of the initial work on deployable WSNs concentrated on their use and application as novel scientific instruments, enabling advances in science in a variety of fields [4–6]. Marked successes in WSN deployments, advances in standardisation of wireless communication technologies and MEMS sensors, and the forming of a wealthy, active community of WSN researchers, have expanded the applicability for WSNs into traditional industrial domains. Increasingly, WSN systems are being proposed and evaluated as potential solutions to measurement and automa-

tion in a variety of monitoring scenarios, even within safety-critical applications [7–11]. The key to WSN desirability lies with their potential of 1) providing, in real-time, spatio-temporal field representations of phenomena, with unprecedented detail and 2) enhanced deployability to conform with the constraints of the systems being monitored (low weight, wireless communication, ease of maintenance, and low cost).

For aero-engines, both properties above are equally important. For example, gas temperature within the engine exhaust, which is critical to the control and assessment of its overall health and performance, is measured by an array of thermocouples installed at a various radial and circumferential positions. Due to weight restrictions, only a single analogue signal is derived, effectively an averaging of the individual measurements, and this is transmitted to the engine control unit by a single heavy duty cable. This approach precludes the determination of a detailed picture of the engine gas temperature and its evolution with time, which could help in identifying potential engine problems early. It also prevents the diagnosis of individual sensor faults, sensor de-calibration, or sensor drift.

The use of wireless technology to create a wireless instrumentation system has the potential to substantially increase the amount of detailed data sent to the engine control unit and hence enable more sophisticated engine health monitoring. Further, replacing cabling with wireless transmission should reduce the monitoring system weight thus leading to improved fuel efficiency and reduced carbon emissions. On-line statistical analysis of data from such a wireless system could also permit a clearer understanding of engine/aircraft health, hence enabling the shift to Condition-Based Maintenance (CBM), whereby maintenance can be scheduled according to actual wear and usage rather than at fixed intervals.

In addition, a wireless system could allow for the sensors in the network to communicate their health metrics with each other, in turn allowing faults and drift to be identified and possibly corrected for in the engine control system. Greater confidence in the accuracy of the measured temperature could, potentially, allow the engine to be run with less safety margin and, therefore, more efficiently (with similar benefits on fuel consump-

tion and emissions).

However, embedding wireless technology into any gas turbine will have some significant challenges to overcome. The *safety-critical nature* of the system being monitored implies that a high degree of safety assurance and certification is required, hence fault management and reliability of any deployed instrumentation are core issues. The *deployment environment* is harsh and precludes the use of most conventional silicon-based electronic systems (for example, the temperatures outside the casing of the engine can reach in excess of 250°C). Furthermore, maintaining the integrity of an RF signal transmission in an environment that is largely composed of metal whilst avoiding interference from (or to) other electronic equipment is difficult. To enable *longevity*, correct autonomous estimation of residual life is needed. Retrofitting existing engines implies *flexibility* at instrument level to make use of existing opportunities for node positioning. Permanently fitted instrumentation precludes the use of batteries alone, hence some means of *energy harvesting* will be required, which needs to also comply with the small physical footprint required. Detailed monitoring may produce too much raw data to be readily understood and therefore *inference* and *summarisation* of the data is required. But such inference needs validation and verification. Also, inferred estimates of state may be uncertain, and this must be clearly represented.

Whilst work on the significant engineering challenges is only beginning, a proof of concept demonstrator has been developed to allow evaluation of the engine monitoring functionality, featuring five wireless nodes and 24 temperature and sound sensors. A back-end system for receiving, storing, analysing, and visualising the acquired data has been built and integrated with the sensing system. The end-to-end system has been deployed as a table top demonstrator and detailed testing and evaluation carried out towards identifying the challenges above and finding initial solutions.

The paper is structured as follows: Section 2 describes the proposed system design, architecture and implementation; section 3 reports on the system evaluation results and section 4 concludes the paper.

## 2 ARCHITECTURE

Fundamentally, the instrument development exercise reported here was governed by a need to: provide a rapid-prototyping set-up for the class of applications at hand; produce a flexible design which allows investigation and evaluation of implementation choices for each component and the system as a whole. The demonstrator produced is not meant as a deployment prototype; it is aimed to allow researchers to emulate and measure the system's response to a variety of conditions in the engine environment and allow plug-and-play, com-




Figure 1: Conceptual Flow for the prototype gas turbine monitoring system

ponent by component whilst maintaining the end-to-end functionality of the system.

