A Distributed Instrumentation System for the Acquisition of Rich, Multi-Dimensional Datasets from Railway Vehicles

Edward James Charles Stewart

A thesis submitted to
The University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

School of Electronic, Electrical and Computer Engineering
College of Engineering and Physical Sciences
The University of Birmingham
9th August 2012

ABSTRACT

This thesis presents work carried out over a number of years within the field of railway vehicle instrumentation. The railway industry is currently moving to be more heavily "data driven". This means that railway organisations are putting policies into place whereby decisions have to be justified based on recorded and citable data. To achieve this, the railway industry is increasingly turning to greater and greater levels of instrumentation to deliver the data on which to base these decisions.

This thesis considers not only this increased requirement for data, but the frameworks and systems that must be put into place in order first to obtain it, and then to extract useful information from it. In particular the author considers the issue of contextualisation of data, where multiple datastreams may be used to provide context for, or allow more accurate and beneficial interpretation of each other in order to support better decision making. In order to obtain this data, the thesis explores, through a series of case studies, a number of options for different instrumentation system architectures. This culminates in the development of a distributed system of embedded processors arranged in an extensible modular framework to provide a rich, coherent and integrated dataset which can then be processed contextually to yield a better understanding of the railway system.

The first case study relates to energy metering and demonstrates that greater data rates allow

additional information such as driving style to be extracted from basic measurements such as energy consumption. The second case study operates in the inhospitable environment of the conductor shoe / third rail interface and uses one measurement technology to validate the results of another. The third considers a combination of multiple sensing technologies to improve explanations of track alignment results. The final case study details an extremely heavily instrumented train and demonstrates a proposed instrumentation framework capable of autonomous, in-service, recording and integration of a vastly expanded sensor set. Analysis of this richer, more closely integrated dataset demonstrates an ability to contextualise the results from one data stream using another, thus improving understanding of a signal of interest.

The thesis represents a combination of both methodological work relating to the implementation of instrumentation systems, and application based activities in terms of using the data obtained. In this capacity, the work presented in the case studies has directly generated one journal paper and contributed to another.

ACKNOWLEDGEMENTS

I would like to thank my boss and supervisor, Professor Clive Roberts, for supporting me through this PhD and allowing me to (attempt to) balance the demands of a life in research with the demands of everyone else telling me to hurry up and finish my thesis.

I would like to thank Dr Steven Quigley for taking me on in the first place; for lending me to Clive for "6 months"; and for understanding when I didn't come back.

I would like to thank Mary Winkles for keeping me on the books, and the multiple heads of school that have come and gone for supporting each and every one of the numerous extensions that I've required.

I would like to thank Dr Paul Weston with whom I've worked on many projects, including several described in this thesis. His technical brilliance and scathing humour are appreciated in equal measure and more than he realises.

I would like to thank Dr John Easton for being a voice of reason; for his willingness to "muck in" when deadlines loom; and for his eternal willingness to go to the pub!

Finally, I would like to thank Dr Suzanne Gildert and my family for their unwavering support and believing in me even when I couldn't see the end.

CONTENTS

Glossary

Bibliography

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

This thesis is about making intelligent decisions based on the data that exists, or will exist within the railway environment. There is a vast, and increasingly growing, repository of railway data becoming available but at this time there is comparably little processing applied to that data in order to support these intelligent decisions. This thesis considers the question of contextualising data such that analysis of particular data streams can be performed more intelligently and thus better decisions can be made as a result. The thesis identifies the need to integrate data streams, and highlights the benefits of doing this at the point of acquisition. An instrumentation framework, capable of obtaining and integrating data from a distributed array of disparate sources, is presented and a sample application described. Only through this closer integration of the data that is being obtained, will full contextualisation of the data be possible. Without this contextualisation analysis of data is, at best, limited to the localised performance of the asset being monitored. By being able to observe and understand larger sections or components of the railway industry, larger scale or "bigger picture" decisions can be made with impacts in efficiency and safety as well as performance.

1.1 The Need for Instrumentation Within the Railway

The key to operating an efficient railway network is information. Information is used in almost every aspect of the railway system from the number of passengers dictating the required frequency of services to the standard measurements by which a set of points is put into alignment.

The key to information is data. Information is the result of processing and interpretation applied to data. Before one can process data, however, one must collect it. Data acquisition is the process of recording raw data from the environment; in this case, the railway environment. To acquire data from an environment, to measure it, one must use instruments. Equipment used to measure the physical properties of an environment or system is known as instrumentation.

Simply obtaining data and then processing it to extract information does not, however, lead directly to an efficient railway. The infrastructure manager or train operator must use the information to inform decisions and then act upon those decisions. Generally speaking, the better informed the decisions are, the better the outcomes of said decisions can be. This then becomes a question of the volume of accessible and relevant information available.

1.2 The Complex Railway System

The railway, as a system, is fundamentally complex. This means that there are a huge number of independent subsystems which must all work together in order to deliver efficient operation. The railway system is divided on almost every level from infrastructure vs. rolling stock, to the interaction of international railway networks, right down to the rivalry between permanent way and electrification work crews and the fact that plain line and switch and crossing welding must be done by different people.

This level of separation within a single industry can clearly lead to both practical as well as

idealistic disconnects between different railway systems and the subsystems that they are comprised of. In an attempt to overcome this, the industry has developed standards. Standards describe interfaces between different systems, setting up rules that the two sides of the interface must adhere to in order to ensure compatibility. Standards can describe terminology, thereby ensuring that different railway companies are able to communicate without confusion. They can also be technical, relating to mechanical or electrical interfaces. A guidance document from the Office of Rail Regulation (ORR) [1] describes how a number of European directives intended to increase compatibility between different national railway systems are now enshrined in UK law. While this may support higher level system integration, there are a lot of lower level systems which still function in a relative state of isolation.

1.3   Information Management

Over the last 10 years or so, instrumentation systems have been becoming both more capable and more accessible to the point that instrumentation is now quite prevalent in the UK railway industry. This increased willingness to instrument both rolling stock and infrastructure, combined with the rise of cheap and reliable storage technologies has lead to huge volumes of data being retained by a whole range of organisations.

More recently companies such as Network Rail have realised that they are accruing huge volumes of data and that they need to put strategies in place to manage the information within that data in order to generate the promised business case benefits from performing the instrumentation.

Part of the problem that companies such as Network Rail are facing is that, as large multi-faceted organisations, their internal subsystems do not communicate sufficiently well what information they are obtaining and thus what is available to the other subsystems. Information exists, but it is segregated in repositories which not everyone is even aware of, let alone has access to.

Overcoming this segregation has become increasingly important in recent years, as exemplified by the proliferation of management terminology such as "breaking out of the silo".

One approach that Network Rail have taken to reducing this segregation is to develop a standard data model. In developing this model, Network Rail have inventoried their existing data repositories to establish a definitive index of what they know, and allow themselves to identify what decisions this data can be used to inform. The data model describes how the different subsystems (or teams) within the company collect and manage data, how they use the data, and how the data could be fed into other subsystems to support their decisions as well. The development of the Network Rail standard data model suggests that, as a company, Network Rail are now starting to think about data and information management and making more intelligent use of the informational resources that they have available. It is significant, however, that this approach is based on making use of existing resource and does not necessarily dictate how new data acquisition exercises are managed.

1.4 Data Integration

On a more immediate and practical level, the increased availability of computational resource and complex instrumentation systems has lead to a situation in which far more advanced data gathering exercises can be undertaken. Increased instrumentation complexity allows greater numbers of signals to be recorded at higher data rates at a single instrumentation site. More recently, advances in communication systems have allowed the development of distributed instrumentation systems in which multiple, otherwise independent, instrumentation sites can synchronise and pool the data that they obtain - essentially acting as a much larger instrumentation network. Such systems rely on synchronisation and data interchange protocols as well as increased communication bandwidths.

Taking this concept one stage further. As long as the synchronisation is sufficiently robust, it is possible to record multiple datasets independently and then align them in post processing. Generally speaking, the systems used for this need to have some common interconnect or signal recorded in each dataset which can then be used as a key to align them. Depending on the required level of integration, this might be a a high frequency timing pulse, a regular clock signal, or just a common control signal present in all of the systems. One approach used at the University of Birmingham is to fit GPS receivers to infrastructure assets, not to identify their position, but to record the satellite time signal at multiple locations. The common time signal can then be used as an alignment key.

A significant amount of the instrumentation found in the railway is used to record data about fixed assets (infrastructure). Also, it is common to interpret data recorded on vehicles in terms of the location of that vehicle when the data was recorded. If data is to be interpreted in space rather than time it may be possible to spatially tag the data and then align it based on the tags at a later date. This requires sufficient accuracy in the spatial tagging mechanisms, but these technologies have also advanced greatly in recent years. On a raw data level, this approach is similar to what the Network Rail model aims to achieve with information.

In a similar way to the railway industry's development of standards to define the interfaces and interactions between it's various subsystems, standards can also be used to define interfaces at instrumentation outputs in order to simplify post processing alignment. On some levels these frameworks may look like data exchange formats, in other cases they are more fundamentally the basic data output and transmission mechanisms.

1.5 Instrumentation Frameworks

Throughout this thesis, the author describes a flow of information: from fundamental measurements of the physical world, through signals and transducers, to data, through processing and analysis to information and finally back to an understanding of the physical world on which decisions can be based. To achieve this, a combination of instrumentation systems and example analysis processes are used.

This thesis will demonstrate that modern electronic and computational resources allow the development of ever increasingly powerful instrumentation systems. By basing these systems on modular frameworks that support distribution of transducers, but also integrated data repositories, it is possible to develop extensive, rich, datasets that incorporate a number of previously disparate data sources using a dynamic range of different transducers. An example distributed instrumentation system based on such a framework will be specified and developed. Analysis of the datasets obtained from the example system will then be shown to lead to far greater understanding of particular signals than analysis of those signals alone, thus providing the ability to make more informed decisions based on the recorded data.

The key hypothesis to be tested is therefore: Computational resource has now progressed to a point that much more advanced instrumentation frameworks can be developed to operate in the harsh environments of, for example, railway vehicle under-frames. These frameworks allow for the development of much more complex instrumentation systems and the integration and synchronisation of previously disparate data sources recorded from different locations across the vehicle. Analysis based on these integrated datasets is more powerful and can yield more information than analysis of the signals individually. Thus, by using instrumentation frameworks that support this more complex data collection regime it is possible to make better, more informed, decisions about the operation of the railway.

1.6   Thesis Structure and Organisation

In this thesis, the author will consider the use of modern instrumentation technologies in developing complex bespoke instrumentation systems capable of supporting "bigger picture" objectives and data / information integration strategies for the railway industry as a whole. The thesis will use a number of case studies to demonstrate how standardisation of outputs can facilitate post processing and allow a dataset to be interrogated to answer research questions outside of those envisaged when the instrumentation was being developed. Consideration will then be given to the integration of a number of data sources to generate a greater wealth of information than processing the datasets individually. This will be done both through the combination of separate datasets, and also through more closely integrated systems. In particular, chapter six describes an instrumentation architecture in which the instrumentation is distributed over an entire railway vehicle but centrally recorded. The processing undertaken on this, much richer, dataset demonstrates how an increased quantity of instrumentation allows advanced interpretation of data from one type of sensor through information gained from another, and how this can be used to verify assumptions and hypotheses as well as answer more complex research questions.

The thesis is structured into four individual case studies. Through these, a range of different approaches to the instrumentation of railway vehicles will be presented. Each case study will include a technical discussion of the development of the instrumentation, along with analysis of the data and a selection of results from the associated study. From these discussions, some element unique to the instrumentation, and its associated benefits, will be identified.

The work described in the first case study relates to the instrumentation of the electrical systems of a DC electric railway vehicle. The work was undertaken by the author with vehicle access provided by a train operating company and funding supplied by the Department for Transport. The specified outputs of the study were to investigate strategies for energy metering equipment

in light of legislation to mandate energy metering on all railway vehicles, and to investigate potential energy savings through improved driving styles. In addition, the case study shows an example of a minimalist instrumentation system and highlights some of the limitations that this creates in the information extraction and data analysis processes. The outputs of the work have been presented in numerous industrial working groups, and an international conference and have led to further projects including the instrumentation of parts of the infrastructure of the DC railway network.

The work described in the second case study was commissioned by the Railway Safety and Standards Board and Network Rail, with vehicle access provided by a train operating company. Instrumentation was fitted to the current collection systems of a third rail powered DC electric vehicle. The specified outputs were the quantification of the behaviour of the shoegear and the forces experienced at the interface between the conductor shoe and the third rail. The trials included repeated passes over a single route section at multiple specified speeds to quantify the consistency of the performance of the shoegear equipment. This case study is included to show the evolution from a minimal dataset to one only slightly more advanced, and how much benefit this can bring to the analysis processes. In addition to presentations to the industrial sponsors, this work has been presented to the manufacturers of the shoegear, at national and international conferences and in the journal of the Institution of Mechanical Engineers. A subsequent research project into using the dynamics of the conductor shoe to identify defects in third rail alignment is a direct result of this work.

The third case study describes an extension to the conductor rail measurement work. The instrumentation developed in the study involved inertial measurement sensors sufficient for the monitoring of the bogie of a railway vehicle in six degrees of freedom. Analysis of the data from these sensors has shown that it is possible to consider the condition of the track from data

collected on a railway vehicle. The study shows a progression in the instrumentation method-ology from previous work carried out within the research group, and demonstrates improved results despite utilising a reduced selection of transducers. The case study shows that even with more advanced instrumentation systems, the field trial arrangements are critical and the total volume of data collected is hugely important. The work has been presented in the journal of the Institution of Mechanical Engineers as part of a larger project into using vehicle mounted sensors to monitor the condition of railway track. In addition, individual track features relating to transition zones have been identified and will form the core of part of a work package in a substantial EPSRC funded project. Discussions are also ongoing with one of the train operating companies to install track condition monitoring equipment based on this technology. This case study provides a good example of information being extracted from a dataset not originally recorded for that purpose.

The final case study re-visits the work from case study 1. The results presented in the first case study were used to ensure a second phase of trials with an expanded instrumentation set and vehicle routing options. The case study explains how the system has been re-specified and re-designed to include instrumentation of many more aspects of the vehicle including inertial systems, driver systems and the traction control system, and was fitted for an extended period to an in-service vehicle. This allowed a facility for continuous data collection and thus vastly increased data volumes. To support this, the instrumentation framework had to be signific-antly re-developed to be much more complex, and to include local storage media and far more comprehensive data analysis techniques bringing in the different types of signals recorded. The case study describes the specification and development of such an instrumentation framework, and shows that while smaller systems may answer specific questions, the integration of multiple (often disparate) measurements is vital to truly understand a system as a whole. With the rich dataset obtained from the instrumentation described in this, fourth, case study it has been

possible to answer many more questions about the operation of the vehicle and the system of which it is a part.

The thesis concludes with a discussion of the increasing complexity of the different instrumentation systems and frameworks presented, the corresponding increase in the volumes of data recorded, and thus the increased levels of information that can be extracted. Strategies for future instrumentation frameworks and how to obtain the most from existing datasets are considered, along with how the lessons learned throughout the projects described are being realised in ongoing and future instrumentation based research.

In addition to the core discussions about the benefits of increasing the complexity of instrumentation systems, and the instrumentation frameworks that are required to allow analysis and interpretation of data in relation to other disparate data streams. The work presented in this thesis shows that through instrumentation systems, and instrumentation systems of increased complexity, it is possible to obtain a greater understanding of the operation of the railway. This understanding can then be used to improve the capability and the reliability of the system. Specifically, the results of the case studies show:

- That through the use of multiple, integrated, data streams it is possible to obtain a greater understanding of the operation of a railway vehicle and thus be in a better position to interpret individual signals recorded during its operation.

- That the balance between performance and the energy use within a vehicle, particularly the traction system of that vehicle, can be more closely understood.

- The benefits that can be obtained from increasing data rates beyond the legislated minimum.

- That an instrumented vehicle can be used as an instrument of measurement for the infra-

structure in both the areas of the power supply network and track alignment.

- That track alignment can be monitored independently of vehicle speed, and using a reduced instrumentation set to that previously reported and shown in the literature.

- That feature identification can be used to extract data about certain locations.

- That the configuration of a vehicle can have an impact on optimum driving style and vehicle performance.

- That, while certain driving styles are more efficient than others, the environment has one of the most significant effects on energy efficiency.

1.7   Industrial Partners

The results to be presented have been obtained from several different research projects, each with different industrial partners. Some of the partners and their contributions to the projects are listed here:

- Angel Trains. One of the UK's Railway Operating Support COmpanies (ROSCO). Angel trains own the Class 508 used in the fourth case study. Their arrangement with the Train Operating Company, and the level of modification, meant that their approvals were required.

- Department For Transport. UK government division associated with the rail industry. Funded and specified many of the outputs of the first and fourth case studies.

- Network Rail. The UK's main line infrastructure owner and manager. Permitted track access for modified vehicles in all case studies. Also partially financed and specified the outputs of the second case study.

- Merseyrail. Train Operating Company for the Liverpool metro system. Provided access to the Class 508 used in the first and fourth case studies.

- Railway Safety and Standards Board (RSSB). Responsible for railway standards within the UK. Involved in specification of the outputs of the second case study.

- London and Southeastern Railway. Train Operating Company for the Southeast region. Provided access to the Class 375 used in the second and third case studies.

CHAPTER 2

REVIEW OF INSTRUMENTATION TECHNOLOGIES AND THEIR APPLICATION

This chapter considers some of the previous work that has been undertaken in the field of railway systems to obtain data, and then to process that data in order to generate information with which to make decisions. This thesis is focused on better understanding data, and the information that it yields, in order to make informed decisions based on contextually explained and relevant data. Part of this process is concerned with data analysis, but a large part of this is based around obtaining and preparing the data and databases in advance of any processing. A significant portion of the work described in this chapter is related to the instrumentation systems that are used to obtain data from the physical world. The discussion that follows will consider different approaches to this, and will highlight some of the key differences between robust, long term, minimally complex systems and those larger systems more actively targeting open research questions.

2.1 Introduction

Instrumentation can be defined as the design, provision or use of measuring instruments. Virtually any physical parameter can be measured. The railway industry uses instrumentation in almost every aspect of vehicle and infrastructure development and maintenance. Some instrumentation is used to assist with development; often, existing systems are instrumented to provide information which is then used to improve future system designs. Other instrumentation is fitted in order to answer specific questions thus helping with troubleshooting of particular issues. A third kind of instrumentation is associated with the field of condition monitoring. In this case, systems are instrumented to provide data on their operational status, which can then be used to direct maintenance. If implemented correctly, this system can be used to identify, and fix, developing faults before they become failures.

In the context of railway vehicles, instrumentation extends both to the vehicle itself and also to the surroundings, with measurements being taken from the vehicle. Conversely it is possible to place instrumentation on the track, such as Hot Axle Box Detectors (HABD) or Wheel Impact Load Measurement Systems (WILMS), to reveal information about the vehicle.

Instrumentation is more than the application of sensors [2]. It involves management of the data that the sensors present. This management can extend to localised processing with the instrumentation outputting alerts or summaries to the user or, in a more fundamental case, recording all of the data for future analysis. To obtain the most information from any given data set, the data needs to be contextualised. This thesis considers contextualisation to be the integration of multiple data streams such that one can be interpreted more intelligently using information from another. Another key data management process involves pre-processing data to ensure compatibility with other data or a data processing technique. For example, most instrumentation systems involve the acquisition of sensor data in the time domain; in some

cases, for this data to be useful, conversion into a spatial domain may be required.

In many cases, the intended data processing may define the complexity of the instrumentation system to be developed. Instrumentation for development often involves full data acquisition so that the data can be repeatedly analysed as new questions are asked. Instrumentation for troubleshooting purposes tends to be concerned with extremes in the measurements or with specific events and as such much of the data is discarded. Condition monitoring instrumentation tends to perform real time processing of the data, discarding the raw signals immediately. Clearly this approach to the handling of the data can have dramatic impacts for the form and architecture of the instrumentation system.

Technology Readiness Levels (TRL) [3] are a method for describing how far along a development path a given technology is. The levels are numbered 1 to 9 where one is extreme blue sky research and nine is a finished and tested product level implementation of a technology. As a general rule, academia operates in the lower 4 or 5 technology readiness levels, while industry tends to occupy levels 5 upwards. The higher up the TRL scale an instrumentation system is operating the more robust and ruggedised it will need to be. High TRL systems that operate for extended periods, or that may become permanent installations, need to be particularly robust and are therefore usually simplified in terms of extent and capacity when compared to systems operating further down the scale. Lower TRL systems tend to be less robust but more complex. Is is here particularly that the the integration of multiple datastreams is essential to maximising the potential outputs of these systems.

The remainder of this chapter will consider some of the components used in instrumentation and also look at how these component parts can be combined together to generate different styles of instrumentation system depending on the application or the intended further processing to be carried out.

2.2 Instrumentation Technologies

An instrumentation system may consist of many different types of technology. At the physical extreme, there must be a component of the system capable of detecting a change in the thing being measured. The system may need to process the information that a change has occurred either to interpret further information from it, or to quantify the change. Once information has been gleaned, this information must be recorded somehow before ultimately being processed to present the fact that the change has occurred in a usable form.

The Cambridge dictionary defines the word sensor as a device which is used to record that something is present or that there are changes in something. Sensors can be applied to detect changes in almost any physical system. They can also, therefore, be used to detect and quantify consistency in said systems. The work described in this document makes extensive use of multiple types of sensors during the development of a number of instrumentation systems. Of particular interest are sensors which can be used to monitor electrical systems, sensors which monitor movement, and to a lesser extent temperature.

Although the word sensor has been used, modern instrumentation systems use standard interfaces to record their inputs. The job of the sensor is therefore to convert signals from their native format to one compatible with the recording components of the system. This family of devices are known as transducers.

One of the most basic kinds of transducers is the voltage transducer. Modern data recording equipment generally requires its signals to be in the form of varying voltage signals over a limited range. Common ranges are bipolar ± 10 V and ± 5 V or unipolar 0-10 V and 0-5 V. Much more sensitive units also exist, for example thermocouple recording equipment with inputs designed for ± 80 mV. Clearly, measuring a large voltage (e.g. kilovolts) with a system calibrated for

10 V inputs is not possible so a transducer must be used to reduce the voltage. In its most basic form, a voltage transducer could be a precision voltage divider circuit; however, voltage transducer technology is far more advanced than this. When measuring high voltages, one of the key factors in transducer selection is the isolation rating. This is the voltage that can be applied to the inputs without causing the voltage at the output of the unit to exceed its specifications.

Voltage isolation is not always required for electrical safety in the form of high voltages. Sometimes the signal being monitored is sensitive to interference or being modified by the instrumentation system. In these cases, isolation amplifiers can be used to electrically separate the signals and the instrumentation systems. In some cases, these isolation systems may be in the form of isolation amplifiers, although often the gain of such amplifiers is 1:1. Isolation amplifiers commonly consist of an equally balanced transformer where the signal is relayed electromagnetically rather than electrically. Alternative forms of isolation include optical isolation whereby the signal being monitored is used to illuminate an infrared LED which then, through a photo sensor, transfers the signal across the isolation boundary. Optical isolation is inherently Boolean in nature and so can only be used in digital systems.