Given that genericity and wide applicability of the demonstrator as a benchmark was a priority, the starting point in the work was the conceptual flow represented in figure 1. Most test-oriented and some permanently fitted instrumentation systems are generally designed to sit within this flow. Conventional instrumentation might integrate into a unique system all or some of the components in figure 1, with a strong domain drive towards both multi-point measurement and full integration and provision of a real-time information at user end.

For the engine monitoring application here, this conceptual flow applies as follows: temperature and sound is sensed (*sense*) at a number of locations within the gas turbine engine. Raw sensor data is noisy and in some cases, sensors may be faulty. For temperature, the *model* smooths the data, making use of assumptions about the likely evolution of temperature over time and other appropriate/application specific contextual information. Ideally, a conversion method would be provided here to enable the shift from sparse sensed data to continuous, field information (note that in most conventional systems this function is performed by the human observer applying their expertise and judgement). An interpolation model could be used here, for example, to derive a field function that fits the sensed data. Real-time visualisation (*visualise*) is enabled, allowing local and global temperature events to be identified as they occur. Furthermore, analysis of data over a time period (*analyse*) can also be performed off-line, after the event. These two information flows allow the human expert to either derive a *maintenance decision* (such as a component or sensor is faulty and must be replaced), or to devise a *test strategy* (such as modifying the engine structure to try to balance the heat distribution). The set of steps and actions can be seen as part of a control




Figure 2: Overview of the FieldMAP framework showing local node components above and base-station or back-end components below.

loop, feeding back changes to improve or maintain the engine.

With regard to the visualisation component, two similar works exist [12, 13]. Tricorder [12] is based around a hand-held mobile device that interfaces with Plug WSN systems [14] to allow easy access to sensed data. In comparison to the work here, Tricorder does not attempt to integrate data from neighbouring nodes to form an overall picture of the sensed phenomena. Rather data values are shown on a node-by-node basis. SpyGlass [13] is a visualisation tool for WSNs that supports a user-defined plug-in to paint a background on which the nodes are displayed. SpyGlass is aimed more at resolving the visualisation engine and less at supporting in-network processing.

Given the distributed, wireless and real-time nature of the instrumentation system proposed here, the sense-model-analyse-visualise flow described above can be further broken down into the schematic shown in figure 2, whose formal description translates to a generic wireless monitoring framework, FieldMAP, developed by the authors to aid the designer of WSN applications.

Whilst aiming to cater for a variety of monitoring applications, the FieldMAP framework [15] not only allows the core requirements of the system under design to be inferred but also shapes the goals and scopes the pro-

duction of the demonstrator. The framework supports the following node functions:

1. *Sense*. Given a set of sensors, form a vector of sample measurements corresponding to a time instant.
2. *Filter and manage faults*. Update a local state vector based on the sensed parameter vector taking account of which sensors are active. The state includes a management part to allow self-diagnostic capabilities.
3. *Detect events*. Based on a user-defined function, decide whether or not an event has occurred based on the current state vector and the last transmitted event. If an event is detected, append the current state vector to the event queue.
4. *Prioritisation*. While the channel is available and the event queue is non-empty, sort the queue according to some priority ordering scheme and transmit the first item.
5. *Schedule*. Based on the last attempted sampling time and the required sampling rate for each sensor, schedule the next sampling time for that sensor.

The base station supports interaction with the real-time information stream using the following generic approach:

1. *Store*. Store received vectors for replay, post-analysis and health management along with the associated transmission and receipt times.
2. *Update location state*. Predict the current state for all sensed locations based on a model. The model is the evolution of the state, based on the time for which the location state was last updated, and the previous prediction for the location
3. *Identify fresh locations*. Given a newly received state estimate, the age of the sample is updated based on the transmission time-stamp and the current time. Fresh locations are those whose ages are within some limit.
4. *Estimate field*. Given the current state estimate, estimate the field using a generalised function. This step results in a field approximation of the observed phenomena built from sparse point measurements.
5. *Colour map*. Produce a 2D visual representation of the field based on a false colour map.
6. *Recall / replay*. On request, data from a previous period can be recalled and optionally replayed, showing the evolution of the field over a period of time

7. *Feature / anomaly extraction.* Identify features or anomalies in the observed phenomena based on a set of rules.