Another electrical property that is commonly monitored is current. In its simplest form, current flow can be measured by monitoring the voltage across a known, precision, resistor and performing a calculation based on Ohm's law. More complex current transducers tend to operate on the principle that where current flows, magnetic fields are developed and in turn can develop voltages. This was discovered by Edwin Hall and is therefore known as the Hall effect. Hall effect current transducers fall into two categories. Open loop and closed loop.

Open loop transducers often come in the form of split core transducers which can be clipped around cables to measure the current flowing therein. They use the basic principles of electromagnetism to develop a voltage proportional to current flow and commonly amplify this voltage

to present it as an output in a standard range. Open loop transducers commonly have voltage outputs and can be unipolar or bipolar with the maximum voltage corresponding to a nominal maximum current value and the minimum either to zero or an equivalent current flow in the opposite direction.

Closed loop transducers are usually more accurate than their open loop counterparts. Commonly they are solid core units which require a cable to be passed through them rather than being capable of being clipped around it. Closed loop transducers draw current from a power source in order to develop a magnetic field that counters the one being developed by the current being measured. In keeping these fields balanced, a current proportional to the one being measured is exhibited as an output.

In addition to electrical monitoring, the instrumentation described in this document is also concerned with the measurement of movement, vibration and rotation. For these measurements, a number of different types of transducer are used.

Displacement transducers are used to measure linear displacements between one point and another. They broadly fall into two categories: contact and non-contact. Contact based displacement transducers rely on attaching part of their structure to either end of the thing being measured. As the two objects become closer together, or further apart, the transducer is physically modified and thus so is its electrical output. Draw string displacement transducers consist of a sprung potentiometer attached to a coiled string. As the string is pulled out, the potentiometer rotates one way; when the tension is removed, the spring rotates the potentiometer in the other direction winding the string back in. Such sensors are commonly capable of operating on a reasonably large voltage range and so are easy to accommodate in many instrumentation systems.

Non contact displacement transducers are attached at only one end of the space being measured

but still require a clear path to the other component. Laser displacement transducers, operating on the triangulation principle, transmit a beam of highly coherent light at a target. Some portion of that light is reflected back to a receiver housed in the same chassis as the source. The position of the light on the receiver, which can be known precisely due to its coherent nature, allows the unit to calculate the distance between the source and the target object. This kind of transducer is moderately complex and requires internal calculation. The outputs are commonly digital but most transducers have the option of also generating an analogue, voltage based, output.

Accelerometers are used to measure accelerations within mechanical systems. With respect to time, acceleration is the differential of velocity, which is in turn the differential of position. In the work described in this document, accelerometers have been used to infer the direction of movement, as well as to compensate for movement effects in other measurements, and also to provide a measurement of the interactions at the wheel/rail interface. Lots of different technologies can be used to produce accelerometers, but most rely on the fact that a body in motion has inertia and will continue in motion until brought to rest by an external influence.

Some of the most common systems operate using a piezoelectric effect. In these systems, a moving component provides the loading force for a piezoelectric element which then produces a small, but measurable voltage. This is commonly amplified within the accelerometer before being presented as an output. All of the accelerometers used in the work described in this document are MicroElectroMechanical Systems (MEMS). MEMS accelerometers tend to rely on a cantilever system whereby, under acceleration, a proof mass depresses a number of the cantilevers. The electro-mechancial aspect of the system is commonly capacitive and affected by the other end of the cantilevers. MEMS accelerometers are commonly used in high vibration environments because the small size of their moving components means that they have a high fatigue resistance. Accelerometers commonly operate in a single axis with three often being

combined to produce tri-axial units. These are usually more accurate than externally combining three accelerometers as a single compensation for temperature and amplifier offset can be used.

Rate gyroscopes produce an electrical signal proportional to an observed rotation. They operate around a single axis and are rated in the number of degrees of rotation they can measure per second. The bandwidth and stability are the other important factors. As with accelerometers, a number of different technologies are used in the manufacture of rate gyroscopes. Among the most expensive and accurate rate gyroscopes are those built around fibre optic technologies [4]. These are prohibitively expensive and not commonly robust enough for use in a railway environment. Instead, the work described here makes use of Vibrating Structure Gyroscopes (VSG) and MEMS VSGs. In these units, a vibrating structure is excited to have a linear velocity. When the unit is rotated, the Coriolis force translates this velocity to another part of the structure. This secondary vibration is measured and it is this signal that is proportional to the rotation.

The signals output by the various different types of transducer are generally in fairly standard formats suitable for modern, commercial, data loggers. A data logger is essentially a data recording device. In the past, chart recorders would have represented varying analogue voltages as a graph against time. Modern data logging equipment operates at far higher speeds and so uses electronic storage media rather than the roll of paper used by the chart recorder. Before data can be stored electronically, it must be converted from an analogue to a digital format. This analogue to digital conversion is performed within the data logger.

Data loggers come in a variety of different forms. Some of them are complete and self contained units developed for a single application. Units like the CR800 from Campbell Scientific [5] are designed to be used in combination with a standard set of transducers. When purchased, many are pre-configured for the end user who has very little interaction with them beyond the

main power switch and a download button. This kind of data logger is traditionally used for environmental monitoring and so operates at a very low sample rate.

Some companies make highly customised data logging equipment that is suitable for use in a railway environment. LEM make an energy metering unit specifically for rolling stock. The EM4T (described in [6] and used in [7]) is suitable for use on both AC and DC systems and can operate at a variety of standard railway voltages. The system can be customised by the user to sample between 1 and 60 times per hour. This is adequate for most industrial metering purposes but little use in most research contexts.

Other data logging equipment is far more customisable. The LGR-5320 series from Measurement Computing [8] is designed to be used with any transducer with a standard voltage output. The unit is configured via a computer using custom software but records locally to a memory card before uploading the results to the computer in the same way that the Campbell Scientific logger does.

One of the most successful companies in the field of data logging is National Instruments. National Instruments make a variety of data logging equipment from smaller units that convert standard voltage inputs to digital signals and transmit them immediately to a PC, such as the NI USB-6009 [9]; to powerful data acquisition, control, processing and storage solutions such as their flagship, modular, CompactRIO system [10]. The CompactRIO system consists of a backplane, a control module and a number of input specific acquisition modules. These can be basic voltage input modules, thermocouple input modules, strain gauge input modules, relay controllers, output modules etc. While the CompactRIO solution is very powerful, the customisability coupled with the rugged nature of the product makes it far more expensive than many of its competitors. It is, therefore, more commonly used in low volume (or high profile) systems.

Systems such as those offered by National Instruments offer the flexibility required to customise an instrumentation system to meet an exact specification. The low volume nature of most high-end instrumentation systems, particularly those used on rolling stock, would suggest it to be the ideal solution. In some situations, however, the footprint of the equipment can be a constraint, in others the data storage capacity is the limiting factor. In these situations, the only acceptable solution is often to build custom data acquisition and logging equipment.

2.3   Targeted Instrumentation Systems

Although this thesis argues the case for complex instrumentation systems delivering a broad range of highly integrated data streams, this is not necessarily appropriate for all instrumentation applications and in fact there are many complete systems available which are built from simpler architectures. This will be demonstrated in the coming chapters where lessons from simpler systems are combined to demonstrate an example complex instrumentation framework and some of the benefits that it can bring. This section showcases some examples of instrumentation systems that use comparably straight forward architectures. The reasons for this selection may be that they are targeted to a particular measurement and do not require other signals to contextualise the data. Alternatively the systems may be operating for extended periods and looking to increase the levels of system robustness. In a similar vein, they might be deliberately simple in order to facilitate integration with existing technologies. Or possibly the systems are simply not that complex because their designer's do not feel that they need to be.

Reading the literature, you may think that the idea of using machines rather than humans to inspect the condition of your railway infrastructure is a modern one. A Modern Railways article from 2004 [11] makes the case for replacing manual track inspection with automated systems on safety grounds. While the motivation may be more modern, condition monitoring of the

infrastructure from a railway vehicle has been around for much longer. Back in 1974, a pair of East coast Americans published a paper [12] in the IEEE transactions on industry applications. The paper describes the use of a mechanical feeler system fitted to a high-railer inspection vehicle. The system was capable of identifying areas where the track was statically out of gauge but did not detect instances of dynamic gauge widening, where the loading of the track is responsible for deformation affecting gauge. The paper then goes on to propose a system in which a regular vehicle is instrumented with capacitive proximity sensors. These would monitor the displacement between the sensor tips and their corresponding railheads. Unfortunately, the system described operated in the time domain and so an operator was still required to insert markers into the data stream at locations of interest. The system, or at least one very much like it, is still, however, the subject of a patent [13].

By the early 1990s, railway inspection had moved on to use optical methods. A paper from Japan, which uses lasers in a system for optical inspection of the area surrounding the track, is described in [14]. The paper proposes the use of a laser to highlight the profiles of objects in its field of vision. Two optical image sensors are then used to identify the profile by detection of the laser line. The use of the laser dramatically reduces the complexity of the image processing required. This is beneficial as the system is nearly twenty years old and the proposed speed of operation was 40 km/h. More recently, variations on the laser gauging system described in [14] have been adopted for use in the UK. An article [15] in the journal of the Permanent Way Institution describes several laser gauging survey systems and the software that manages them and makes their outputs useful for a modern railway network.

The system described in [14] is designed to be mounted on a truck, capable of operating on both road and rail systems. It is, of course, not the most famous railway inspection vehicle in Japan. That title would undoubtedly go to Dr Yellow the 900 series inspection vehicles operating on

the Shinkansen lines. [16] suggests that, prior to 2000, most optical track inspection systems operating in Japan relied on simultaneous measurements from three points along the length of a vehicle. The article, published in 2003, explains that this three-point measurement is not suitable for use on the Shinkansen lines because, for structural reasons, the additional bogie cannot operate fast enough. Instead it summarises work in which displacement sensors are positioned at both ends of both bogies of a car and three of them are used depending on the direction of travel. The system relies on there being no flexibility in the car body, which is strengthened and has a laser shone along its length with detectors to measure deflection. These measurements are used in the data processing to compensate for the lack of the third bogie. The new, two-bogie system is also described in [17]. The paper suggests that the two-bogie version of the system can operate at speeds of up to 210 km/h. The paper then goes on to explain that even 210 km/h is slow enough to be disruptive to passenger services and describes a version capable of operating at 275 km/h. It explains the development of the two-bogie system for use on the East-i inspection vehicle. The system uses a modified E3 bogie to overcome some of the optical restrictions which applied in previous versions.

Around the time that the Japanese were implementing optical obstacle detection systems, work was being done in the UK to upgrade the central line of the London underground. Instrumentation was used to establish the quality of the track and its capability to support the increased frequency of service. This instrumentation, described in [18], makes use of an optical system called Automated Video Inspection (AVI). AVI does not really automate visual inspection, but it allows the user to consider several visual inputs simultaneously and without the need to be physically present. The paper also talks about a pair of proprietary ride quality and track assessment systems called Macminder and Mactrack. Macminder is a box of accelerometers, mounted in the vehicle body, to assess ride quality for passengers. Mactrack is a similar box of accelerometers but with some of the signal filtering removed and the option of connecting

additional, external, accelerometers. It is designed to monitor track quality. The paper describes the use of both AVI and Mactrack, with external accelerometers mounted on the axleboxes, to measure the condition of the central line track. By combining the two technologies, it is possible to perform visual inspection only of areas highlighted as potentially problematic in the inertial measurement data.

Of course integration of the two proprietary systems, AVI and Mactrack, was not trivial. It was thought that exporting the datasets might be an option but this apparently was not the case. In the end, alignment of data was achieved by pointing a fifth AVI camera at the Mactrack live output screen. Clearly greater integration of systems was required. A paper from British Rail research [19] describes another system for the measurement of track. In this system, measurements in the vertical plane are performed using inertial measurement techniques, while optical systems are still used for measurements in the horizontal plane. Displacement transducers and accelerometers are used to measure vertical track irregularities, while projectors and line scan cameras are used to monitor gauge. Curvature and cant measurements are performed using gyroscopes. The system described in the paper was designed to incorporate all of these technologies and so does not require any additional integration processing. Real time thresholding and alert analysis is performed with the signals also being recorded for further analysis.

Measurement systems such as those described in [18] and [19] have the potential to lead to enormous improvements in track maintenance efficiency and so much investment has been made in their development. One of the more recent systems deployed by Network Rail is referred to as the New Measurement Train (NMT). The NMT was built on a tight budget and is primarily a converted Intercity 125. An RSSB report [20] explains that the NMT entered service in June 2003 and provides measurement of 98% of the network. The data that it produces is used to populate the Trackmaster database but it can also identify exceedances in real time.

While Network Rail elected to produce their own track-recording vehicle, many other nations chose to simply purchase commercial equivalents. In January 2000, the Austrian Federal Railway (ÖBB) took delivery of an EM-SAT 120 track survey car and an EM 250 high-speed track recording coach, both manufactured by Plasser & Theurer [21]. The EM 250 is capable of operation at speeds of up to 250 km/h (155 mph) and the data it provides are used in the planning of ÖBB's track maintenance operations. The EM-SAT 120 is a self-contained system used to determine correction values, with are then used by tamping machines. The system replaces a manual process and operation is limited to a few km/h by the tamping machines that it accompanies.

Getting the appropriate approvals to operate track-recording vehicles on, particularly European, networks can be a lengthy, difficult and expensive process. [22] suggests that it can be up to 5% of the cost of the vehicle and that this is made more significant as the number of vehicles typically required is low. One approach to this is to combine several vehicles. For example, instead of operating separate tamping and track survey vehicles, it may be possible to combine these. This is the subject of an article in the proceedings of the Permanent Way Institution (PWI) [23]. The article refers to the use of Absolute Track Geometry in the upgrade of the West Coast Main Line (WCML). Two systems were trialled, one being the EM-SAT 120 from Plasser & Theurer, and a PALAS system mounted directly on the tamping machine.

As [11] explains, dedicated track inspection vehicles such as those described in [20] and [21] undoubtedly record far more information, far more accurately than manual inspection alone. They can operate safely among other services and provide less variation in the reports that they generate. This is significant as the standardisation of reporting allows a more effective maintenance schedule, with tighter tolerances, to be developed. These vehicles combine optical and inertial measurement techniques to perform, among other things, ballast profiling, six-foot

clearance and vegetation detection, in addition to standard track geometry analysis. The data from the measurement vehicles can be combined with trackside measurements to generate a far more complete picture of railway health, however, this approach is not suitable in every case. The paper explains how on typical rural branch lines, the inspection process is often combined with minor maintenance which would still need to occur, even if the inspection component of the task was automated.

The vehicles described in [11] are designed to operate, more or less, continuously. This brings its own problems, particularly in the use of optical systems. The paper refers to having to remove dead flies and brake dust from the lenses. Image processing is also computationally expensive and so a lot of work has been undertaken to move track inspection techniques away from optical methods. Papers such as [17], [16] and [24] tell us that the systems for track inspection used on the Shinkansen lines already make use of accelerometers and displacement transducers. A paper presented at the World Congress on Railway Research (WCRR) in 1997 [25], suggests that this may not be adequate to see short wavelength track defects such as rail surface roughness. The paper puts forward the case for the use of axlebox accelerometers, higher sample rates and frequency filtering during processing. With these measures in place, the paper claims to see track roughness around welds, corrugations and loose sleepers.

A Chinese system called a Track Geometry Real-time Inspecting System (TGRIS) is described in [26]. The paper uses a system based entirely on inertial measurement and displacements. No optical systems are used. The system is based around measurements made from an inertial platform housed within the vehicle body combined with measurements from displacement transducers linking the body with the axleboxes. Although some comparisons are drawn with other commercial systems, it is hard to gauge how the accuracy of the system suffers from the mounting location within the vehicle body. The paper suggests that the analogue to digital conversion

should be performed using 16 bits although this is derived from the predicted range rather than the visible signals. The paper also suggests that the central processor timer should be triggered at 1 kHz. This seems somewhat unnecessary for systems operating within the vehicle body.

Sometimes, it is possible to chart the progression of a system from a low TRL to a higher one. An example of this can be found in [27]. The paper describes how a collection of small / medium sized businesses in Germany worked together to develop a system for monitoring bogies and their interaction with the track, the aim being the early detection of accident threats. The paper describes the instrumentation of a bogie with piezo-electric accelerometers, acoustic probes and displacement sensors before measurements were taken both in a laboratory and then in the field. The development process, highlighted by the progression from static to dynamic laboratory tests and then two sets of field trials, culminated in a patent with further work proposed to develop more complex laboratory test rigs and prototype commercial instrumentation. Another patent for an inertial measurement based monitoring system where a Global Positioning System (GPS) is used to locate identified defects is described in [28].

Another good example of progression through the TRL levels can be found in [29] and [30]. The papers, first from the German Institute for Safety Technology (ISTec) and then from ISTec and Deutsche Bahn, describe work towards condition monitoring for the German InterCity-Express (ICE) trains. The papers build on the work in [31] and describe a progression of tests from rolling rig experiments, through simulated faults to full field trials. The described instrumentation primarily consists of tri-axis accelerometers and temperature sensors. The paper explains that the initial rolling rig experiments were used to develop the instrumentation before the first trials of the onboard Measuring and Assessment System (MAS) were performed. Few changes were made at this stage with the rolling rig instrumentation consisting of 13 accelerometers and the MAS system consisting of 12. The difference being the removal of an additional accelerometer

mounted on the centre of the bogie frame.

Extending this development a stage further, the results of the rolling rig and MAS trials allowed development of the RailWay COndition MOnitoring System (RW COMOS). This system consists of just 8 tri-axis accelerometers of which 6 are mounted directly on the bogie. These sensors are a subset of those used in the rolling rig and MAS tests. The development from research to product development is therefore clearly visible.

Fault simulation occurred at the rolling rig and MAS testing stages. The faults simulated on the rolling rig were reasonably extensive but focused on the wheelset. They included tread wear, non-circular wheels, flat spots, bearing lubrication failures, worn guiding sleeves and malfunctioning anti-roll dampers. Simulated faults in the MAS tests were more limited consisting primarily of tread wear, sleeve wear and anti-roll damper malfunctions. This is still an impressive list for trials involving an actual vehicle.

While many of the papers discussed have used railway vehicles to monitor the condition of the track, and [29], [30] and [31] have suggested that it is possible to consider the interaction between the track and the vehicle, it also possible to use similar instrumentation techniques to ascertain things about the vehicle itself. A paper from the East Japan Railway Company [32] considers the use of instrumentation to monitor the condition of bogies without periodically dismantling them. It suggests that vibration sensors used in the vehicle body can be used to monitor ride quality, and detect hunting and derailment. It concedes that body mounted instrumentation is not suitable for preventative bogie maintenance.

Another paper that describes the purposeful derailment of a vehicle for testing is [33]. In this case the paper is from Austria and is concerned with the detection of a derailed freight vehicle. The proposed system is installed in the locomotive and uses vibration and longitudinal force sensors to detect a derailed freight car elsewhere in the train.

As previously mentioned, most instrumentation systems record data in the time domain, sampling at a regular interval. To use the data, it must often be converted to a spatial domain which can be difficult if the vehicle is not travelling at a constant speed. The traditional approach to this problem is to record a tacho signal from the vehicle along with the data. Often, a toothed wheel attached to the axle is combined with an inductive sensor to generate pulses at spatially regular intervals. These pulses are used to calculate the distance that the wheel has rotated. The system must be re-calibrated whenever the wheels are maintained. The problem with this system, as described in [34], is that the rotation of a single wheel only corresponds to the movement of the vehicle if that wheel is not slipping. In some cases the toothed wheel is fitted to a non-motored axle, however, on some vehicles all axles are motored and this is not possible. In these cases, where wheel slip can obscure the position, alternative systems are used.

Automatic train control systems, such as those used by the Docklands Light Railway (DLR) in London [35] and the SkyTrain in Vancouver [36], report train position to a central computer which then advises the vehicle of movement. In the case of the DLR, communication is via an inductive loop laid between the running rails. The position is established using an axle driven pulse generator combined with Automatic Train Protection (ATP). Automatic train control is also used on the London Underground. On lines such as the northern line, all axles are motored and so wheel slip and thus loss of position are common occurrences. In these cases, manual control is used to re-align the trains.

Since the 1980s, Japanese railways have also used automatic train control systems. They too use axle generators but combine them with trackside transponders so that distance is measured only from the nearest transponder. This reduces the potential for errors induced by wheel slip but does not eliminate them completely. A paper published in 2001 [34], describes the use of commercial range sensors, in place of single point transponders, to continuously re-align the axle

generators.

Axle based tachometer units, although susceptible to wheel slip, operate to reasonable precision. Underground vehicles often have thirty teeth around the wheel diameter while metro vehicles, such as the Class 507 stock operated in Liverpool, have thirty or ninety depending on configuration [37]. Another technology which can be used to identify the position of railway vehicles is the Global Positioning System [21]. The problem with the use of GPS is its accuracy. The way that GPS works means that noise caused by atmospheric effects can introduce inaccuracies of up to 10 m and errors due to objects obscuring the signal path can be anything up to 20 m [38] or even more when reading are being taken while in motion. In comparison to a tacho system where each wheel revolution is divided into ninety sections, this is rather poor. GPS can however be used to give an approximate position which is used to direct inspection teams or track workers. This is the application described in [31] which is a paper primarily concerned with the use of track inspection vehicles rather than manual patrols on safety grounds. The paper describes the difficulties associated with manual track work and the first two phases of the EURAILSCOUT GB project. The first phase involved the conversion of a train to perform "virtual inspection" where signals from five cameras and a commercial geometry measurement system were passed to a "patroller" for consideration. The second phase used computer processing to replace the "patroller" and considers overhead power systems as well as track and rail condition. Both systems mark their results using GPS signals so that further work can be scheduled. The paper suggests that this is to an accuracy of ± 1 m.

The positional accuracy described in [31] is clearly far better than traditional GPS systems can guarantee. [39] describes the POS/TG systems operated on the EURAILSCOUT UFM-120 vehicles. The article suggests that accelerometer based positioning systems are inherently compromised by their associated processing [40], that contact based systems cannot operate

at high speeds and that manual positioning is inherently inaccurate. The POS/TG system presented in the article uses a combination of GPS and inertial measurement systems to improve accuracy. The article suggests that even greater accuracy could be achieved with the inclusion of differential GPS systems. This idea is also being incorporated into the EM-SAT 120 track survey vehicles as mentioned in [41].

Another example of basic GPS accuracy being adequate for track inspection and maintenance is provided by a system described by the US department of transportation in [42]. The paper considers the measurement of vertical roadbed stiffness known as track modulus. This is particularly significant on freight lines where there the axle loads are so much greater. The measurement is performed optically, using multiple lasers and a camera. Data is recorded along with latitude and longitude measurements obtained from a GPS. These are then overlaid on maps to assist with the identification of the cause of, and the direction of track workers to resolve any discrepancies.

While GPS is a powerful technology that can be used to direct maintenance, it can only do so once the GPS tagged data has been downloaded. One approach is to combine the GPS with communications systems to upload the tagged data in real time. This is described in [43] which describes the commercialisation of the systems developed in [44] and [45]. The paper is written by a company and essentially describes a product that has been developed from work funded by the US department of transport. A counterpart paper describing the higher-level objectives of the work has also been published by the department of transport itself [46]. The commercialisation relates to the Sensor System Boxes (SSB) which can be used with dedicated or existing sensors. The SSB performs threshold based event extraction which it tags with GPS information and uploads via a Geographical Information System (GIS) to a data centre. Communication is bi-directional allowing the data centre to request a snapshot of any or all of

the sensor inputs at any time.

While it may be acceptable for temporary instrumentation to consume large quantities of space and power, anything being installed in the longer term has to be efficient. Axle monitoring is a good example of this as space considerations underneath a vehicle are important in terms of maintaining gauging. Furthermore, the environment is electromagnetically demanding and so smaller equipment, which can be mounted closer to the sensors, reduces noise susceptibility. Also the environment is mechanically harsh (typically 100 g) and so smaller, lighter components are beneficial. The data loggers presented in [47] are mounted directly on an axle to maintain proximity to the strain gauges that they service. The units used include signal conditioning, filtering, data storage, 16 bit analogue to digital conversion and wireless communications and are enclosed in an envelope not exceeding 27 x 22 x 8 mm.

Sometimes the requirement for miniaturisation is directly linked to the operational capabilities of the instrumentation. An example of this can be found in the area of pantograph instrumentation. Pantographs are instrumented to identify the forces and accelerations at the point of contact with the overhead power supply cables. These things are also affected by the aerodynamic performance of the pantograph and so any instrumentation should have a minimal effect in this area. One system [48] involves the miniaturisation of the instrumentation such that all cables and components can be mounted inside the body of the pantograph thus minimising the aerodynamic impact.

2.4 Complex Instrumentation Systems

In the forthcoming chapters, this thesis will describe a number of instrumentation systems based on moderately straightforward architectures. Each of these, while adequate to the task for which it was designed, will have limited extensibility. As such, they are generally similar to the systems

described previously. This thesis will ultimately present a more complex and extensible instrumentation system based on a framework which allows a much more complex and integrated approach to data collection and management. This section will describe a number of instrumentation systems which are each in some way similar to this, more complex, instrumentation strategy.

Complex instrumentation systems are commonly used for large scale instrumentation projects where little is known about the target prior to the instrumentation being installed. By using an expanded sensor set, the level of data available is increased and so is the possibility to identify the key signals and the relationships between them. Complex instrumentation systems are therefore used primarily in larger instrumentation projects. These project tend to involve some low level TRL components which means that the work involved in the projects is often published in greater technical detail than those higher TRL level systems which tend to be focused on business outcomes.

Examples of this can be seen in papers such as [49]. The paper does not really present any background beyond the fact that accelerations are measured on railway vehicles, but the level of technical detail is far greater than any of the papers discussed in the previous section. The system is divided into a series of sub-systems which are each then described in detail. System block diagrams and internal communications schematics are presented along with bullet lists of technical specifications and the key components. A summary of operation is also presented.

Some papers take the level of technical detail they present even further. One paper [50], which describes trackside monitoring to detect axle loads through strains in the rail, starts by presenting system layouts, block diagrams and component descriptions. It then goes on to give a full schematic diagram, including component values, for its signal conditioning and analogue to digital conversion modules. The paper then describes the procedures used for calibration and

validation of the instrumentation and a state diagram for the algorithms being used. As with [49], the systems described in the paper have been used in a railway environment, but this is only mentioned briefly in the introduction and conclusion sections.

Some technical papers not only describe the practices and components that they use in their instrumentation systems, but also the architectures and how they are used in the context of the railway. These systems tend to be slightly more developed but are still generally low on the TRL scale. One example of such a system is the Health Cards system. This was produced by Queensland University in Australia and has been developed over a number of years. The system, described in [51], uses the data from multiple tri-axis accelerometers and Micro-Electro-Mechanical Systems gyroscopes to establish a measurement of vehicle body inertia with six degrees of freedom. These measurements are used to generate event flags based on exceeding a three-dimensional spherical threshold. The generated events are then transmitted back to a control centre.

The system takes measurements from the vehicle body as instrumentation of bogies is comparably more difficult to achieve. As the instrumentation is post the secondary suspension, the results will be smoothed and so the accelerometer and gyroscope quality is less significant. This provides access to comparably inexpensive options. ADXL202/10 accelerometers and ADXRS150 gyroscopes, both from Analog Devices are used. Each car has a single node to which all of its accelerometers and gyroscopes connect. The nodes are managed by Rabbit 3000 processors, operating at 40 MHz.

The individual Health Card nodes communicate as part of a network, however, only the events are transmitted as bandwidth is comparably low. There is a master node in the locomotive which communicates with the control centre. The communications network is achieved using the same connections that power the nodes and power line transceiver devices. The connections

themselves are part of the installed train bus.

One way in which modern instrumentation architectures differ from historical ones is the method of reporting the data that they acquire. It is not uncommon now, particularly in systems with on board processing, for the system to upload some or all of the recorded data to a remote location. As described in both [51] and [52], later versions of the Health Card system do not interact with the existing train bus. Instead, the nodes are solar powered, with battery backup, and communicate over a Bluetooth network. As the available power is limited, and the Bluetooth communications expensive in power terms, some concessions have been made to the system. The gyroscopes have been replaced with additional accelerometers mounted at the corners of each car. This consumes less power while allowing the three axes of rotation to be calculated from the differences in acceleration over the length of the car.

While the Health Card system allows the nodal processors to generate events for each car and then only communicates the events, some systems are forced to collate the data from each subsystem before processing can occur. This is the case when using accelerometers to monitor and control the levitation gap of a maglev train. In this situation, it is vital that the recorded accelerations are transmitted to the central processor both quickly and accurately. This communication is the focus of [53]. The final system does not use an error handling bus, therefore the Baud rate is reduced to give the maximum possible bus reliability. The system uses a thirteen-bit framing message and twelve-bit data messages with a resolution of 0.0068 g/bit. Although the choice of processor is never revealed, it is unlikely that the twelve-bit limitation comes from the processor. More likely it is the native precision of the accelerometers. No suggestion is made that the use of twelve bits was a design decision although it is possible that the twelve are the top bits of a larger signal where the lower bits are predominantly noise.

Another example of a hugely complex instrumentation system is described in [29]. The 13 tri-axis

accelerometers used in the system generated 39 channels of signals alone. The system also used other sensors such that the acquisition system had to cope with 70 channels of data. In a paper from America [44], the described instrumentation is based around a PC/104 computer. The paper makes a case for the instrumentation of vehicles based on a federal track safety standard (49CFR213.333), which states that daily measurements must be made for vehicles travelling above 125 mph. The PC/104 computer is used more for analysis and communications than for managing sensors as the instrumentation is reasonably straightforward. Sensors consist of vertical and lateral accelerometers mounted on the body and a lateral accelerometer mounted on the truck frame. The processing involves the use of thresholds to raise alerts for various sections of track. Two cars per train have been equipped with the system which communicates with central data centres. The choice of two cars per train is primarily for redundancy and so the systems operate independently.

A similar, although more complex, system is described in [45]. This time, the system is concerned with the condition of freight vehicles. Instrumentation consists of thermocouples, accelerometers and displacement transducers, to monitor bearings, wheels, trucks and brakes in order to detect defects and derailments. As with [44], the system uses PC/104 computers, this time mounted along with signal conditioning and communications systems in a Sensor System Box (SSB) contained within each car. In addition to communications with a central data centre, the individual SSBs also communicate. This allows a local user front end to be implemented using a laptop computer within the locomotive. The system can then inform the user, in real time, of alerts raised by signals exceeding thresholds.

A few years later, the same group who published [45] published an extension of the work [54]. The paper extends both the range of instrumentation and also its complexity. The previous version was largely confined to acceleration measurements, the system described in this paper

also includes sensors to monitor bearings, brakes, wheels and trucks as well as actuators for remote control of hand brakes, angle cocks and cut levers. In previous versions of the system, all of the sensors in a car were linked directly back to the main Sensor System Box where analogue to digital conversion occurred. In the new version, small embedded processors are associated with each sensor or location containing multiple sensors. These processors are responsible for the analogue to digital conversion. The digital versions of the signals are then passed back to the SSB using a Controller Area Network (CAN) bus, which links all of the processors together. This means that cables containing analogue signals, which are more susceptible to noise can be shorter, which is good for accuracy, and the main processor contained in the SSB can be less powerful, which is good financially. This illustrates the approach of testing with off the shelf equipment and then developing a product with a more commercially viable solution. The local user interface, previously a laptop within the locomotive, is now a more formal arrangement with certain events, such as hot bearings, derailed wheels and hunting, reported there automatically. The system also appears to be moving away from transferring all of the data and now relying on the local processing to communicate only the alerts.

A PhD thesis from Australia [55] describes instrumentation applied to coal infrastructure wagons. The thesis is concerned with the change in wheel / rail interactions over time and uses an extensive, if not disparate instrumentation set. Eight strain gauges are used in conjunction with a longitudinal body mounted accelerometer, a bogie yaw transducer, two brake cylinder pressure sensors, a speed monitoring device and GPS. While continuously sampling data from this many sensors would generate an extremely rich dataset, the practical implications are difficult to contend with and the subsequent data processing would be extremely computationally intensive. As the study is concerned with changes over time, it instead only records data over a specific section of the network, identified by the GPS. In addition to the targeted data logging, the system is also in communication with the railway operator through a two-way mobile phone

based communications link.

If the number of sensors, or the sampling rate, is low enough, it may be possible to record all of the data or even transfer it to a processing centre in real time. This is the subject of [56] which uses only three, single axis, accelerometers to monitor track condition from the vehicle. The system uses Code Division Multiple Access (CDMA) to transfer the three channels of acceleration measurements to a central processing centre which uses multiple threshold feature extraction to separate vibrations caused by poor track quality from vehicle vibrations.

In most cases, the dataset is too large to be transmitted to a processing centre. In these cases, the data is logged [55] or locally processed either by local computing [57] or embedded processors [51]. In the case of local processing, the progression towards embedded processors may require the systems to be modified based on the reduction in computing resources. This is the case in [58] where an optical system for track monitoring is being redeveloped to use embedded processors. The paper mainly focuses on the modifications required to make the algorithms compatible with the different processor architecture and capabilities. It ultimately presents two solutions, a fast option and a robust option.

This trade-off between accuracy and speed is the topic of [57]. In this case, neural networks are used as a second stage of processing for a Charge-Coupled Device (CCD) based optical track inspection system operating on the Milan underground. The authors claim that the available processing power is the only limitation that their system places on vehicle speed. The same would apply to the wheel / rail interaction monitoring described in [59] and [60], but in this case the authors divide the processing over multiple computers to increase the possible operating speed.

Depending on the application, it may be preferable to forgo real time processing in favour of the additional accuracy and simplicity available in post processing. An American study, presented

in [61], suggests that bogie mounted instrumentation is difficult to install and maintain. It proposes a system whereby measurements are taken from the vehicle body. Such measurements do not give as much information but the paper suggests that they are acceptable if combined with a model of the vehicle which includes its secondary suspension components. This obviously requires far more processing capability than would be sensible to install in the vehicle and so off site post processing is definitely required.

One of the major issues with post processing of recorded data is aligning it with the infrastructure. This is particularly noticeable when a railway vehicle is used as a tool of measurement and the first question that the Infrastructure Manager asks is "where was this data recorded". Although it is not concerned with data analysis, rather vehicle location for signalling, [62] presents a summary of several alignment techniques, some of which may be applicable to this task. Tachometers are a staple of the railway industry and are often used despite the known problems of wheel slip. Often they are used in conjunction with transponders which provide a known reference at given positions thereby allowing errors due to slip to be reset. Inertial Navigation Systems (INS) comprising accelerometers and gyroscopes can be used to identify the magnitude and direction of movement respectively. Acceleration based systems suffer particularly badly from noise and so are best used alongside other techniques or in environments with frequent stopping. Doppler effect sensors can be used as an alternative technology for identifying vehicle speed although commercial systems do not perform tremendously well in railway environments. Finally GPS systems can be used to provide position and movement information. These systems have limitations of accuracy which vary with the environment they are used in but do provide all the required information from a single source.

If instrumentation systems are fitted to the infrastructure and used to monitor vehicles, communication of data and alignment between datasets in post processing is substantially more

attainable. Such an approach would also allow multiple commercial systems to be used together in a coherent manner. In a project funded by the Australian government, multiple datasets were brought together to improve the overall level of understanding of railway vehicles [63] [64]. Initially the two datasets had to be combined to a single repository before the integration could be undertaken to allow alignment of the data based on the vehicles that passed through the two systems.

For vehicle-borne systems, integration at the time of acquisition is still preferable as the data is often continuous rather than divided discreet train passage events which are more easily aligned. In Australia, the GPS approach has been used, along with a tilt sensor and a gyroscope, to generate curvature and gradient maps for existing railway track [65]. These maps are used to build simulations of the railway for driver training. The GPS system has been shown to provide adequate information when four or more satellites are being tracked. This is not the case the majority of the time due to environmental features and vehicle speed. The tilt sensor and gyroscope are used to fill in where the GPS system fails. The obtained accuracy, while not perfect, is considered adequate once the post-processing step of reducing the data to straight and curved track sections has been taken. This is not exactly contextualisation of data, but rather the use of one sensor to support the processing of another. It is therefore still valuable as it shows how multiple sensors can be used to augment the results obtained from a master data source, in this case the GPS.

While [65] uses a tilt sensor and a gyroscope to complete a dataset otherwise populated by a GPS, [40] and [66] build the dataset entirely from inertial measurement components. The work is not intended to develop infrastructure maps, but to identify track irregularities. The papers describe instrumentation fitted to both metro and Diesel Multiple Unit (DMU) vehicles, although only a subset of the instrumentation is used in each case. In addition to the developed instrumentation,

the papers are primarily focused on the data processing which must take place to render the results useful. The work in [40] is centred on obtaining a metric for vertical alignment. This is performed using data from a pitch-rate gyro and also requires some measurement of vehicle speed. The work presented in [66] is similar but uses a yaw-rate gyro to identify lateral alignment. While the authors do suggest ways that the findings of the papers could be used by permanent way engineers to facilitate track maintenance, they do not go as far as developing such systems.

Using just a tachometer and a yaw rate gyroscope, the key subset of the instrumentation described in [66], it is possible to generate a map of a vehicle's passage from the data that it records [67] [68]. The identified sections of travel can then be aligned with infrastructure maps to give precise alignment to features located in any associated recorded data.

While [61] proposed the use of a vehicle model to interpret data recorded on the body of the vehicle more accurately, it is possible to use the same vehicle model in a more proactive way to improve the ride quality. A paper from the China Academy of Railway Science (CARS) [69] proposes the use of a lateral semi active control system to improve ride quality and performance in high-speed rail links. The system is based around a series of acceleration and vibration sensors, controllable dampers, and a control system implemented using a PC.

The concept is taken further by another paper [70], which suggests that, while lateral accelerometers have been used for years as the inputs for such control systems, the technically significant component is the processing. The paper proposes the use of a fuzzy logic based control system to manage the tilt level of tilting trains.

Control systems can be applied to other aspects of a vehicle beyond the bogie. One paper [71] proposes the use of active suspension on pantograph systems to increase the reliability of the contact between the pantograph head and the catenary. In the paper, the position of the pantograph head is measured using position sensors, and the paper suggests that these may be

upgraded to accelerometers for reasons of robustness.

Another paper that involves the instrumentation of pantograph systems is [72]. The paper describes a series of instrumentation used to monitor the health of the entire traction system of a railway vehicle operating on the French high-speed network. The paper summarises the use of different kinds of instrumentation, from optical and thermal imaging through to mechanical vibration monitoring. Actual instrumentation isn't described in detail, however, as the core of the paper is related to scale modelling and signal processing. A more practically realised solution is presented in [73]. In this case, the paper describes a system used to monitor the health of the traction system on the Korean high-speed lines. The system is modular consisting of commercial sensor systems combined together using a commercial data acquisition system and LabVIEW running on PC. The modular nature of the system in [73] lends itself to individual discussions of the modules. This is the case in [74] which is a discussion of the practicalities of monitoring induction motor performance using temperature sensors.

While papers like [72] and [73] concentrate on a large system from quite a high level, some research takes a single component of a system and considers it in great depth. This is the case with [75] and [76] which describe a variety of instrumentation and analysis work done on a single train door. The objective of the work was to develop a fault detection and warning system for rotary operated train doors such as those used on the London Underground network. The paper describes the monitoring of airflow, air pressure, actuator rotary displacement, door position and operation time. A detection system, based on radial basis function neural networks was then used to perform the analysis. These smaller scale instrumentation systems can still be complex in nature and still benefit from advanced, robust, instrumentation frameworks. Such systems would not, however need to be distributed and so some of the benefits of a modular structure would be lost.

2.5 Conclusions

This chapter has presented a discussion of data acquisition technologies, the instrumentation architectures that make use of them, their use in the railway industry, and the data processing techniques that are necessary to extract useful information from the recorded raw data. Different reasons for applying instrumentation systems have been considered and thus such systems have been broadly classified as complex exploratory systems, or targeted detection mechanisms concerned with a narrower range of inputs.

The development of instrumentation for the railway industry has also been charted through time. In some cases this has been evident as application leading technology, where systems have generally become more powerful as increased computational capacity or sensor precision has become available. Alternatively, technology has been shown to lead application with developments in sensor technology bringing new instrumentation opportunities to the railway industry.

The consideration of the evolution both from tightly focused instrumentation systems to larger more powerful systems, and also from complex exploratory systems to more refined dedicated monitoring equipment has highlighted substantial differences in system architectures. Several different backbone systems have been encountered in node based systems, while monolithic data acquisition architectures have also been observed. Some consideration has been given to post acquisition integration of data. This appears to be a viable solution for discreet data such as event based alerts of rolling stock condition, but is less well suited to continuous data obtained on board the rolling stock itself.

CHAPTER 3

CASE STUDY 1 - ENERGY MONITORING

3.1 Introduction

This chapter presents the first of four case studies into the instrumentation of railway vehicles. The instrumentation to be presented was developed as part of a system for monitoring energy consumption on a Class 508 metro vehicle in Liverpool. The project was managed and conducted by the University of Birmingham with support and backing from the Department for Transport (DfT) and the local train operating company, Merseyrail. High-level project objectives from the DfT included establishing the specifications for retro-fit energy metering systems. Merseyrail's interests were based around using the system to map energy demands within their network with a view to identifying problem areas and targeting remedial action.

The access to the vehicle required to complete this project was provided specifically for the completion of the trials. The instrumentation system developed for this work was therefore limited in scope to include the minimal range of signals to complete the objectives specified in the original project documents as provided by the DfT. The dataset collected is therefore both thorough in terms of the requirements, while limited in terms of additional signals with which to interpret or contextualise the results. This case study demonstrates the wealth of

information that can be obtained through modern instrumentation systems, but also highlights the limitations associated with restricting the scope of the instrumentation system.

3.2 Motivation

Approximately 40% of UK railway track is electrified, with approximately 60% of passenger kilometres being travelled on the electrified network [77]. At present, billing for the electricity used is calculated using a model based system. Periodic meter readings are taken in electricity feeder stations and then compared to the total outputs of the models. The costs associated with any discrepancies between the models and the readings are distributed pro-rata among the train operating companies making use of that feeder station. This process is known as the wash up [78].

The wash up system does not encourage individual train operating companies to perform efficiently in terms of their energy use. Any savings that the train operating company might make will not necessarily be represented in the models, and the wash up process would then distribute their savings among the other train operating companies using the feeder station.

One solution outlined in the rail transport submission to the committee on climate change [78] is that all railway vehicles should be fitted with energy meters. This would allow more accurate billing for electricity based on the usage within the vehicle. This solution has been adopted by the UK government with energy metering equipment to be mandatory on all new rolling stock and to be retro-fitted to all existing rolling stock [79]. Energy metering implementations are to be in line with GM/RT2132 [80], the railway group standard implementation of the European railway energy metering standard BS EN 50463 [81] , and the energy management systems standard, BS EN 16001 [82]. A second edition of GM/RT2453 (currently in draft) will update the mandatory data requirements for rolling stock to include energy metering data.

In addition to basic billing procedures, metering the energy usage of railway vehicles has many further potential benefits. Network Rail now offer a reduction in the energy tariff for vehicles that can demonstrate suitable levels of regeneration. Virgin Trains carried out a five month trial in which they fitted energy metering equipment to a vehicle in their Pendolino fleet (390049). A report of the trial in Railway Gazette International [83] suggests the average proportion of energy consumed that was then returned by regeneration to be 16.8%. This is in line with the 16.5% tariff reduction for regenerating vehicles.

Energy metering could also be used as a tool for analysing journeys in relation to driving styles. Different train operating companies employ different systems for managing driver behaviour with some more prescriptive than others. By analysing the energy consumption on a given route, it may be possible to advise the driver of a more energy efficient approach to various sections of the journey. The Railway Gazette International article suggests that over a four week period there was a variation of 33.5% in the energy consumption per kilometre on the London to Manchester route with no additional driver advisory systems beyond the timetable.

Another area where energy meters may feature in the future is in comparisons of energy consumption between regular journeys and those suffering delays. With energy metering implemented, the additional cost incurred by the train operator due to a delay would become apparent. This cost could then be included along with delay minutes as part of the settlement system.

Finally, energy metering on board railway vehicles could lead to journey analysis for energy efficient timetabling. This could take two forms. Firstly, energy costs for each track section could be included in the development of the timetable such that regions of track with high energy costs were allocated more time and so a lower energy requirement. Secondly, the timetable could be better designed for regeneration. This is particularly important in DC networks where energy storage is uncommon and regeneration is only truly of benefit if another vehicle is available in

the same electrical section to accept the regenerated energy.

The standard governing energy metering on railway vehicles (GM/RT2132) states that energy data should be recorded at five minute intervals. It does not prohibit data being recorded at higher rates providing it can be aggregated into five minute reference periods. Meeting the minimum aspect of this requirement prohibits many of the additional benefits of energy metering described above. Many commercially available energy meters, such as the EM4T from LEM [6], while therefore acceptable by government standards, do not provide the level of data required to lever any substantial further business benefits. The work presented in this chapter involves the development of an energy monitoring system. This is different to an energy metering system in a number of respects. Firstly it operates at sufficiently high data rates to identify how and where energy is being used within the system rather than just how much energy the system is using; and secondly it records additional information thereby allowing this greater level of analysis.

3.3 Metering Different Traction Systems

The UK railway network is divided into a number of separate franchises, each operated by a Train Operating Company (TOC). There are currently 24 train operating companies operating in the UK [84]. Each train operating company is eligible to select their own rolling stock, usually through leasing companies. This means that there is a great variety of rolling stock in use on the UK railway network.

Possibly the broadest division for rolling stock is between locomotive and multiple unit systems. Locomotives consist of a single power car at the head of a train with un-powered wagons or coaches making up the remainder. Multiple units distribute the traction system throughout the train. Due to the higher axle loads and stresses put through the wheel/rail interface, locomotives are rarely used where multiple units are a feasible alternative. They are primarily used for freight

operations with most passenger services operating multiple unit solutions.