The framework helps to structure the system in such a way as to make it relatively straightforward to adjust components such as the filtering or feature extraction. For example, initial benchmarks have shown the large gap between current power utilisation on the node and likely available power through energy harvesting systems currently available. Further, this study informed the need for a number of concurrent measures to be applied to any prototype for the application at hand, such as choice of microprocessor, radio and communications technology, application of energy aware design methods at both hardware and software level coupled with careful design of boundaries between the analogue and the digital system components.

An instantiation of the framework described above has been implemented and is reported here in the context of the motivating aerospace application. External phenomena are sensed in two modes: audio and heat. (The audio component is a place-holder for future investigation of vibration monitoring—sound sampling with microphones has some similar characteristics to accelerometer based vibration sampling in terms of format and data rate.) The sensors are attached to wireless processing nodes that communicate with the base-station. Mode dependent sub-components acquire temperature and sound, while a generic filtering system applies either a pass-through or Kalman filter. The base-station is responsible for storing data, displaying the field phenomena in real-time and interacting with the user.

Nodes were implemented using a Gumstix Verdex 400xm-bt processing board [16], with a custom-built expansion board to provide audio and I2C connectivity. The base station was implemented using a desktop PC which runs the visualisation software and interfaces to the back-end database. The Gumstix Verdex board includes an Intel XScale PXA255 400MHz processor, 16MB of flash memory, 64MB of RAM, a Bluetooth controller and antenna, and 60-pin and 92-pin connectors for expansion boards. The Gumstix boards run GNU/Linux and are powerful enough to perform the in-network processing required. Communication is provided by an on-board Bluetooth radio. The Bluetooth stack allows point to point communication using RFCOMM and L2CAP protocols with extended Bluetooth functionality such as the Service Discovery Protocol built-in.

Temperature sensing is provided via the Analog Devices ADT75A chip [17], which performs sampling and conversion internally before delivering temperature values via an I2C bus. The ADT75A provides temperature resolution of 1/16 degrees Celsius via a 12 bit ADC, and is rated for operation between -55°C and +125°C.




Figure 3: Visualisation of gas temperatures (left) and associated sample point time-series in terms of temperature (above right) and rate of change of temperature (below right).

On the sensor node, temperatures are gathered using an in-house written I2C library to query each temperature sensor in turn. The sensing library detects attached sensors at start up. The audio sensing module uses the on-board AC'97 processor and periodically extracts RMS levels of the audio signal, providing an indication of the sound level detected. Time stamps are appended to the data at the time of collection.

Upon collection, Kalman filtering is performed on the data, based on the assumption of a constant rate of change of temperatures. An advantage of this filtering approach is that it provides smoothed estimates both of the sensed temperature and its rate of change.

On the base station frames are received and processed into individual samples. Each frame is a recording of the environment at a given instant and can contain multiple samples (two volume measurements, for example, for the audio sensing mode). The time stamp of the original transmission is used to identify stale data. The received samples are logged in an external database, and any required data transformations (such as calibration offsets) are applied.

To generate the field visualisation of the two phenomena, three user selectable interpolation components have been developed and integrated into the end-to-end system (Inverse Distance Weighting (IDW) [18], QSHEP2D [19, 20] and Nearest Neighbour, the last simply giving the same value to a point as that of its nearest neighbouring sensor). The values associated with the phenomena are displayed graphically using a false colour map (see figure 3).

The user interface displays the interpolation results and history of a selected sensing point. It performs the following functions: periodically (or as often as possible) requests gathering of new samples; initiates the logging of data to a database; produces the interpolation of current data; displays interpolation results and historical

time series for user selected nodes; displays the results of the fault management component by marking on the display the faulty sensors and discarding their data; allows configuration, including: whether to display sound or temperature; selection of sensors for historical data display; interpolation block size; desired frame rate.