Another major division is the power source. There are currently three types of power source in operation in the UK: diesel engines, AC electrification at 25 kV, and DC electrification at 750 V or 1.5 kV. Options also exist for the final element of the traction system, with mechanical drive systems taking their place alongside both AC and DC motors.

Between the power source and the final elements exist a number of other components which can be combined to produce different styles of drive train and hence different types of vehicle. Purely mechanical vehicles make use of hydrodynamic gearboxes, while vehicles with all or partially electrical drive trains may also make use of alternators, rectifiers, inverters, rheostat controllers or electrical choppers. Table 3.1 lists some of the most popular combinations in use on the UK railway network.

Table 3.1: Drive trains by vehicle type

Vehicle	Power Source	Drive Train
Diesel Locomotive	Diesel Engine	Alternator | Control | Electric Motor
Electric Locomotive	25 kV AC	Transformer | Rectifier | DC Motor
Diesel Multiple Unit	Diesel Engine	Hydrodynamic Gearbox | Wheels
Diesel-Electric Multiple Unit	Diesel Engine	Alternator | Rectifier & Inverter | AC Motor
Electric Multiple Unit	25 kV AC	Transformer | Rectifier | DC Motor
		Transformer | Rectifier | Inverter | AC Motor
		PWM Input Converter | Inverter | AC Motor
	750 V / 1.5 kV DC	Rheostat Controller / Chopper | DC Motor
		Inverter | AC Motor

Metering these different combinations of vehicle configuration, power source and drive train requires a number of different approaches; no one meter is likely to function as a universal system. The energy input into a diesel system can, to some extent, be monitored by monitoring the consumption of diesel fuel, or the mechanical operation of the engine. This is possibly the simplest approach to metering. Greater accuracy can be obtained if the vehicle has an electric component.

Metering an electrical system involves recording the currents flowing through some part of the system, and the voltage at that point in the system at precisely the same time. Combining the instantaneous current and voltage measurements gives a single measurement of power. Some monitoring applications may then require further analysis of this instantaneous power. Energy metering applications are more likely to record cumulative energy use over extended periods of time.

The location in a vehicle's power train that the current and voltage transducers are fitted determines the validity of the measurement. Fitting the transducers earlier in the power train ensures that the losses in the system are included in the measurement. Ideally, in vehicles with electric power sources, the transducers should be fitted as soon after the current collection equipment as possible. In vehicles using pantograph systems to collect power from overhead line equipment, this is usually quite easy as movement in the overhead cables caused by the passage of the vehicle mean that it is uncommon for a second pantograph to be used. This means that there is a single point of energy supply at which the transducers can be fitted. DC powered vehicles making use of third rail systems, where a conductor rail runs alongside the running rails to deliver power through shoegear, can be more complex as there is often a DC bus within the train linking several pickup shoes together. In this case, it is often necessary to fit transducers (particularly current transducers) to each set of shoegear.

3.4 Instrumentation

The Class 508 railway vehicle instrumented for its energy consumption was an electric multiple unit vehicle. Class 508s were manufactured in 1979/1980 and underwent a full refurbishment between 2002 and 2004. They are now used almost exclusively on the Merseyrail network where they work alongside a fleet of Class 507 vehicles. The 507 and 508 stocks are extremely similar with the main differences being the in car lighting and the air compressors. The instrumentation was developed to be fitted to either a Class 507 or Class 508 with the final decision being taken by the Train Operating Company rather than the university.

The Class 508 is, as mentioned, an electric multiple unit operating from a 750 V DC third rail power source. The instrumentation was therefore based on voltage transducer and current transducers. Full operation of a Class 508 requires 656 kW [85]. A standard configuration consists of an "A" car and a "B" car with a trailer car between them. The "A" and "B" cars are distinguished, among other things, by one carrying the compressor and the other the Motor/Alternator (MA) set. The two ends of the train receive common demands from the driver but traction system operation is otherwise completely independent. Each end has its own power pickup points, control systems and motors. Hence, the power consumption at both ends of the vehicle should be metered. In this case, the instrumentation was only fitted for a "one-day" trial so the decision was taken to approximate energy consumption as twice that measured at the "A" end of the vehicle.

The architecture selected for this instrumentation project was done so based on the minimal nature of the trials as dictated by the adherence to the requirements specified by the DfT. The system used a two point sensor / recording architecture as all of the signal sources were located together but not in an easily accessible location. Due to the safety considerations of instrumenting high current / high voltage systems, and the harsh electrical environment that the

sensors would be located in, remote analogue to digital conversion at the point of recording was deemed inappropriate. Instead, all of the sensors used were connected to a single analogue to digital converter managed by a single microprocessor located in close proximity to the sensors. This was then linked, through a fibre optic interface to a data recording unit. The fibre optic was selected for its resistance to electromagnetic interference as well as safety reasons.

3.4.1 Overview

The instrumentation developed for energy metering on the Class 508 railway vehicle considered the traction and auxiliary energies to be independent. These systems, however, were both fed from the same 750 V supply using a common bus. It was therefore possible to use a single voltage transducer combined with a pair of current transducers to measure the energy consumption. The instrumentation and transducers were all fitted in one of the equipment racks hung from the underside of the vehicle and highlighted in Figure 3.1. Note that the figure is of a Class 507 but that the equipment rack concerned is a common component with the Class 508.

Figure 3.1: Class 507 with equipment rack containing instrumentation highlighted

All of the transducers selected presented analogue outputs and so signal conditioning and analogue to digital conversion were used to convert the signals to more manageable digital equivalents. A central microprocessor managed the analogue to digital conversions and time-stamped the data before relaying it to a laptop contained within the vehicle body via a fibre optic serial link. Sampling rates for the signals were chosen based on a combination of the predicted rates of changes of the signals, and the processing capability of the microprocessor managing the conversions. The key consideration here was based on the objective of the study. For energy metering systems, sampling rates can be very low, however, for the energy monitoring system there was an intention to evaluate the perfomance of the traction system. This required the sampling to be sufficicently faster than the operation of the traction system. This is a camshaft based configuration and thus electromechanical with the fastest mechanical components operating in maybe 0.1 seconds. The selected sampling rate was massively faster than this, but ensures that any unexpected traction system effects will also be observed. The laptop was used to display the values from the transducers in real time, as well as recoding the data for further analysis. A diagram showing an overview of the instrumentation, including the division between the components in the equipment rack and the car body, is shown in Figure 3.2.
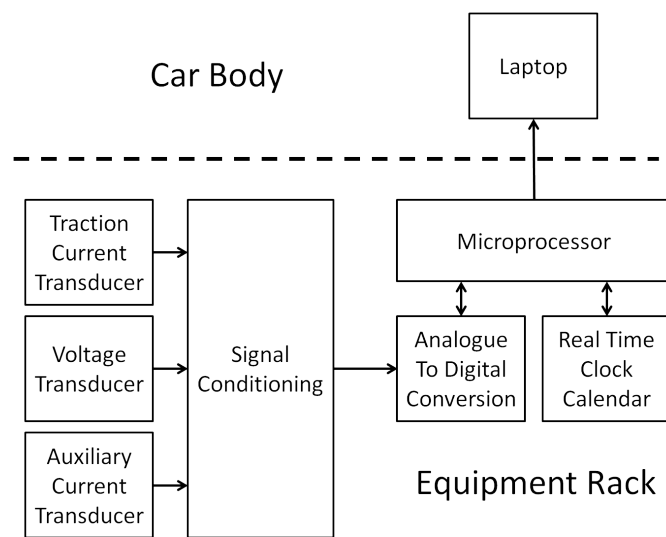
Figure 3.2: Overview of instrumentation system

3.4.2 Sensing technologies

The Class 508 is an electric multiple unit which uses a 750 V DC third rail as its power source. The 750 V is a nominal value as specified in BS EN 50163:2007 [86], the railway applications standard for the supply voltages of traction systems. The standard allows for lower and upper voltage thresholds of 500 V and 900 V respectively. Non sustained versions of these thresholds exist at 400 V and 1000 V, the lower difference being due to poor weather performance on the third rail network in the South of England.

The voltage transducer selected was the DV 1200/SP2 from LEM [87]. The transducer is from the LEM traction rated range which means that it has been approved in accordance to BS EN 50155 [88] (electronic equipment used on rolling stock), BS EN 51212-3-2 [89] (electro-magnetic compatibility) and BS EN 51024-1 [90] (insulation coordination). The nominal voltage for the transducer is 1200 V however it will operate between ± 1800 V and so covers the expected range. The output signal is in the form of a current which must be converted to a voltage for use with standard analogue to digital conversion equipment. This is achieved by passing the output current through a precision measurement resistor and measuring the voltage across it. The device is accurate to ± 11 V over its nominal range and within the relevant environmental conditions, so high precision resistors must be used in this conversion process.

A Class 508 can nominally draw 656 kW. At the nominal 750 V this translates to 876 A, although this is likely to be distributed over multiple power pickup locations. The traction current transducer selected was the LTC 600-SF [91], again from LEM's traction rated range. It is a closed loop current transducer which makes use of the Hall effect. This means that the transducer has a solid core and the wire that it is to be measuring the current in must be disconnected during installation. The transducer is designed for a nominal current of 500 A and has an operating range of ± 1500 A. Again the transducer has a current output and an accuracy

of ±8 A within its nominal range and the relevant environmental conditions, so high precision measurement resistors must be used.

The auxiliary current measurement was taken in a wire connected to one end of a 100 A fuse. The selection of the current transducer was therefore made to accommodate this range plus some transient spiking. The transducer selected was the HTR 100-SB from LEM [92]. Although not from the traction current range, the transducer meets a number of standards including UL 94 V0 [93]. This is a fire standard important due to the underground nature of parts of the Merseyrail network. The HTR 100-SB is an open loop Hall effect transducer with a split core for easy application. It is designed for a nominal current of 100 A but will operate within a range of ±200 A. The transducer produces a voltage output, which is acceptable as the analogue to digital converter is located near to the transducer. This also substantially reduces the power requirement of the transducer. Being an open loop sensor the accuracy is somewhat reduced compared to the transducer used for the traction current. The measurement range is, however, also substantially reduced giving an acceptable overall accuracy of ±2 A.

3.4.3 Electronics

The outputs of the voltage and current transducers are all analogue signals. These are passed to the analogue to digital converter before being packaged by the microprocessor for transmission to, and recording by, the laptop. The three transducers have two different types of outputs, current and voltage. The analogue to digital converter accepts only voltage inputs. The current outputs, from the traction current and voltage transducers, must be converted to voltages by being passed through a resistor and the associated voltage drop measured. The value of this resistor must be precisely known to maximise accuracy. The resistance value must also be carefully selected to optimise the range of the "output" voltage based on the current. Finally

the resistor must have sufficient power rating to accommodate the combination of current and voltage from the transducer. A number of small, accurate, high power resistors were used in parallel in the instrumentation.

The remainder of the signal conditioning consisted of Electromagnetic Compatibilty (EMC) protection diodes, low pass filtering around 8 kHz using simple Resistor / Capacitor (RC) networks, and the use of INA 128 instrumentation amplifiers on each channel for buffering.

The analogue to digital conversion was provided by a single LTC1859 [94] from Linear Technology. The LTC1859 is an 8 channel, 16 bit analogue to digital converter capable of operating at 100,000 samples per second. This speed essentially means that the converter is capable of producing results as fast as the microprocessor requests them. This is important as the current and voltage measurements must be taken as closely as possible in time. The device has a Serial Peripheral Interface (SPI), which uses two data lines and a single clock line. The SPI connection can be used to interrogate the device, specifying the channel to be sampled, the sampling mode, i.e. differential between two pins or single ended referenced to a common point, and the range of input signal that can be accommodated. The acceptable ranges are given in Table 3.2. As with any analogue to digital converter, it takes time for the internal workings to switch range and so, it is preferable to use signal conditioning to present all signals within the same voltage range. The instrumentation operated the analogue to digital converter in its ± 10 V range.

Table 3.2: Analogue to digital converter input ranges

Minimum Voltage [V]	Maximum Voltage [V]
0	5
0	10
-5	5
-10	10

The microprocessor selected to manage the system was the PIC 24HJ128GP502. The "H" series from Microchip are basically Digital Signal Processor (DSP) offerings with some of the dedicated signal processing peripherals removed. This means that they can operate more quickly than traditional PIC microprocessors and hence manage the high speed sampling required by the instrumentation. The maximum speed for the microprocessor is 80 MHz and it was operated at this speed. The 24HJ128GP502 has two serial ports and two SPI interfaces; one of each was used to communicate with the laptop and the analogue to digital converter respectively. It also has a single I^2C interface. This is less customisable than the other peripherals as it can only be assigned to two pairs of pins on the device. It was used to communicate with the Real Time Clock Calendar (RTCC) to obtain date and time information which was used to appropriately package the data.

The selected Real Time Clock Calendar was the DS3231 from Maxim. When interrogated by the microprocessor, it presents full calendar and time information, including leap year compensation, in binary coded decimal format over the Inter-Integrated Circuit (I^2C) bus. It is accurate to two parts per million, which is approximately five seconds per month, and can use a battery backup to maintain accurate time when the instrumentation is switched off. Most importantly, the DS3231 can output an extremely accurate pulse at a rate of 32768 Hz. This can be used in the rest of the instrumentation as a master timing signal.

3.4.4 Firmware

The PIC 24HJ128GP502 is made by Microchip. It can, therefore, be programmed in assembly language, the native language of all Microchip processors. The 24 series processors are, however, reasonably complex units with a large number of incorporated peripheral devices and so such processing involves a lot of manual configuration of registers. In addition to assembly language,

Microchip also produce a compiler based on the C programming language which can be used to manipulate their processors. This language was selected for use in developing the instrumentation as it provides a number of helper libraries which can be used to configure the peripherals automatically.

The design of the instrumentation specified that the current and voltage signals should be recorded independently for post processing into energy values. It also stated that the microprocessor should produce an energy result at a lower data rate which could be used for validation. This would allow any errors due to misalignment of the current and voltage signals to be identified. To achieve this, the system sampled at a higher data rate than either requirement and averaged the results over the relevant timeframes. This oversampling process provided an additional level of smoothing within the data, while not compromising the specified data rates. The master timing signal used in the system was generated by the Real Time Clock Calendar chip. This, very accurate, 32768 Hz signal was faster than even the oversample rate and so a hardware timer within the PIC was used to generate an interrupt based on four cycles of the timing signal. The frequencies used in the system are shown in Table 3.3.

Table 3.3: Frequencies used within instrumenation system

System Component	Frequency [Hz]
Processor clock	80,000,000
Timing signal from RTCC	32,768
Sampling rate (interrupt rate)	8,192
Current / voltage data rate	256
Energy data rate	1

Following initialisation and setup, the firmware developed for the microprocessor can broadly be divided into two sections. The first of these is contained in an interrupt routine. This will trigger, and thus the code be executed, every four cycles of the master timing signal regardless of what else the processor is doing. The interrupt code is concerned with the analogue to digital converter and recording data from the transducers as this must occur at precise intervals. It is also responsible for the averaging involved in the oversampling process. The second section is a main loop which operates whenever the interrupt code is not in process. This section is less time critical and is concerned with interrogating the RTCC and relaying the data produced by the first section to the laptop. A buffer system is used to allow the two components to interact.

The buffer system is based on packets of data. Each packet has a header which identifies the type of data that it contains and the length of the packet. The interrupt portion of the code generates the packets and stores them in sequence in the buffer. When a clock message is required, it inserts an empty clock message into the stream to be populated later. It does not matter if the clock messages are a fraction of a second inaccurate, just that they are in sequence. A similar system involving packets is used for the main transmission to the laptop. This stream, however, also includes framing messages inserted once per second. These are specific messages which could not otherwise exist in the data stream and can therefore be used to identify specific points in the data and hence align the data extraction.

The two sections of code for the interrupt and main loop routines are summarised in the flow diagrams in Figure 3.3 and Figure 3.4 respectively.
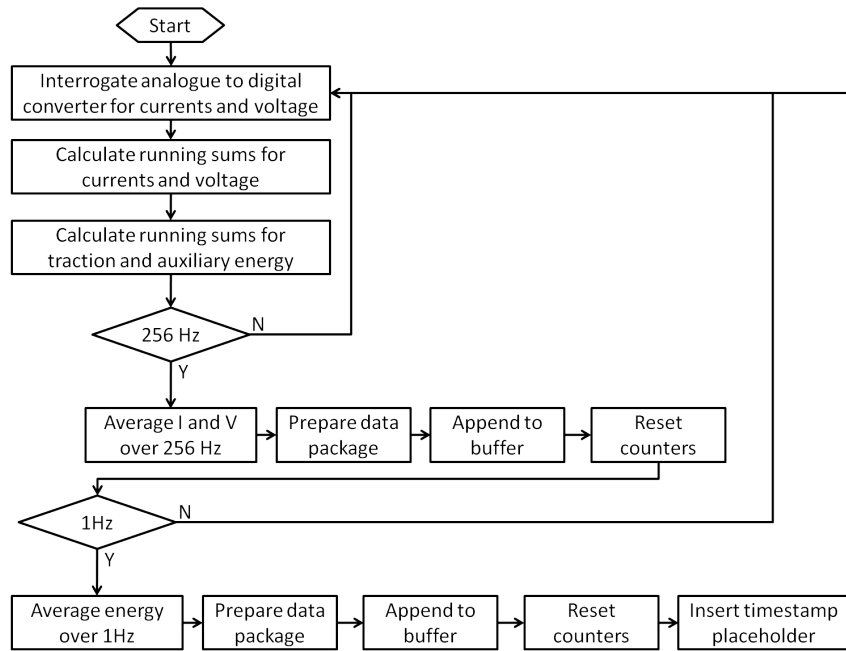
Figure 3.3: Microprocessor interrupt routine
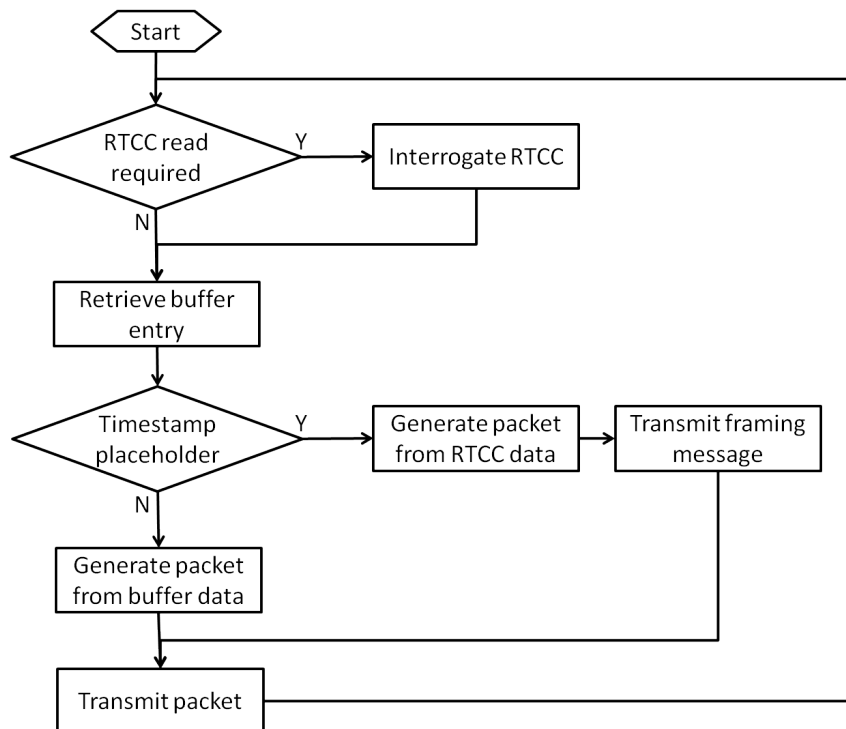
Figure 3.4: Microprocessor main loop

3.4.5 Software

The data generated by the instrumentation was transmitted to a laptop located in the vehicle body. This laptop had two functions, the first was to store the data for further analysis, and the second was to display a subset of the information in real time.

The software running on the laptop was written in C#. The data was supplied to it using a USB serial interface. The software extracted the serial stream into a stream of raw bytes held in local memory. This stream was then stored on the laptop's hard disk as a backup. The software also analysed the byte stream and used the framing packets to align it such that it could extract the data packets and build a second structure from them. This structure was then interrogated, using the header information, to extract the data from each transducer, as well as the pre-computed energy values, alongside the relevant time stamp information. Due to the moderately low data rate and simple structure, it was possible to record these expanded values to a text file ready for processing. The same data structure was interrogated once a second, by another thread. This thread was responsible for the graphical interface to the program, providing a real time display of the results that could be viewed during the trials.

The post processing was not done on the vehicle but instead using more powerful computers back in the laboratory where larger sections of the dataset could be loaded into memory. Energy consumption was calculated for each section of the trial route from the individual transducer signals, and verified against that calculated on the vehicle. The post processing analysis was carried out using MATLAB.

3.4.6 Calibration

Laboratory based calibration of sensors measuring up to kV or kA is difficult unless you have sufficiently large and calibrated sources. The solution used in the development of this instru-

mentation system was to select a number of large, but not huge, sources and then use them in calibration processes. Initially, a number of large (100-300 V) voltage sources were verified using calibrated commercially available metering equipment. Beyond the range of the largest unit, these were combined to generate larger voltages. A calibration curve was drawn between the measured voltages and the recorded results. The equation of this curve was then used in the processing of the recorded data. A similar process was undertaken for current, but instead of combining multiple sources, the same source was used though multiple turns of wire wrapped through the Hall Effect transducers. This has the effect of multiplying the current supplied by the unit, and so an extremely precise source measurement is required.

3.4.7 Mechanical construction

The instrumentation developed to monitor the energy consumption on a Class 508 railway vehicle was designed to fit into the equipment racks underneath the vehicle body. It was not exposed to the outside of the vehicle and was mounted after both the primary and secondary suspensions. This reduced many of the complexities of the mechanical construction.

Safety calculations

As with all instrumentation, in developing this system three main aspects of construction were considered:

1. Electrical

2. Structural

3. Fire

A Class 508 is a comparably aged vehicle with most of its control systems operating using 110 V logic. Although the main power supply for the system was 110 V, this was quickly converted to 24 V and then to the voltages required by the individual components of the instrumentation. The majority of interference which causes concern on railway vehicles is at high frequency. The only high frequency component in the system was the microprocessor which operated at 3.3 V and so was unlikely to affect the control systems. Nevertheless, the system was designed with the current standards for electromagnetic compatibility, BS EN 50155 [88] and BS EN 50121-3-2 [89], in mind. Furthermore, all of the instrumentation, with the exception of the transducers, was encased in metallic boxes for additional shielding.

The railway group standard GM/RT2100, structural requirements for railway vehicles, specifies the loads that equipment mounted to various parts of a vehicle body must be able to sustain. The Class 508 is a passenger vehicle and does not qualify as a rigidly coupled rake. Therefore, any equipment mounted in or to the vehicle body must be able to sustain loads of up to five times the force of gravity. The mounting for the energy metering instrumentation made use of six M8 bolts. Even in shear, these would be more than adequate for the load of the energy metering equipment which could be demonstrated to be fixed using a single bolt. The loads for each axis, which can be found in table Table 3.4, indicate that even a combined load in multiple directions would be sustainable.