The demonstrator, with four fixed nodes and one probe node was deployed on a cold jet pipe, with the sensors mounted in a radial pattern within the pipe in a similar pattern to existing thermocouple harnesses, sampling at 4Hz both for temperature and audio. The temperature sensors were wired in sets of four to the five nodes and the resulting visualisation of the temperature maps over the monitored surface area, field profiles of the temperature rate of change and the history of individual sensor values were observed in a variety of controlled experiments for evaluation purposes.

### 3 EVALUATION

The demonstrator was evaluated both against the application specific goals (i.e. its potential to offer a replacement solution for the current harnessed thermocouple systems in use) and in terms of its end-to-end performance from a networking viewpoint. The user end (information extraction and visualisation component) was evaluated in terms of the accuracy of resulting interpolations and its usefulness for engine technical experts. The system has proven to be sufficiently robust to allow continuous monitoring without data loss. A detailed assessment of power requirements for continuous monitoring using this system has been performed and is reported elsewhere.

#### 3.1 *Networks and end-to-end system performance*

The Bluetooth based wireless network performance has been analysed in terms of latency, data throughput, bandwidth and communications range. The maximum transmission range has been found to be consistent with the Bluetooth specification. Under normal operation, the sample rate used within the system requires an average of 43 kbps throughput. The system has a mean end-to-end transmission time of 0.27 seconds. The maximum latency experienced by the system is around 8 seconds but this is only during the initial system start up when discovery of attached sensors delays acquisition somewhat.

#### 3.2 *Interpolation and visualisation*

While the data generated by the sensors provides known points in the interpolation, it is important to be able to (i) assess the effect sensor failures on the field representation of the phenomena and (ii) to have

a means to measure the quality of values predicted between these points leading to the field representation of the sensed phenomena.

To this end, leave-one-out cross-validation was used to find the error in the field estimate. In this scheme, one sensor point is removed when forming the field estimate. The estimated value for the left-out sensor can then be compared with the sensed value for that point. Experimentation with a heating-cooling cycle was used to provide data for cross validation. The IDW and Nearest Neighbour interpolation algorithms provided similar results with a MSE of 2.07 and 2.24 ( $^{\circ}\text{C}^2$ ), respectively. QSHEP2D gave the highest MSE of 4.35.

The above cross-validation approach may be subject to the problem that the regular arrangement of the sensor nodes causes some feature of the phenomena to be missed. Therefore a further test was used by including an additional “probe” sensor to compare the sensed and estimated values at arbitrary points in the field. Combined with the cross-validation tool, the added node provides a method to determine the accuracy of predicted temperature values between nodes. The results from using a probe sensor mirror those found during normal operation with IDW (0.57) and Nearest Neighbour (0.83) out-performing QSHEP2D (1.00).

### 4 CONCLUSION

The paper reported on the development of a wireless instrumentation system demonstrator dedicated to flow temperature measurement in aircraft gas turbine engines. The development work was guided by a more generic wireless monitoring framework, FieldMAP, devised by the authors, which enables the delivery of user relevant information rather than data as the outcome of the monitoring process and enables event based decisions and diagnostic. A rapid prototyping approach was taken in the development of the demonstrator. The demonstrator’s main role is to allow benchmarking of the various components and needs of a future end-to-end monitoring system suitable for deployment in the engine environment—sensors, communications stack, power needs, processing needs and finally information representation. The modular structure imposed by the FieldMAP framework enables rapid replacement of the demonstrator system functional and hardware components whilst ensuring that the end-to-end monitoring system flow is maintained. Future work aims at accommodating appropriate deployment environment constraints as needed within the demonstrator components, starting with the replacement of the current temperature sensors with harsh environment thermocouples and replacement of the wireless sensor network nodes with low power devices.

## ACKNOWLEDGEMENTS

This project is funded as a three and a half years EP-SRC CASE Studentship through the Integrated Products Manufacturing KTN, UK, industrial sponsorship from Vibro-Meter UK. Additional expertise is provided by TRW Conekt, UK.