Table 3.4: Minimum sustainable loads for equipment mounted on the vehicle body

Axis	Load
Longitudinal	$\pm Mass(kg) \times 5 \times 9.807 ms^{-2}$
Transverse	$\pm Mass(kg) \times 1 \times 9.807 ms^{-2}$
Vertical	$\pm Mass(kg) \times (1+2) \times 9.807 ms^{-2}$

The most stringent standards that the equipment had to meet were those relating to fire. Approximately 6.5 miles of the Merseyrail network is in tunnels and in these environments, the fire standards are particularly strict. The master standard relating to fire is BS EN 6853 [95], of which category 1a is the most stringent. It was this standard that the instrumentation had to meet. To achieve this, the instrumentation was contained in metal enclosures where possible, any cabling was selected from a limited number of certified options, and all electronics were sprayed with non-flammable conformal coatings.

3.5   Trials

The Merseyrail network is divided into two lines, the Wirral and the Northern lines. The Train Maintenance Depot (TMD), at which the equipment was fitted, is located at Birkenhead North which is on the Wirral Line. The trial took place on the Wirral Line which is shown in Figure 3.5.

The trials were conducted on the 12^{th} November 2009 using vehicle 508131. The trial route encompassed almost all of the Wirral Line with the exceptions being the section from Hooton to Ellesmere Port, the part of the underground loop east of James Street, and the spur to New Brighton. The trial route can be broken down into five sections as indicated in Table 3.5.

Table 3.5: Trial sections

Start	End
Birkenhead North TMD	Hoylake
Hoylake	James Street
James Street	Chester
Chester	James Street
James Street	Birkenhead North TMD
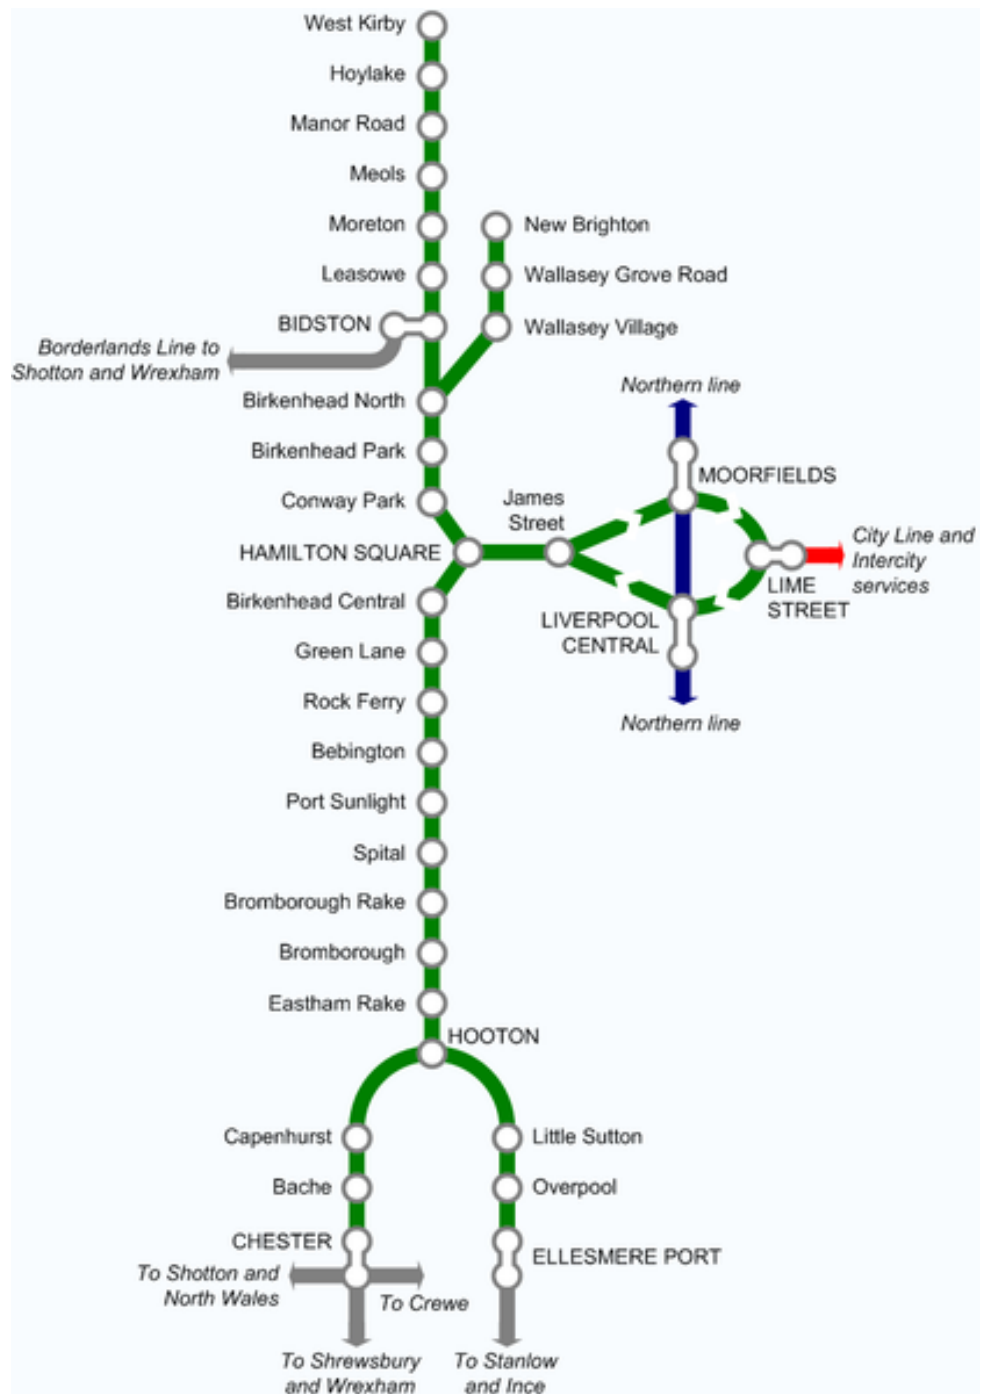
Figure 3.5: Map of the Merseyrail Wirral line

3.6 Results

The instrumentation used the high precision timing signal from the real time clock to trigger the acquisition of signals from the transducers. These signals were packaged with timing data and transmitted to a laptop computer for recording and processing. This processing was conducted using MATLAB. One element of the processing was to consider the data from the transducers in the spatial as well as the time domain. To do this, the data had to be converted from one domain to the other. A tacho signal can be used to establish distance travelled along the railway track. It is therefore the key signal required when converting from signals in the time to the spatial domain.

An additional tacho signal, originating on a different part of the vehicle, was obtained and recorded on the laptop at the same time as the signals from the transducers. The tacho signal consisted of an analogue voltage which was converted to a digital signal for recording.

The results presented in the remainder of this section are based on the data recorded. This data was taken at only one end of the vehicle and so may only represent half of that used by the vehicle for the whole journey. Any results presented based on the assumption that the energy usage can be doubled to represent both of the traction cars will have this process indicated.

3.6.1 Birkenhead North to Hoylake

The first section of the trial involved negotiating the exit of the Birkenhead North train maintenance depot and travelling along the most northern part of the Wirral Line towards Hoylake. Hoylake is not the end of that section of line but the penultimate station and the location of the last feeder substation along the route.

The energy use for this first part of the trial is presented in Figure 3.6. The energy use has

been plotted against distance, with the speed of the vehicle visible on the same figure. Following the initial departure from the depot, the vehicle increased speed consistently over a distance of approximately two kilometres eventually reaching a speed of just under 60 mph. At this point the traction system was disengaged and the vehicle continued to coast for almost a further kilometre, some braking may have been used, particularly at the end of this section, to bring the vehicle to a halt. A second, large acceleration then brought the vehicle speed up to 20 ms^{-1} before the traction system was again disengaged and the vehicle allowed to coast for the last two kilometres of the trial section.
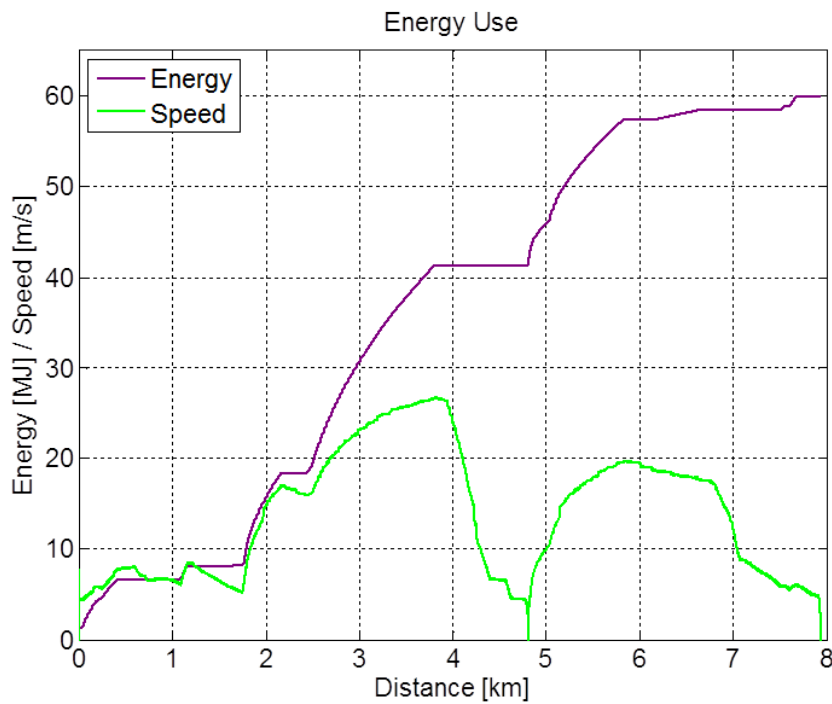
Figure 3.6: Energy use for the first trial section

The energy use displayed in Figure 3.6 is calculated from the power used by the vehicle which is in turn established from the current and voltage measurements. A plot displaying the power information is presented in Figure 3.7. The figure confirms that, once out of the depot, the 8 km trial section was travelled using two large powering phases and two coasting phases, one of them quite significant.
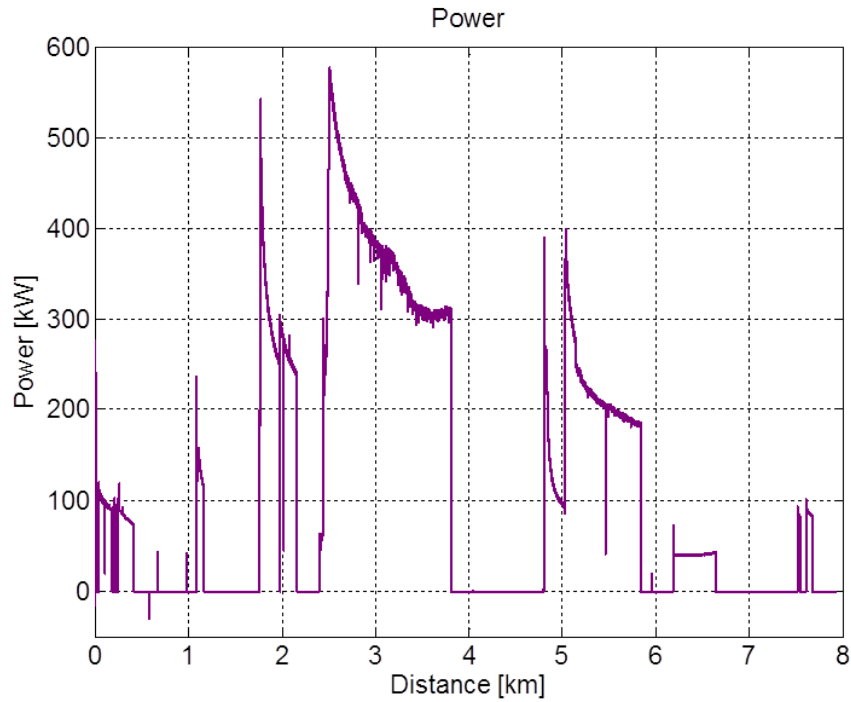
Figure 3.7: Power use for the first trial section

This data represents the first time that the vehicle had been operated following major maintenance work. It is likely that the approach of powering and coasting is used to test various systems on the vehicle during this initial testing phase.

3.6.2 Hoylake to James Street

The second section of the trial took the train back towards Birkenhead North and then on to James Street. This involved passing underneath the river Mersey as James Street is an underground station on the East side of the river.

The energy use and power results for this section of the trial are shown in Figure 3.8 and Figure 3.9 respectively. The results from the first trial section have been included to facilitate comparisons but it should be noted that, while over the same region of track, the runs are in opposite directions. The speed profile, again displayed in the energy figure, indicates that the

return journey was conducted at more consistent speed than the outwards journey, and in fact did not stop part way through despite slowing significantly in the same region. The smaller steps in the energy result, combined with the more frequent variation in speed, suggest that the return journey was undertaken using a slightly different driving style. Instead of large periods of tractive effort followed by periods of coasting, the second part of the trial uses finer control of the power to maintain a more constant speed. This is reflected in the power plot where, towards the end of the section, the vehicle uses far less peak power but generally uses power more regularly.
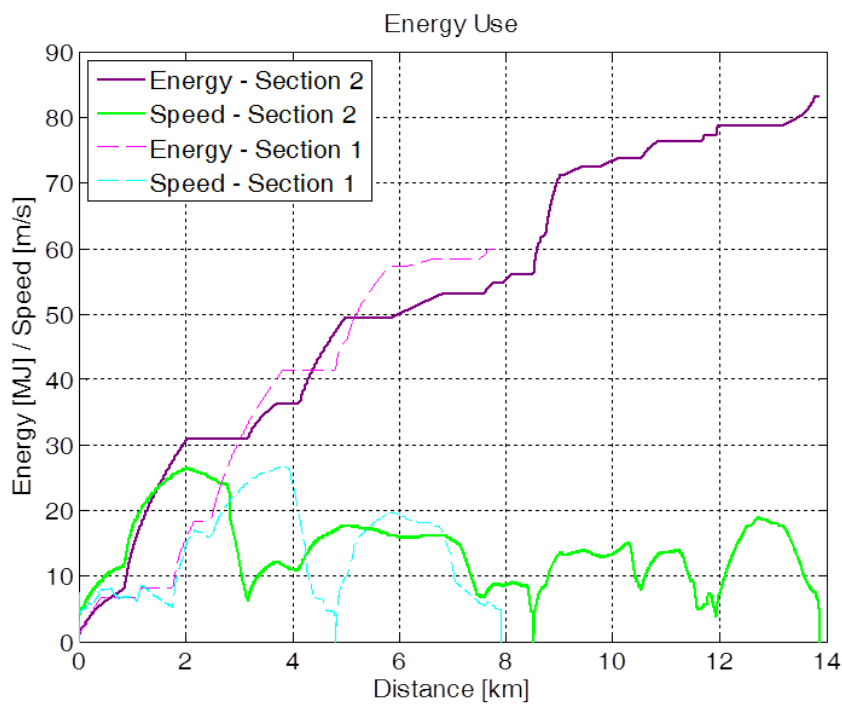
Figure 3.8: Energy use for the second trial section, first section overlaid

The first trial section was approximately 8 km long. The results for the second trial station indicate the vehicle coming to a stop around 8.5 km. This stationary period was at Birkenhead North station which is near to, but not precisely co-located with, the train maintenance depot. The acceleration away from the station is visible in the power graph.
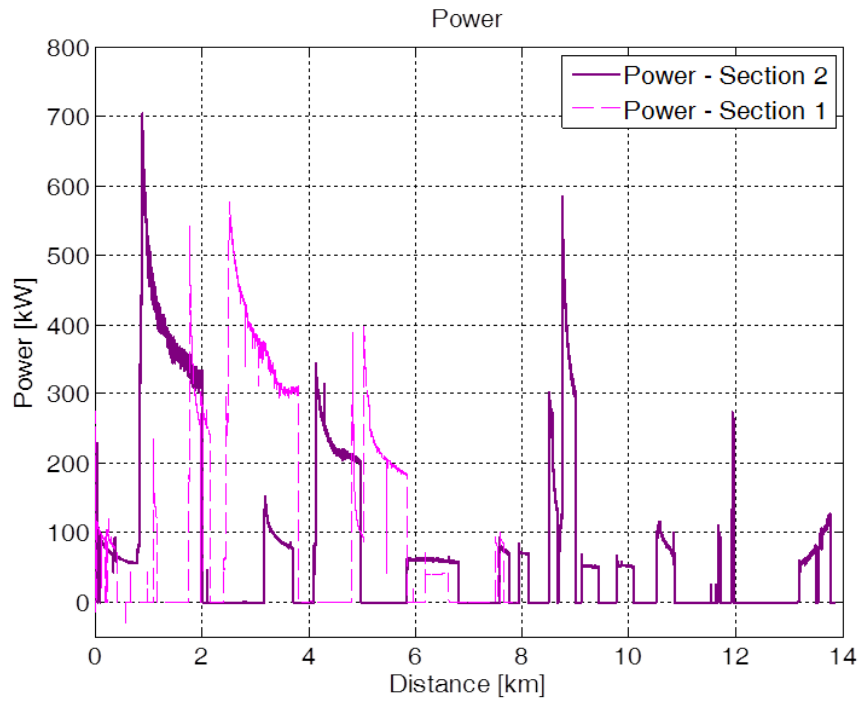
Figure 3.9: Power use for the second trial section, first section overlaid

From 12 km onwards, the vehicle appears to accelerate while consuming very little energy. This is confirmed in the power graph which shows no consumption between 12 and 13 km. This section of the data represents the portion of the trial where the vehicle passed under the river Mersey. The burst of power just before 12 km gives the vehicle enough momentum to enter the tunnel section where it accelerated due to gravity down the hill under the Mersey. The vehicle uses power again, after 13 km, to enter the station which is uphill from the lowest point of the tunnel but still under ground.

The different driving styles displayed in the first two datasets allow some comparison of the efficiencies of each style. A summary of the energy use for the first section of the trial and the corresponding return section in the second trial are presented in Table 3.6. The table contains average power information as well as energy use per kilometre and total energy use for the section.

Table 3.6: Comparison of trial sections one and two

Route	Average Power [kW]	Energy / km [MJ/km]	Total Energy [MJ]
Birkenhead North TMD - Hoylake	62.9	7.6	59.8
Hoylake - Birkenhead North Station	83.0	6.6	56.1

The energy figures in the table would suggest that a more consistent driving style requires less energy per kilometre than driving using bursts of power interspersed with periods spent coasting. This observation, however, does not include the fact that during the first trial section the vehicle stopped part way through the journey. This acceleration from stationary period would likely require greater energy use than simply maintaining momentum. Likewise, the average power result for the second trial section is likely exaggerated by the stationary period. Removing the time spent stationary from the (time based) average the figure changes to 69.1 kW. This is still below that of the second trial section and is likely to be due to the efficiencies of the traction system during its different phases.

This comparison is far from perfect. The two sets of data were taken in different directions over the same route without the gradient information being recorded. Also, the section of the line around the train maintenance depot and towards Birkenhead North station is not common. Suggestions for the differences are limited by the lack of understanding of the traction system. A more advanced system in which the workings of the traction system could be recorded in the same dataset would help to better explain this.

3.6.3 James Street to Chester

The third section of the trial again involved reversing direction to leave the underground section of the Merseyrail network without navigating around the "loop" under the centre of Liverpool. The section covered the journey from James Street to the most southern point on the Wirral Line in Chester.

The energy plot, again with speed overlaid on it, is shown in Figure 3.10. The very start of the plot suggests that the vehicle accelerates without using any energy. As with the section at the end of the second trial, this is related to the vehicle accelerating due to gravity on the steep section of track that passes underneath the river Mersey. After approximately 2 km, the vehicle used a disproportionately large amount of energy for the acceleration that it displayed. This can be attributed to climbing out of the tunnel section underneath the Mersey.
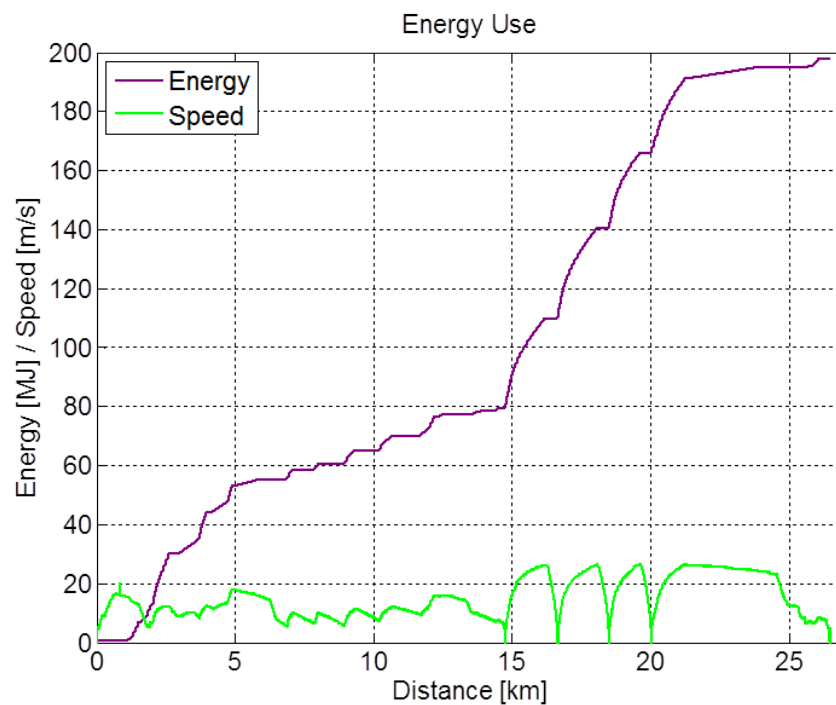
Figure 3.10: Energy use for the third trial section

Other than the section under the Mersey, the first 15 km of the figure suggest that the vehicle is being operated in a normal fashion using small injections of power to maintain a reasonably consistent speed. This is similar to the behaviour demonstrated in the second section of the trial, between Hoylake and James Street.

Between approximately 15 and 20 km, Merseyrail elected to perform a series of brake tests as part of the procedure to bring the vehicle back into service following the extended maintenance that it had undergone. These consisted of bringing the vehicle to an abrupt halt before powering up to approximately 60 mph to allow the brake test to be repeated. This full power / maximum braking cycle was repeated three times. Full braking on a Class 508 vehicle involves braking at 11.5% of the force of gravity.

The data collected during the brake tests allowed analysis of the traction system operating at near-maximum capacity. The traction portion of each brake test cycle was approximately 1.5 km in length. Each of these cycles used approximately 30.5 MJ of energy. This is a little over half of that required for the full 8 km journey from the train maintenance depot to Hoylake. The average energy use per kilometre during the 1.5 km acceleration phase was 20.3 MJ/km. This is nearly three times the average energy per kilometre used over the whole trial section which was 7.5 MJ/km. The average power during a single brake test cycle was 244 kW. Again this is approximately three times the level recorded for the full trial section (84 kW).

The region of the trial section, between 14 and 22 km, where the brake tests occurred is presented in Figure 3.11. Instead of showing the power or the energy used, the figure displays the traction current (blue) and line voltage (red) recorded by the transducers. The speed is also shown on the figure (green) although it has been multiplied by a factor of ten to make it easier to see. The four distinct repetitions correspond to the three brake tests and the acceleration back up to line speed following the final one.
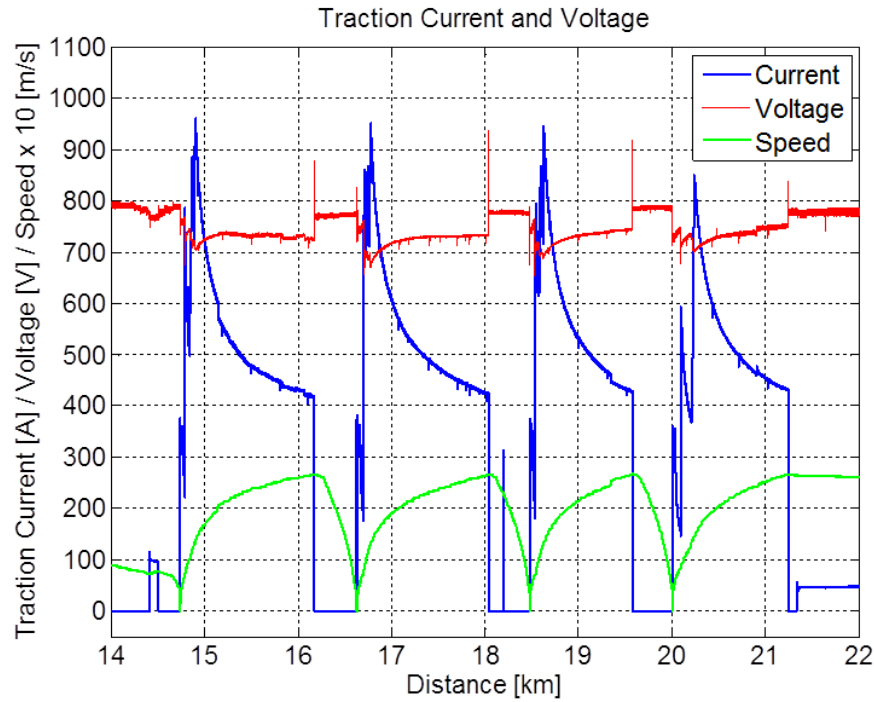
Figure 3.11: Traction system operation during brake testing

Each acceleration can be divided into a number of sections. Initially the current drawn rises sharply to approximately 360 A before reducing in steps to somewhere around 200 A. This corresponds to the traction system operating in series mode. Once the traction system has finished notching through its series options it switches into a parallel operating mode. At this point the current drawn rises dramatically and another sequence of notching occurs. The final notching in parallel mode instigates a field weakening process which sees the current rise again to as much as 950 A before gradually reducing to a little over 400 A. At this point the power is disengaged so that the brake test can occur. In the final acceleration phase, the current drawn drops to zero and then resumes at a much lower level to maintain the speed.

The red line on Figure 3.11 represents the voltage between the power pickup equipment and the vehicle earth point. Assuming the pickup equipment to be in contact with the conductor rail, this then represents the line voltage. The sections of the figure between and after the

acceleration sections indicate a standard line voltage level in that area of approximately 780 V. This is reasonably typical of that seen throughout all five sections of the trial. When the vehicle draws its maximum current load, the line voltage can be seen to drop to around or even just less than 700 V. When the traction system disengages, as demonstrated immediately after the 18 km point, the line voltage can be seen to spike. This is related to the inductive nature of the system.

3.6.4   Chester to James Street

The fourth section of the trial took also took place between James Street and Chester, but in the opposite direction. Similar to the reversal at Hoylake, the energy consumption between the two stations should therefore be comparable. As with the trials between Birkenhead North and Hoylake, no gradient information was recorded and so the comparison is slightly flawed. This effect will be particularly noticeable as the routes involve the tunnel under the river Mersey and because James Street station is underground. The two trial sections also differ because of the inclusion of brake tests on the way to Chester, however, these are somewhat offset by the inclusion of line voltage tests existing the station.

The energy results from the second trial section between Chester and James Street are shown in Figure 3.12. The line voltage tests essentially consisted of hard accelerations to draw large currents and therefore to test the resilience of the line voltage. These can be seen in the first 5 km of the figure. A station stop and the subsequent acceleration from standing occurred at approximately 13 km. The total energy used during this trial section was 169 MJ. This is approximately 15% lower than the 198 MJ used in the other direction. The difference is likely due to there being three brake tests and only two line voltage tests. Each brake test uses approximately 30 MJ of energy.
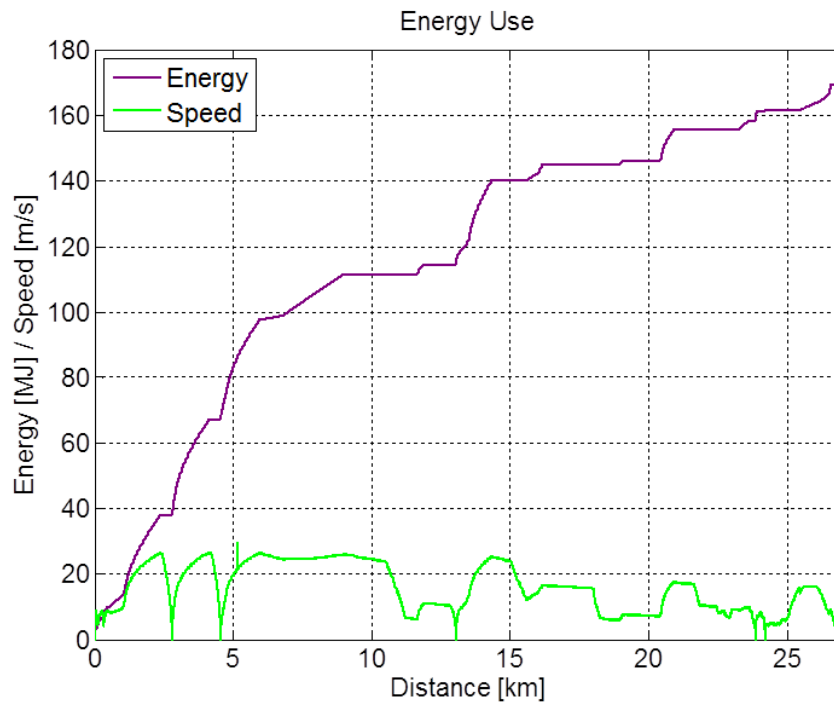
Figure 3.12: Energy use for the fourth trial section

Upon exiting Chester, Merseyrail took the opportunity of having instrumentation fitted to the conductor rail system to test the line voltage under the load of a vehicle. This involved the train using full acceleration to draw as much current as possible while watching to see what voltage drop occurred. Using the full accelerating power of the vehicle maximised the current draw however, in regular service two three car trains may be combined to form a six car unit which would potentially be even more demanding.

The two line voltage tests and subsequent acceleration to line speed are shown in Figure 3.13. The figure shows the current as well as the line voltage and the speed. The speed has been multiplied by a factor of ten to put it on a comparable scale. The largest voltage drop visible is approximately 100 V to 650 V. Although this is substantial, it is still within the specifications of the third rail system. The voltage drop seems to reduce with the second and third accelerations; these are closer to the Mollington feeder station that supplies the power to the conductor rail in

the Chester area. The line voltage tests were carried out as far from this substation as possible, the limitation was the driver's sighting distance to the signals coming out of Chester station.
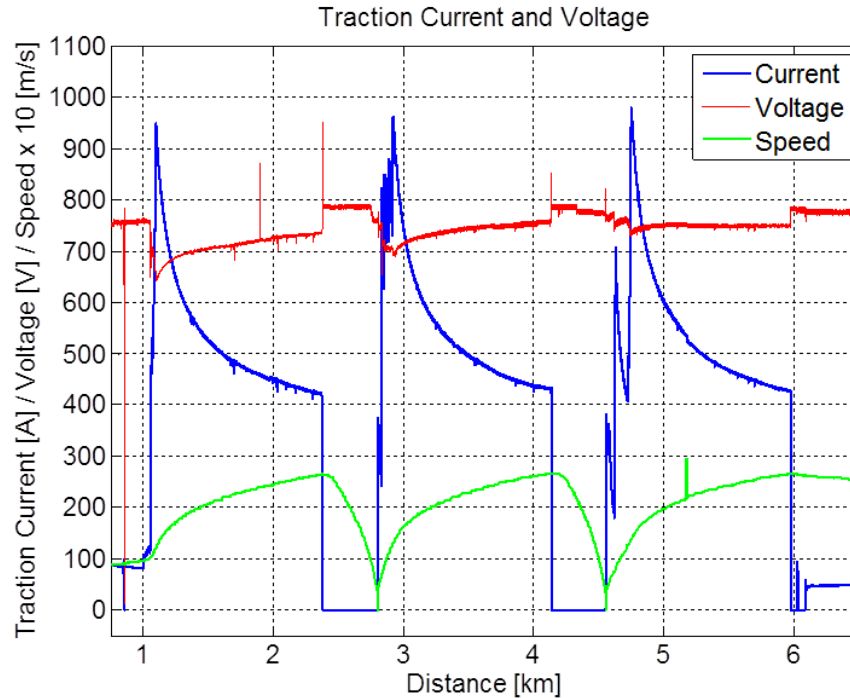
Figure 3.13: Traction system operation during line voltage testing

The inclusion of the line voltage tests on the Chester to James Street trial section allows an interesting comparison of three different styles of acceleration and the energy that they use. Exiting Chester station the vehicle was performing line voltage tests and so was accelerating at maximum capacity. The railway was not closed for the trials and so the vehicle had to fit in around other traffic. This involved a station stop approximately 13 km from Chester. The acceleration from this stop was typical of a regular station departure. For similar reasons, the vehicle had to queue to enter the area surrounding James Street station. This was under the river Mersey and so acceleration into James Street station was assisted by gravity in the tunnel section.

Although these effects can be seen in Figure 3.12 they are more clearly apparent in Figure 3.14

which shows the same data plotted against time rather than distance. The stationary periods are more apparent and can be seen at the start of the figure, 1000 seconds and 2200 seconds. The voltage testing acceleration appears to use as much as 38 MJ in the first acceleration and then 30 MJ subsequently. The standard station departure uses approximately 25 MJ, while the gravity assisted acceleration appears to use almost no energy.
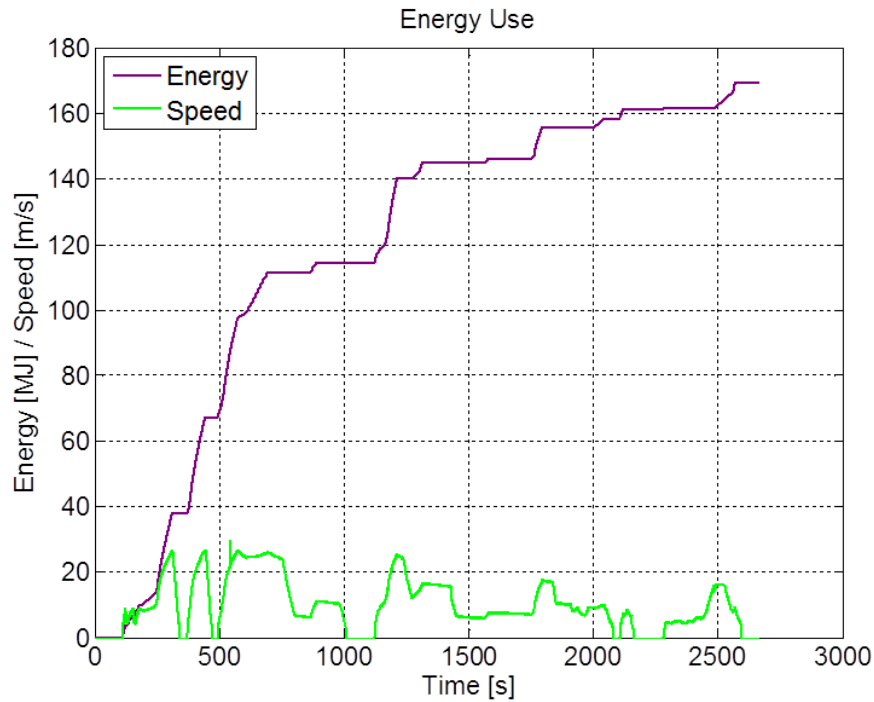
Figure 3.14: Energy use for the fourth trial section, plotted against time

The gravity assisted acceleration can be seen more closely in Figure 3.15. In this figure, the speed has been shifted by 140 ms^{-1} to improve visibility. Initially the train needs to use a small amount of energy, as little as a few hundred kilojoules, to start moving. The traction system is then disengaged and the vehicle allowed to coast for approximately 150 seconds on a gradient that nether increases or decreases the vehicle speed. At 2450 seconds, the gradient increases and the vehicle rapidly picks up speed. When the gradient reduces again, the traction system is activated to maintain the speed before being deactivated allowing the vehicle to come to a halt.
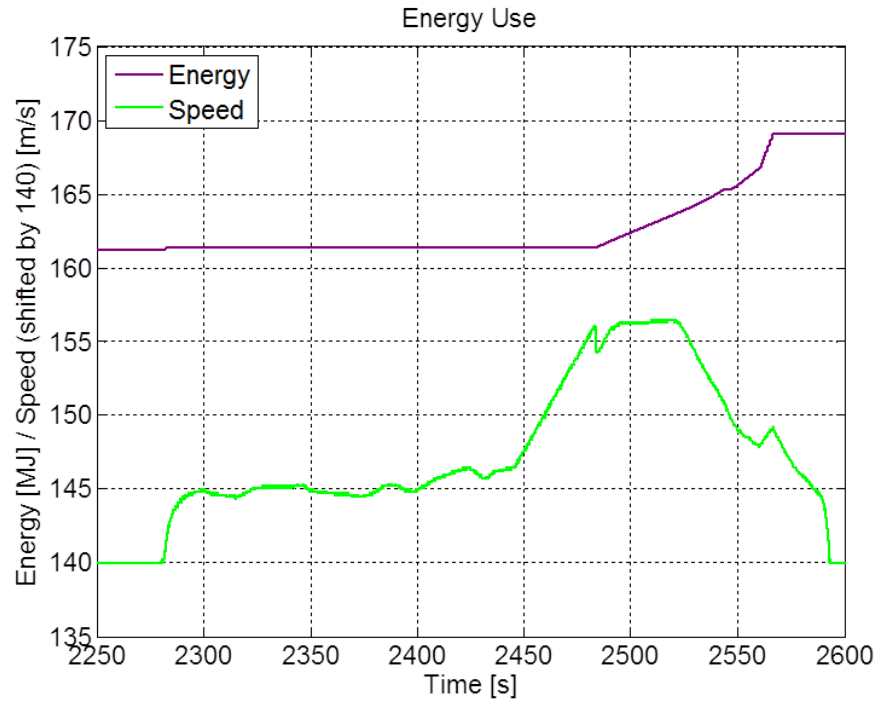
Figure 3.15: Gravity assisted acceleration

3.6.5 Comparison of trial sections

The previous sections have focused on specific points of interest, searching within large quantities of data. It is also possible to reduce the dataset to provide summary information for each section of the trials and thereby allow some comparison between them.

The data presented in Table 3.7 was obtained using simple filters to allow consideration over each second of the main datasets. This reduces spikes in the data but maintains sudden loading such as the inductive spikes that occur when the traction system is disengaged. The minimum operating line voltage is identified by excluding the data from gapped sections in the conductor rail, where the voltage falls away to zero.

Table 3.7: Voltages observed during vehicle operation in each trial section

Trial Section		Line Voltage [V]	
Start	**End**	**Minimum**	**Peak**
Birkenhead North TMD	Hoylake	712	785
Hoylake	James Street	710	788
James Street	Chester	676	798
Chester	James Street	609	794
James Street	Birkenhead North TMD	702	793

In all cases, the recorded values for the maximum and minimum voltages are within the tolerances defined in the conductor rail specification. The three trial sections that did not involve drawing power from the Mollington feeder station (near Chester) vary from the nominal 750 V line voltage by an average section maximum of 43 V. In all cases, the variation from the nominal is larger in the minimum than the peak values. The results from the section including the Mollington feeder station display peak values broadly comparable with the other sections but minimum values two or three times more separated from the nominal than the other sections. Despite being within the specification, this supports Merseyrail's theory that voltage supply is comparably poor in this area.

Similar processing can be applied to the current signals to extract the data presented in Table 3.8. Once again, the data has been considered over one second intervals to provide some smoothing without being destructive. The "average traction current during motion" is a time based average with the included data selected based on the tacho signal. It includes periods of coasting but removes data where the vehicle was stationary. This is necessary as a number of checks were made on the traction system at the exit of the train maintenance depot in the first data set. The "average traction current during traction" relates to data selected only when current was

being drawn by the traction system. In every case, the peak current occurred as the traction system switched into field weakening mode during acceleration.

Table 3.8: Currents observed during vehicle operation in each trial section

Trial Section		Traction Current [A]		
Start	**End**	**Average (Motion)**	**Average (Traction)**	**Peak**
Birkenhead North TMD	Hoylake	100	242	755
Hoylake	James Street	95	184	953
James Street	Chester	124	250	945
Chester	James Street	105	268	959
James Street	Birkenhead North TMD	95	189	879

The peak traction currents from trial sections two, three and four suggest an average peak current of 952 A during hard acceleration. The results from trial sections one and five can probably be discounted from this average as section one was the first testing after the maintenance work and section five was essentially manoeuvring back into the depot. Neither of these involved any hard acceleration. The lowest peak current draw during heavy acceleration occurs in the trial section containing the brake tests. With the exception of the line voltage tests, this was expected to be the heaviest acceleration encountered during the trials. The brake tests were conducted on flat ground; it is possible that the greater gradients encountered in trial section two may have had a significant contribution.

The brake tests which occurred during the third trial section had a marked effect on the average traction current when the vehicle was in motion. The usual braking pattern means that the vehicle travels much further when coming to a halt than under the 11.5% of gravity retardation involved in a brake test. The harder braking, and thus shorter coasting distances meant that

the average traction current was far above the averages recorded in the other trials. The total variation in average traction current during motion was 29 A over all five trials, this dropped to 10 A with the exclusion of the trial involving the brake tests. The results from the other trials were all reasonably consistent despite the extended coasting periods exhibited during the first section of the trials.

The use of coasting in the first section of the trials is emphasised by the average traction currents exhibited when the traction system was engaged. The 242 A average is much higher than that of the second trial section, which was driven without as much coasting, despite the two having similar averages during periods where the vehicle was in motion. The driving style presented in the second trial section did not involve artificially large coasting periods or unusually hard acceleration and so is presumably the most typical of a typical duty cycle. This is encouraging as it has the lowest average current demand.

The energy used by the traction system for each of the five trial sections is presented in Table 3.9. The total energy used is of interest but of limited value due to the differing lengths of the trial sections. Approximate lengths are presented to facilitate comparison. Considering the average energy use per kilometre (with accurate section lengths) allows a more direct comparison of the trial results. In addition to the results presented in the table, a separate calculation reveals the average energy use per kilometre between Hoylake and Birkenhead North station to be 6.6 MJ/km. The stability of the voltage creates an almost linear relationship between the power and traction current results. The average power recorded over all five trial sections was 60.6 kW.

The total energy use for all 77 miles of the trials was 540 MJ with an average energy use per kilometre throughout all five trial sections of 7.0 MJ. The values presented in the table for all of the individual trial sections were within 1 MJ of this value. The section most different from the average was the trial between Hoylake and James Street which was thought to be the most

typical of a regular driving style. While the driving style may be conventional, the stopping pattern wasn't and this may explain the lower values. The energy consumption between the two longer runs differed by approximately one hard acceleration phase and this gives a variation of 0.7 MJ/km even over that longer distance. This highlights the significance, in energy use, of starting and stopping a railway vehicle. The final trial section involved a lot of starting and stopping as the vehicle was returned to the train maintenance depot. Despite the average speed being low, the high energy requirements of the early phases of the acceleration result in this trial section having a high average energy use.

Table 3.9: Traction system performance observed during each trial section

Trial Section		Length	Average Power	Traction Energy	
Start	End	[km]	[kW]	Average [MJ / km	Total MJ
Birkenhead North TMD	Hoylake	8	38.2	7.6	59.8
Hoylake	James Street	9	62.9	6.0	83.1
James Street	Chester	26	83.2	7.5	197.6
Chester	James Street	26	63.5	6.3	169.0
James Street	Birkenhead North TMD	8	51.1	7.7	58.1

The results obtained from the auxiliary current transducer revealed the current consumption of the auxiliary systems to be reasonably deterministic. The Class 508 is a reasonably simple vehicle by modern standards and the requirements of the auxiliary system are limited. The cab heater was shown to draw 5 A while the heaters in the vehicle body and the centre trailer draw 30 A each. The peak current draw identified across all of the trials was 74 A. This occurred during the Hoylake to James Street section.

Using a nominal value for the line voltage, the maximum detected current drawn would equate

to a power use of 55.5 kW. The "A" car supplies the power to the heating in the trailer car, and so this figure corresponds to two thirds of the train rather than half of it as is the case with the traction current figures. Also, the reduced stopping pattern has been shown to have a huge impact on the power requirements of the traction system. Nonetheless, it is still interesting to consider the 55.5 kW used by the traction system in relation to the average traction system usage demonstrated to be 60.6 kW.

3.7 Traction Profile

Although the standards for energy metering only require industrial energy metering systems to record data at very low rates, the instrumentation developed for this project samples and stores data at far higher frequencies. This allows close inspection of the data which is of particular interest when analysing the traction system.
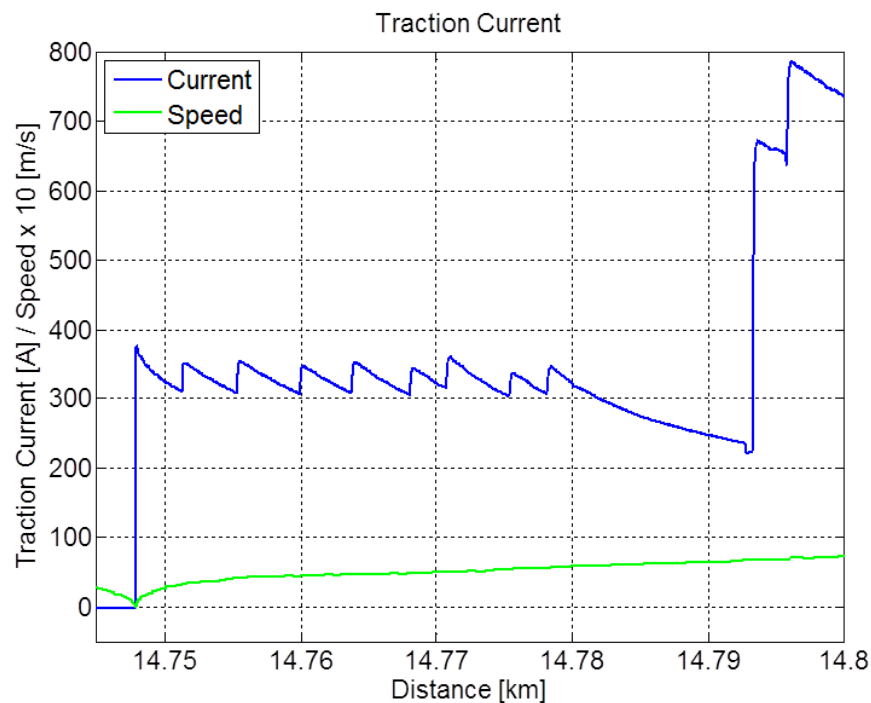
Figure 3.16: Traction system notches visible in current waveform

A close view of the data from the traction current transducer is shown in Figure 3.16. The figure shows part of the tractive effort waveform recorded during one of the brake tests. The waveform clearly shows the notches as the traction system switches between modes. Further studies might allow this data to be aligned with the operation of the traction system.