## REFERENCES

- [1] Daniel Fokum, Victor Frost, and Minden Gary. An evaluation of sensing platforms used for sensor network reseach. Technical Report ITTC-FY2008-TR-41420-04, University of Kansas, Information and Telecommunication Technology Centre, December 2007.
- [2] Jason Hill, Mike Horton, Ralph Kling, and Lakshman Krishnamurthy. The platforms enabling wireless sensor networks. *Commun. ACM*, 47(6):41–46, 2004.
- [3] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. A survey on sensor networks, 2002.
- [4] Philo Juang, Hidekazu Oki, Yong Wang, Margaret Martonosi, Li S. Peh, and Daniel Rubenstein. Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with zebnet. *SIGOPS Oper. Syst. Rev.*, 36(5):96–107, 2002.
- [5] M. Allen, L. Girod, R. Newton, S. Madden, D. T. Blumstein, and D. Estrin. Voxnet: An interactive, rapidly-deployable acoustic monitoring platform. In *Information Processing in Sensor Networks, 2008. IPSN '08. International Conference on*, pages 371–382, 2008.
- [6] K. Martinez, P. Padhy, A. Riddoch, HLR Ong, and JK Hart. Glacial Environment Monitoring using Sensor Networks. In *Real-World Wireless Sensor Networks*, 2005.
- [7] J. Kemp, J. Brusey, and E. I. Gaura. Instrumenting bomb disposal suits with wireless sensor networks. In *Proc. 5th Intl. Conf. on Informatics in Control, Automation and Robotics (ICINCO)*, 2008.
- [8] J. Kemp, E. I. Gaura, J. Brusey, and C. D. Thake. Using body sensor networks for increased safety in bomb disposal missions. In *Proc. IEEE Intl Conf. on Sensor Networks, Ubiquitous, and Trustworthy Computing (SUTC 2008)*, pages 81–89, Washington, DC, USA, 2008. IEEE Computer Society.
- [9] Link C. Jaw. Recent advancements in aircraft engine health management (EHM) technologies and recommendations for the next step. In *Proc. ASME Turbo Expo 2005*, number GT2005-68625, Reno-Tahoe, Nevada, June 2005.
- [10] Philippe Dubois, Cyril Botteron, Valentin Mitev, Carlo Menon, Pierre-André Farine, Paolo Dainesi, Adrian Ionescu, and Herbert Shea. Ad hoc wireless sensor networks for exploration of solar-system bodies. *Acta Astronautica*, 64(5-6):626–643, March 2009.
- [11] J. Riihijärvi L. Hanna, K. Koumpis and P. Mähönen. Reconfigurable, ubiquitous, networked embedded systems and their use in emergency response applications. In *Proc. 5th European Conference on Mobile Solutions, Mobile Europe*, September 2006.
- [12] Joshua Lifton, Manas Mittal, Michael Lapinski, and Joseph A. Paradiso. Tricorder: A mobile sensor network browser. In *Proceedings of the ACM CHI 2007 Conference - Mobile Spatial Interaction Workshop*, April 2007.
- [13] Carsten Buschmann, Dennis Pfisterer, Stefan Fischer, Sándor P. Fekete, and Alexander Krölller. Spyglass: a wireless sensor network visualizer. *SIGBED Rev.*, 2(1):1–6, January 2005.
- [14] Joshua Lifton, Mark Feldmeier, Yasuhiro Ono, Cameron Lewis, and Joseph A. Paradiso. A Platform for Ubiquitous Sensor Deployment in Occupational and Domestic Environments. In *Proceedings of the Sixth International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 119–127, April 2007.
- [15] James Brusey, Elena Gaura, Daniel Goldsmith, and James Shuttleworth. FieldMAP: A spatio-temporal field monitoring application prototyping framework. *IEEE Sensors*, May 2009. (to appear).
- [16] Gumstix Inc. Gumstix.org. Available online <http://www.gumstix.org/>; Last Accessed 11/08.
- [17] Analog Devices. <http://www.analog.com/en/prod/0> accessed 25th Feb 2008.
- [18] Donald Shepard. A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 1968 23rd ACM national conference*, pages 517–524, New York, NY, USA, 1968. ACM Press.
- [19] Robert J. Renka. Algorithm 660: QSHEP2D: Quadratic shepard method for bivariate interpolation of scattered data. *ACM Trans. Math. Softw.*, 14(2):149–150, 1988.
- [20] Robert J. Renka. Multivariate interpolation of large sets of scattered data. *ACM Trans. Math. Softw.*, 14(2):139–148, 1988.