3.7.1 Results over five minutes

The standards for energy metering state that results for energy use should be recorded every five minutes. The results presented above involve data points being recorded 256 times per second. It is possible to reduce the recorded dataset to mimic what would be recorded had the instrumentation adhered to the standards. The result of this process being applied to the traction energy use is presented in Table 3.10. Where a dataset is not a precise multiple of five minutes long it is not included in the table, but is summarised in Table 3.11.

Table 3.10: Energy use presented in five minute intervals

Time	Energy Use by Trial Section[MJ]				
[Seconds]	Section 1	Section 2	Section 3	Section 4	Section 5
1-300	0.03	36.26	12.77	34.19	16.31
301-600	0.02	18.44	41.87	64.34	36.01
601-900	9.11	21.38	5.92	15.65	3.38
901-1200	32.16	6.97	9.75	21.16	
1201-1500	18.51		9.33	9.43	
1501-1800			60.73	10.91	
1801-2100			53.68	2.44	
2101-2400				3.29	

From the data presented in the table, it is possible to identify when the vehicle has spent time

Table 3.11: Energy use for incomplete five minute sections

	Trial Section				
	Section 1	Section 2	Section 3	Section 4	Section 5
Duration [s]	65	38	276	263	236
Energy [MJ]	0.00	0.00	3.56	7.63	2.36

stationary and when a five minute period has involved heavy acceleration. For instance, the 300 and 1500 second brackets of trial section 3 involve traction while the 600, 900 and 1200 second brackets were mainly spent at approximately constant speed. Beyond this, it is difficult to tell much about the behaviour of the train. There is sufficient information, however, for billing processes to be based on energy use. For comparison purposes, Figure 3.17 shows the energy use plot against time for the trial section between Chester and James Street with the five minute version of the data shown in Table 3.10 overlaid onto it. While the general shape is the same, the detail is clearly missing.
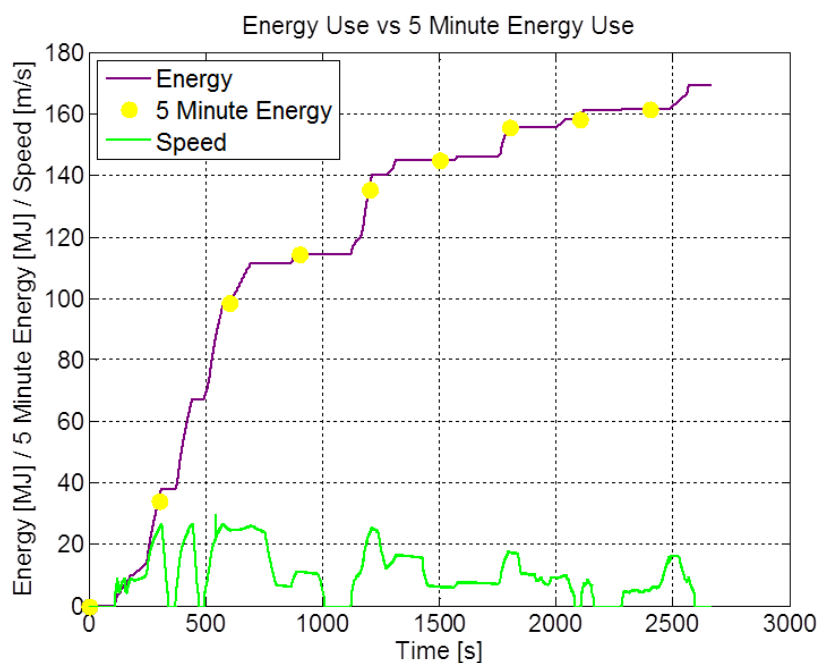
Figure 3.17: Energy use during the fourth trial section with overlaid data at five minute intervals

3.8 Conclusions

This case study has presented work involving the instrumentation of a Class 508 railway vehicle to monitor energy usage during a single day of trials. The trials took place as part of a standard re-conditioning exercise and therefore included specific system functionality tests, such as braking and high traction, as well as periods of general operation. This unusually wide range of vehicle operation has allowed analysis to be focused on some of the subsystems of the vehicle, and the interactions between the vehicle and the railway system as a whole. This kind of investigation is not usually possible due to the difficulties involved in scheduling the appropriate track possessions and / or vehicle activity.

The instrumentation developed for and described within this case study made use of a minimal set of transducers, as specified by the project remit. The system was based on a structure of a single transducer interface module servicing a number of transducers and communicating with a central data acquisition unit. The case study has shown that through non-invasive modifications, such as increasing the sampling rate, it has been possible to generate a larger dataset and then extract a greater level of information through post processing. Much of the analysis has made use of an additional tacho signal to allow reference to be made to speed and to divide the data into sections. Even with this minimal sensor set, the integration of the tacho and energy data has been essential in the extraction of information. Without the tacho, much of the energy data would be without context other than the notes made during the trials attempting to relate output files to the physical world.

The nature of the trial runs undertaken by the vehicle during the data acquisition phase of this work have lead to a greater understanding of the different strategies available for accelerating a Class 508 railway vehicle from a stationary position. Overcoming inertial in this way is a huge component of energy consumption within the operation of a railway vehicle, and so this has a

huge effect on the energy efficiencies of different driving styles. Despite the lack of a standard stopping pattern, a significant level of driving style analysis has therefore been undertaken and fed back to both the DfT and Merseyrail. Further analysis would require larger datasets and different (more automated) analysis techniques to cope with the increased volume of data. Initial investigations also suggest that the position within the network, possibly in relation to vehicle loading and track gradient, can have a significant effect on energy consumption. In order to verify this hypothesis, the instrumentation would need to be expanded to record additional parameters relating to the operating conditions of the vehicle. In this context, an alternative to the PC based data acquisition component of the system would be preferable from a reliability perspective.

The case study has also shown that it is possible to use an instrumented railway vehicle as a tool of measurement for the network infrastructure. Results have been presented relating to the line voltage and particularly its resilience to loading around the Chester area. This was investigated at the request of the train operating company. The investigations have shown that the supply in that area is somewhat deficient in comparison to other supplies on the network. However, they have also shown that it remains within the specification of the system when loaded as heavily as possible by a three car train.

Although not part of the original project specification or instrumentation design, the increased sampling rates have also allowed an analysis of the behaviour of the traction system. The operation as it switches through modes and between series and parallel operation has been identified in the traction current waveforms. A particularly close look at these waveforms has revealed notching believed to be associated with the camshaft changing mode although further investigation is required to confirm this before it could be used as a measure of such.

The instrumentation of the auxiliary systems has revealed a surprisingly extensive maximum

power requirement. The energy metering systems mandated by government and now being retrofitted to all railway vehicles do not differentiate the traction and auxiliary systems and so could not demonstrate this potential efficiency saving to the TOCs. While the auxiliary usage seems large when referenced to the traction requirement, it is worth remembering that the traction requirements shown are based on the very limited stopping patterns demonstrated in the trials and the difference would be much greater in regular service.

In addition to the traction / auxiliary power usage, a brief comparison has been made on the traction side between the energy monitoring data recorded by the instrumentation systems developed for this academic study and those required for the metering systems mandated by government. The difference in the levels of precision found in the two datasets is sufficiently significant that almost none of the analysis performed on the academic dataset could have occurred using the mandated one. It is acknowledged, however, that the data requirements of the mandated system are very different from the academic ones and recording with this level of precision in a native railway context would lead to huge data management difficulties for the industry.

Despite efforts to consider the data in ways other than those originally mandated, limitations arising from both the instrumentation architecture and scope, and the programme of field trials have been identified. Extensions to the work presented in this case study would initially not require changes to the instrumentation, and would instead be focused on repeating the trials with more realistic stopping patterns. The instrumentation system could then be extended to record supporting data, such as the gradient information and possibly more information about the traction system. This data would be used in order to contextualise some of the results and verify the operation of the traction system. Such an expansion of the instrumentation would, however, likely prove incompatible with the simple architecture used in this project where sensor

outputs are directly fed back to a single data acquisition point. Multiple passes over the same sections of track would also be of benefit as it would allow features of the track and of the traction system to be separated.

CHAPTER 4

CASE STUDY 2 - CONDUCTOR SHOE MONITORING

4.1   Introduction

This chapter presents a second case study in vehicle-mounted instrumentation. The area of interest was the interface between electrified third rail and the conductor shoes which railway vehicles use to draw power from it. Apart from the obvious challenge of instrumenting something operating at 750 V DC, the instrumentation had to operate in the extremely high vibration and shock environment of the bogie / third rail interface. Physical connections to the high voltage interface have been made for instrumentation purposes before, but only for displacement measurements and not to operate at line speeds as this system does.

The high level objectives for the project were based on the contact forces between the conductor rail and the shoegear. These were set jointly by the Infrastructure Manager and the Train Operating Company, but the implementation of the system was left to the author and the University of Birmingham. The instrumentation was therefore configured with additional transducers and to allow the maximum information to be extracted from the recorded data, more than was required for the original project. In this chapter it will be shown how these additional sensors allow a greater level of understanding and interpretation of the mandated measurements.

The architecture used in this work was based on the two-site data acquisitiuon / data recording system emplyed in the first case study. This was largely due to the proximity of the instrumentation to the high voltage conductor rail system. It was also selected based on the sensors being clustered in one location, and the requirement for a live visualisation of the data by an operator remote from that location.

4.2   Third Rail and Conductor Shoe Systems

The very first vehicles to use electric traction relied on carrying their own batteries. These demonstration models, such as Robert Davidson's 1842 offering [96], were ultimately unsuccessful due to the battery technology of the day, but illustrated some of the potential for railway electrification. In 1851, an American named Hall demonstrated that it was not necessary for the vehicle to carry the power source by moving the batteries to a trackside position and using the running rails to supply the power to the vehicle. This system also failed due partly to the battery technology and partly to the efficiencies of power transmission, but again it was a valuable demonstrator of principle.

By the 1880s, Siemens, Edison and Daft had all demonstrated electric vehicles using lineside power supplies. The Daft system from 1883 was the first to use a dedicated third conductor rail. The locomotive, called Ampere, drew electricity from a central third rail to power a 22.3 kW motor [97].

The first British third rail system was constructed in 1890 by the City and South London Railway. It was a tube railway operating underneath the River Thames and worked on 450 V DC [98]. The Waterloo and City line followed in 1898, and Liverpool's Merseyrail link between Liverpool and Birkenhead in 1903.

For historical reasons, many third rail systems in the UK operate a "top running" procedure where the conductor shoe is pressed down onto the head of a third rail. Alternative systems such as "under running", where the shoe presses on the bottom of a suspended third rail, and "side running" where the contact is on the side, are popular in other parts of the world. These systems provide better protection from ice, debris and for track workers, however, the under running principle is particularly hard to implement around junctions and turnouts while maintaining sufficient contact area for the train. It is therefore not entirely suitable for the highly dense UK third rail network.

The London and South Western railway system, which operated out of London Bridge and Waterloo, evolved to become what is now known as the Southern Railway [97]. This is a 650/750 V DC system which makes up a large portion of the 4000 km of top running third rail network that Network Rail has in the UK today [99]. The majority of the network supports line speeds of up to 145 km/h, with the flagship Eurostar link supporting 160 km/h.

4.2.1 Top running, beam referenced, non retractable, Electrostar series III shoegear

Although a number of different vehicle types operate on the DC third rail network to the south of London, the majority of vehicles fall into the Electrostar family. The Electrostars cover the Class numbers, 375, 376, 377, 378 and 379. At the time of writing, the author has experience of fitting instrumentation to the shoegear found on both Class 375 and Class 377 vehicles.

The Class 375 and Class 377 vehicles instrumented both make use of Electrostar series III non retractable shoegear; specifically, parts M20070-04-L or M20070-05-L. This type of shoegear is produced by Brecknell Willis who also produce conductor rail and complete overhead / pantograph systems.

The series III non retractable shoegear, shown in Figure 4.1, consists of an iron shoe, a shoe

bracket, a fibreglass arm, a torsion spring, a rubber bump-stop, electrical cables and a frame. The shoegear system also makes use of an electrical arc shield, a laminated shoebeam, and a down stop bracket to suspend the system from this shoebeam which are not shown in the figure.

Figure 4.1: Series III non retractable shoegear

A new shoe is 25 mm thick and will be used until, on inspection, the thickness is found to be less than 14 mm. This is considerably lighter than the conductor shoes found on older Class 442 or Class 508 vehicles.

The frangible joint is designed to be the point that the conductor shoe separates from the vehicle in the event of a contact or a sufficiently large force being transmitted to the shoe. The frangible joints are designed to break at 25 kN. Anecdotally, lower intensity but longer term forces caused by impacts at lower speeds can occasionally lead to a tearing of the fibreglass arm and thus a less "clean" separation of the shoe.

The fibreglass arm is the primary insulator that separates the conductor shoe and rail from the train. It is a laminated structure with some limited flexibility. The torsion spring is configured

to ensure that the nominal contact force between the conductor shoe and the head of the third rail is 250 N with a tolerance of 50 N, i.e. 300 N maximum contact force [100]. It consists of a metal spring wound around a plastic sheath which encases a metal pivot rod.

The rubber bump stop ensures that the shoegear does not rise too high towards the arc shielding or the bogie of the vehicle. The down stop bracket ensures that the conductor shoe hangs from a laminated shoebeam, fitted between the axles of the vehicle, in the event that there is no conductor rail present to support it.

4.3   Motivation

Prior to the work presented in this case study, the author was involved in a series of laboratory based trials to consider the effect of ice formation on conductor rail, and the suitability of conductor shoe equipment for its removal. The suggestion from the sponsor was that simply increasing the down force would allow the conductor shoes to remove the ice. The trials did not support this hypothesis. One output from the trials was the realisation that little was known about the interface between the conductor shoe and the third rail. While standards do exist, they are only for the static condition and other than manufacturer's assurances, no procedures are in place to verify that the operation is within these standards beyond standard gauging practices during routine maintenance operations. At the end of the trials, many questions were being asked in relation to the operational contact forces and displacements both from Network Rail from an ice clearing perspective, and also from the Railway Safety and Standards Board who wanted to generate new standards for dynamic conductor rail / shoe interactions.

The work presented in this case study comes from a need to better understand the interface between the conductor shoe and the third rail and the dynamic performance of this interface. Of particular interest were the height of the conductor rail, in relation to the running rails, and the

contact force between the conductor shoe and the conductor rail. Both of these are described by standards documents.

The standards document RT/E/C/27010 [100] states that the height of the conductor rail, above the level of the running rails (Hight Above Rail Level) (h.a.r.l.) should be 76 mm with static tolerances of -4 mm (72 mm) and +10 mm (86 mm). The standard also includes dynamic tolerances of -20 mm and +40 mm, but these cover the majority of the shoegear's available travel and so are of limited interest in a study of regular shoe / rail interactions. A down-stop is used to set a lower limit on the shoe's vertical position in the event that it is not in contact with the conductor rail, for example at a gapped section. This limit is 53 mm.

The shoegear has been manufactured to provide a contact force of 250 N ±50 N as specified in [100]. Again these tolerances are for a static system. Dynamic results, for example at ramp end impacts, will be outside these values but are not significant in a study of general shoe / rail interactions.

Although both the height and force parameters are specified in standards, little is known about the actual operation of the systems or how well the standards are adhered to. Insights gained from a greater understanding of the shoegear / third rail interface will go forward into future shoegear design as well as further development of standards for third rail, similar to those already in existence for pantograph / overhead interactions [101].

4.4 Instrumentation

4.4.1 Overview

The standards that describe shoegear operation are written in terms of displacements relative to the running rail, and contact forces between the conductor shoe and rail. The technologies

selected for use in the instrumentation needed, therefore, to reflect this. Displacement transducers were selected for the distance measurements but it was not possible to fit a direct force transducer between the conductor rail and the shoe. Instead a system based on strain gauges was selected. Strain gauges were fitted to the fibreglass laminate arm located between the torsion spring and the conductor shoe. These were used to obtain a value for force, which could then be used in an estimation of the contact force. The instrumentation was therefore based on strain gauges and a displacement transducer.

The architecture selected for the development of the instrumentation system used in the Class 372 conductor shoe work was similar in nature to that from the previous case study. The instrumentation system required the acquisition of data from outside the vehicle, in reasonably close proximity to the energised conductor rail. This combined with the requirement for the strain gauge amplifiers to be close to the strain gauges prohibited a direct connection between the transducers and the recording system which had already been selected and was located in the vehicle body. The system was therefore divided into two parts. The first of these was a sensor management and data acquisition unit mounted on the bogie. This was based around an embedded microprocessor. The second was a media converter and data-recording unit to be located in the vehicle body. This was based on a laptop computer. The two units were linked fibre optically to minimise interference and electrical contact between the two regions of the vehicle. Given the split site architecture, it was decided to minimise connections between the bogie and the body and so the NCDT was also connected to the strain gauge management system despite being better suited to direct interrogation. This approach, however also improved the level of integration between the displacement and force data streams.

Figure 4.2 shows the standard configuration of a set of shoegear mounted to the bogie of a vehicle. The sections of the figure marked in red, have been added to the original drawings and

represent the additional instrumentation that was fitted. The same instrumentation is shown in plan view in Figure 4.3. The sections of instrumentation shown in the lower half of Figure 4.2 are strain gauges which were used to measure the strain in the fibreglass arm. The figure shows four gauges, but in fact gauges were located at eight positions on the arm, four being masked by the four that are visible. An alternative view showing the underside of the fibreglass arm and the layout of the lower four strain gauges is shown in Figure 4.5. The strain gauges were coated in a liquid rubber compound to give them physical and electrical protection while still retaining flexibility. They were connected to the main body of instrumentation by very fine enamel coated wire which was intended to vaporise in the course of a "third rail event", i.e. anything involving them becoming electrically connected to the third rail.

The instrumentation in the upper portion of Figure 4.2 consisted of a Non Contact Displacement Transducer mounted on the outside of a metallic box containing the acquisition and processing components of the instrumentation.
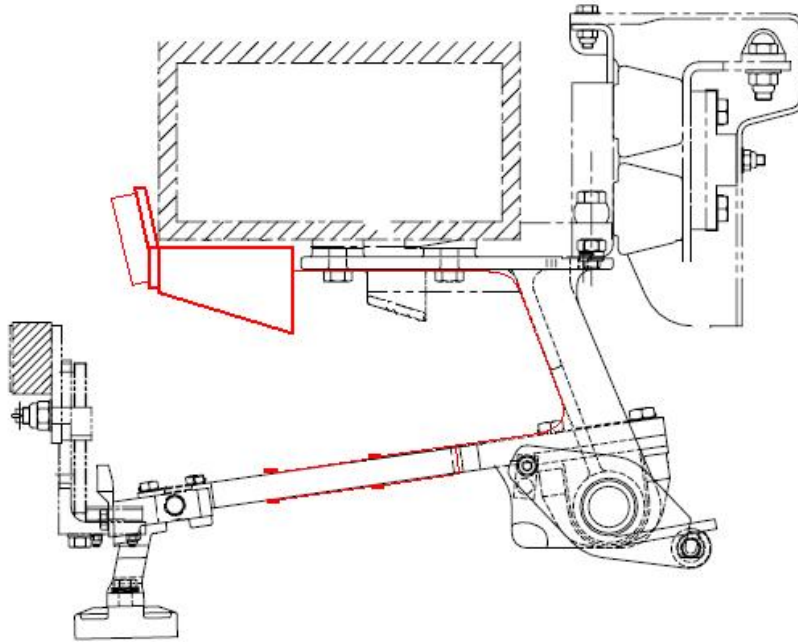
Figure 4.2: Side view of the series III conductor shoe with instrumentation indicated
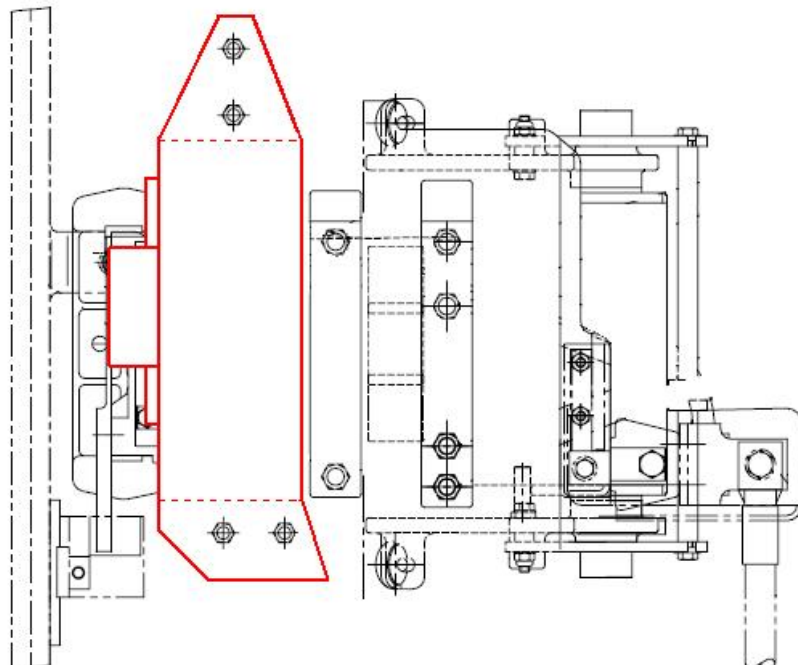
Figure 4.3: Top view of the series III conductor shoe with instrumentation indicated

The instrumentation is shown installed in Figure 4.4. The silver box in the centre of the figure houses the acquisition and processing systems and was linked to the recording equipment using the blue connections going off to the right of the figure. The blue box on the front of the instrumentation is the non-contact displacement transducer. This was angled to measure directly to the tip of the fibreglass arm. As the trial was only for a single day, it was not felt that the transducer needed any additional protection from the elements. The installation of the strain gauges can be seen in Figure 4.5.
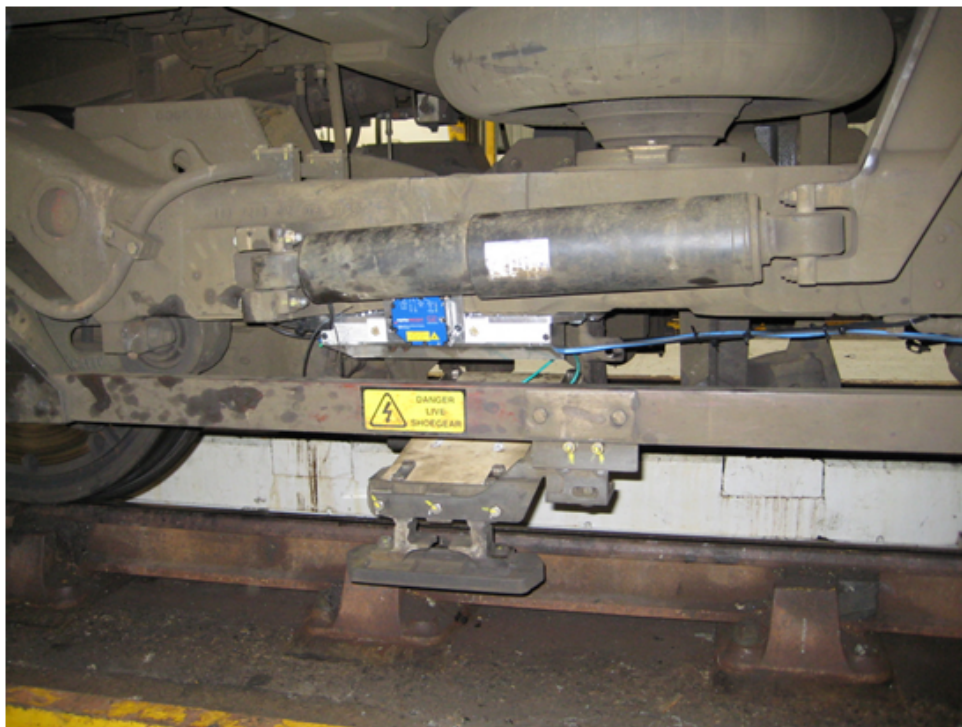
Figure 4.4: Installed instrumentation
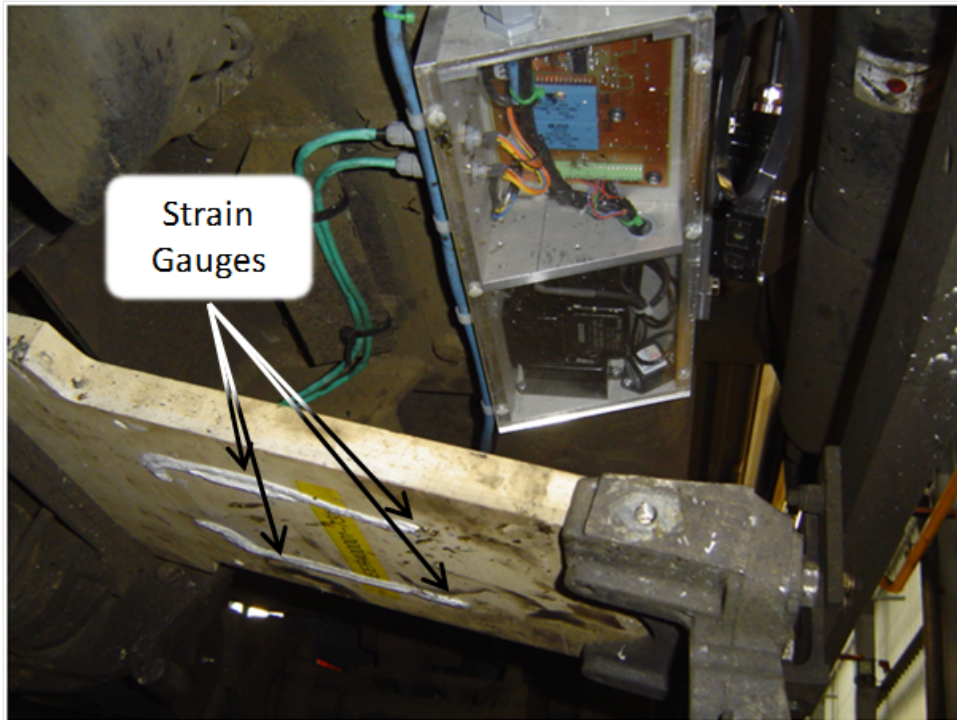
Figure 4.5: Close up view of installed instrumentation

4.4.2 Sensing technologies

The primary aims of the conductor shoe instrumentation work were to measure the height of the conductor rail relative to the running rail, and also the force between the conductor shoe and the third rail.

Previous instrumentation work has measured conductor shoe displacement using drawstring sensors [102]. These have been used previously by the University of Birmingham for fixed asset monitoring [103] but there was some concern over their robustness when mounted on a bogie travelling at line speed. Instead, for this project, a laser based non-contact displacement transducer was selected. This gives equivalent performance to the draw string sensor while not having any moving components which may be likely to fail when subject to accelerations of up to ± 20 g [104], as found on the bogie.

The non-contact displacement transducer was mounted on the main equipment box and directed as to measure the distance to the tip of the fibreglass arm. Using basic geometry, it was possible to convert the measured displacement, the dimensions of the shoe mounting bracket, and the thickness of the shoe (measured before and after the trials) into a conductor rail height relative to the bogie.

This measurement between the conductor rail and the bogie is based on static dimensions for the suspension system of the vehicle. It is therefore not technically correct when the vehicle is in motion because of dynamic effects which cause the suspension system to expand and contract. It would be possible, through the use of displacement transducers fitted between the axlebox and the bogie, to identify this additional displacement and incorporate it into the calculations. Rather than being used to contextualise the main displacement signal, this would constitute part of the processing chain required to obtain it.

Making this measurement, however, would add significant overhead to the instrumentation system. Considering previous work where displacement transducers have been used, [40], [105], it appears that for axlebox movements of approximately ±1.5 mm the associated bogie movement would be no more than approximately ±2-3 mm. These measurements were taken using a vehicle with a different suspension configuration, but suggest a region for the range that this additional displacement might occupy given standard bogie movement of ±4 mm on a Class 375 vehicle [106]. Given this limited assumed movement in comparison to the predicted conductor rail variation of 14 mm in the static case and a range only limited by the physical limitations of the shoegear arrangement in the dynamic case, it was decided that the movement of the primary suspension could be considered non-significant. In this case, the measurement of the conductor rail height relative to the bogie can directly be converted to a measurement relative to the running rails. The author accepts this as a limitation of the system brought about by the

tradeoff of precision against cost and complexity.

Measurement of the force between the conductor shoe and third rail is considerably more difficult. While it is possible to use pressure transducers to take a direct measurement under static conditions, it is impossible to use this technology when the vehicle is in motion. Instead, it is possible to generate an estimate for the contact force by combining the measured force with other forces within the system. The measured force is then combined with the weight of the shoe and shoe bracket to give an estimate of the contact force while the shoe is (vertically) stationary. This approach takes measurements in a directly coupled system to the conductor shoe, but not of the interface itself. It is therefore not perfect. It does not include the forces required to (vertically) accelerate the mass of the shoe. This means that the estimate of contact force becomes less accurate with increasing shoe accelerations but is reasonable when the acceleration is small. In a practical context, this means that the estimate is adequate for monitoring the general condition of the interface, but the outputs from the system are not accurate when the shoe is subjected to shock inputs such as those found at ramps or joins.

An arrangement of 8 strain gauges arrayed in 2 strain gauge bridges was used to obtain an a measurement of the force throughout the fibreglass arm. Two bridges were used for redundancy due to the extremely harsh mechanical environment. The gauges were aligned along the length of the fibreglass arm and wired such that they formed a leading and a trailing bridge while the vehicle was in motion. This is shown in Figure 4.6. The results from the two strain gauge bridges (A & B) were averaged in order to produce a value for the average strain the fibreglass arm. The strain gauge measurement system was calibrated by fixing one end of the arm at its normal pivot point (having removed it from the torsion spring) and applying known loads to the free end while recording the strain values. The results were used to create a calibration curve, the equation of which was then used in a software based calibration process. The results from

the strain gauge were combined with the calibration curve and the weight of the conductor shoe and its bracket to produce the estimate for contact force.
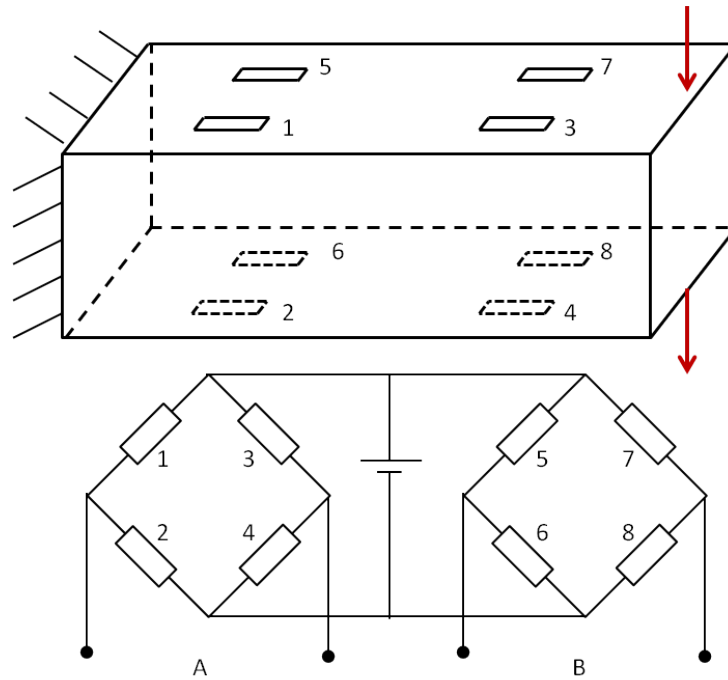
Figure 4.6: Strain Gauge Arrangement

The use of two strain gauge bridges allows an average force within the arm to be identified, in addition to providing redundancy, this protects the measurement from unequal force distribution due to inconsistencies in the fibreglass arm. The presence of the second bridge also presents options for the configuration of the bridges allowing measurement of lateral or torsional flexibility. By selecting the orientation in which to configure the strain gauge bridges, it is possible to compare the results of the two to obtain an estimate for torque. This is particularly interesting around rail transitions where the pitch of the shoe can have a significant impact on its behaviour.

4.4.3 Electronics

The various sensor inputs were conditioned and sampled using custom designed and built electronics mounted on the bogie of the vehicle. A block diagram of the electronics is shown in

Figure 4.7 with the significant components listed in Table 4.1. The input from the non-contact displacement transducer was conditioned using a basic anti-aliasing filter before undergoing analogue to digital conversion and then being sampled by the microprocessor. The two strain gauge bridges were conditioned and amplified using dedicated strain gauge management integrated circuits from Analogue Devices before, again, undergoing analogue to digital conversion and then sampling by the microprocessor.
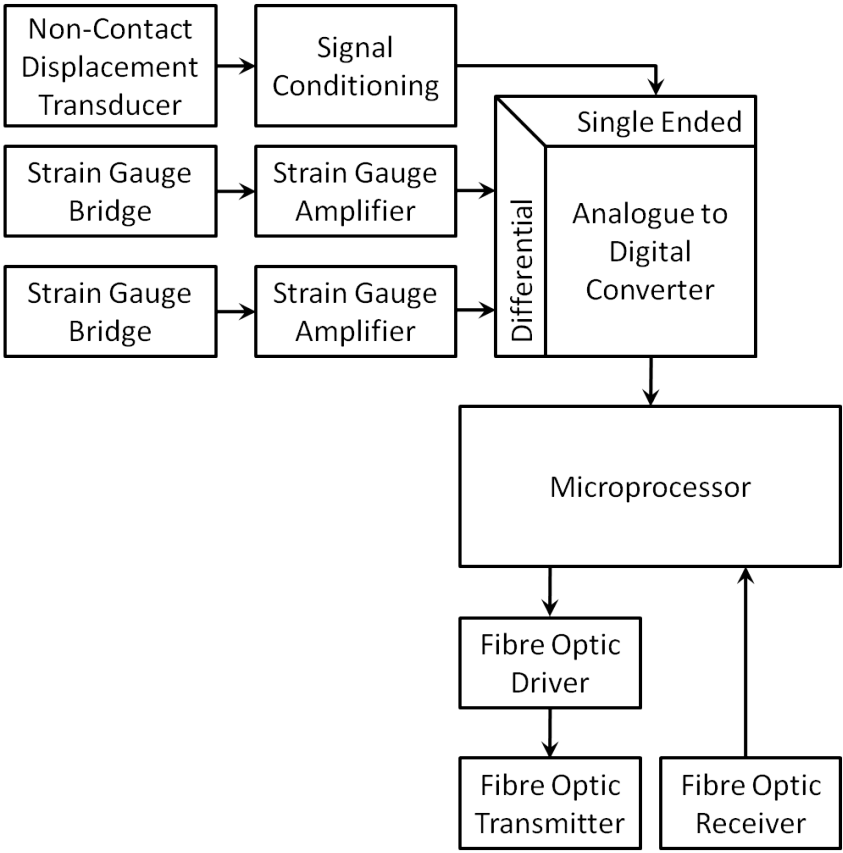
Figure 4.7: System block diagram

Table 4.1: System components

Component	Quantity	Description
Strain gauge	8	Strain gauge (4 in each of 2 bridges)
1B331AN	2	Strain gauge conditioning and amplifier module
ILD1700-250VT	1	Non-contact displacement transducer
LTC1859	1	Analogue to digital converter
PIC18F2680	1	Microprocessor
SN7451BP	1	Fibre optic driver
HFBR-1521Z	1	Fibre optic transmitter
HFBR-2521Z	1	Fibre optic receiver

The analogue to digital conversions were performed using an LTC1859 integrated circuit. This is a 16 bit analogue to digital converter capable of operating at one hundred thousand samples per second. The converter can be configured to operate in a number of ranges and in both differential and single-ended modes where it supports four and eight channels respectively. Communication with the unit is via an SPI interface.

The embedded processor selected was the PIC18F2680. The microprocessor was used to manage the analogue to digital conversions, package the recorded digital signals with suitable framing messages and transmit them to the PC component of the system using a fibre optic serial link. An external crystal oscillator with a frequency of 6 MHz was used to give a microprocessor operating frequency of 24 MHz.

The signals from the strain gauges required conditioning and amplification, this is particularly significant in the rail environment where interference from things like third rail arcing is particularly significant. This strain gauge management was performed by high quality strain gauge amplifiers from Analogue Devices mounted as close to the strain gauges as possible.

4.4.4 Firmware

The firmware for the microprocessor was written in assembly language. The code structure was built around a state machine which allowed the system to be held in a reset mode, or allowed to operate normally using an external signal from the user. When running, the non-contact displacement transducer and other ancillary inputs, such as temperature, were sampled at 750 Hz. The requests for these signals were interspersed with the signals from the two strain gauge bridges, which were sampled at 3000 Hz. The SPI messages that needed to be sent to the analogue to digital conversion unit, corresponding to different channel requests, were stored in a lookup table to make the sequencing simpler. The signals from the sensors were sent to the serial link as soon as they were received from the analogue to digital converter. The serial link was operating at a sufficiently high baud rate (approximately 1 Mbps) to transmit all messages without the need for a queuing system. After all of the channels had been sampled, the microprocessor inserted a framing message into the serial byte stream. This was a unique series of values that could not be obtained from the sensors, and could therefore be used to identify the start of each set of sensor data in the event of a synchronisation failure between the two ends of the communications link.

4.4.5 Computer interface

The interface between the electronics mounted on the bogie and the computer-based data logging system mounted in the vehicle body was in the form of a bi-directional fibre optic serial link. This link was in the form of a binary control channel from the vehicle body to the bogie, and a serial data channel from the bogie to the vehicle body.

On the bogie, the embedded microprocessor used a built in serial peripheral to generate an un-managed serial stream. This was passed to a dedicated driver integrated circuit (SN75451BP),

the output of which was connected to a fibre optic transmitter. The output of a fibre optic receiver was connected directly to the microprocessor to present the user control line.

In the vehicle body, a commercial USB to serial converter was used. A small interface box containing a fibre optic transmitter / receiver pair and a driver for the transmitter was produced to manage the media transition to the fibre optic link.

4.4.6 Software

The project made use of two pieces of software in addition to custom analysis processing carried out in MATLAB at a later stage. The first piece of software was used during the data acquisition to store the raw serial stream to a hard disk. The second piece of software was used to process this data stream into a format suitable for analysis. This two-stage approach was necessary to ensure the integrity of the serial stream was maintained.

The software that was operated on the laptop to capture the serial stream and store it to a file for later analysis, was written in C#. The priority for the software was to capture the data stream and store it, along with a time-stamp, on the hard disk in a "raw" format. Additional functionality was in the form of a decoded, calibrated, live display system in the form of numerical readouts and strip charts. This provided some level of immediate feedback, and thus confidence, during the trials. The second aspect was of a lower priority and so was operated using a separate thread.

This first piece of software was also responsible for issuing control commands to the electronics mounted on the bogie. These were incorporated into the logging control system such that single data acquisition sweeps could be displayed for testing, but any continuous data stream was always recorded.

The second piece of software was used to process the raw serial stream into usable data. This

process was essentially a more thorough version of the live display system, which stored its outputs to a file rather than displaying them. The processing consisted of identifying the framing messages and then, using the deterministic nature of the sampling sequence, passing the serial bytes to a series of format conversion and calibration routines. Data was time-stamped according to the framing message associated with it. If the framing messages were found to be the wrong number of bytes apart, the serial stream was assumed to have "slipped" due to noise in the system disrupting the communications. In this case a re-alignment routine was used to retain synchronisation and discard erroneous data.

4.4.7 Mechanical construction

The instrumentation developed for the project was fitted to a dedicated Class 375 train. The train was not in passenger service and as such it was possible to disconnect one set of shoegear from the 750 V bus. This meant that the instrumented shoegear was not collecting current, which allowed a degree of flexibility with the positioning of the instrumentation. In the final version, the arc shielding was removed to make space for the instrumentation. This meant that the same mounting points could be used. The instrumentation was designed to fit into the space previously occupied by the arc shielding. The size of the instrumentation was governed by these mounting points and a requirement not to interfere with the possible path of the fibreglass arm, as illustrated in Figure 4.2.

Safety calculations

Safety considerations for the instrumentation were broadly assigned to three categories:

1. Fire

2. Structural

3. Electrical

As the instrumentation was largely housed outside the vehicle body, and the vehicle was not used in tunnels, the first of these categories was comparably easy to satisfy. The custom electronics was of insufficient mass to sustain any substantial fire and was contained in a metal box. The fibre optic cable was of a type with a flame retardant sheathing, and the electrical cable used was railway approved.

The second category was largely concerned with the instrumentation remaining attached to the vehicle. As the trials were scheduled to be completed in one day, full fatigue calculations were not required. The instrumentation was fitted in place of the arc shielding, but as the instrumentation was physically smaller than the arc shielding, not all of the bolt holes could be used. Despite extensive measures to reduce the mass of the instrumentation as shown in Figure 4.8, it was not possible to reduce the mass to a point where the ratios of mass (arc shield mass to instrumentation mass) and bolts were equivalent. Instead, an external consultancy was used to perform "bolt calculations" which were used as technical documentation in support of the instrumentation remaining attached.

Electrical considerations were largely related to the modification of the shoegear and the removal of the arc shielding. By removing the current collection cables, the shoe was (electrically) isolated from the vehicle and thus did not form arcs when disconnected from the rail. Hence, the removal of the arc shielding was acceptable as the only major risk was of a foreign object, collected from

the line, being trapped between the shoe and the vehicle body. This risk was minimised by the shoegear being the trailing shoe in 50% of the movements carried out during the trials.
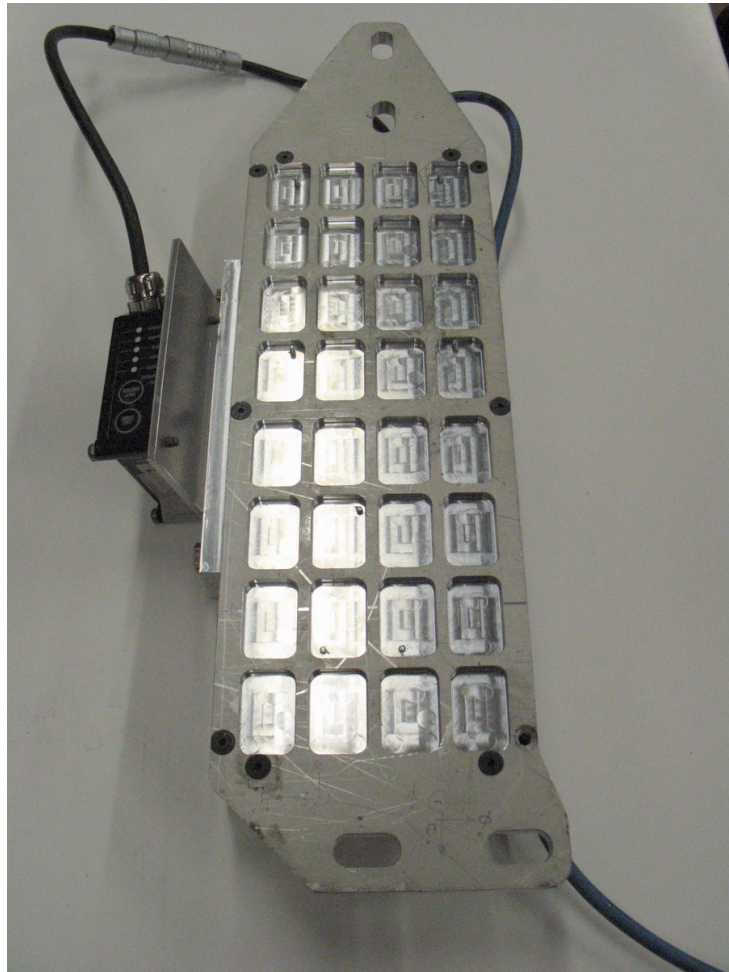
Figure 4.8: Efforts to reduce the mass of the system

The application of strain gauge wires to the fibreglass arm presented a risk as it put an electrical connector in proximity to the conductor rail. This was minimised by encasing the wires in thermally conductive rubber with electrical insulation properties of >8 kV / mm. The wires were also enamel coated and sufficiently thin that they would have acted as fuses should any contact have occurred.

The link between the bogie and the vehicle body was another area of concern. The signal path was implemented using fibre optics but the power for the instrumentation was supplied from

the vehicle body. Batteries could not have been used due to the space constraints on the bogie. The power supply unit used in the vehicle body was fused and insulated to mitigate as much of the risk as possible.

4.5 Trials

The trials took place on the 21st of March 2008. They consisted of three runs on a 16 km stretch of the up line between Birchington on Sea and Chestfield and Swalecliffe. The trial section is highlighted in red on Figure 4.9. To ensure the correct speed throughout the whole of the trial section, and to fit in with the timetable, the vehicle path extended as far as Margate in the down direction and Faversham in the up direction. The three runs were carried out at three different, constant speeds as shown in Table 4.2.
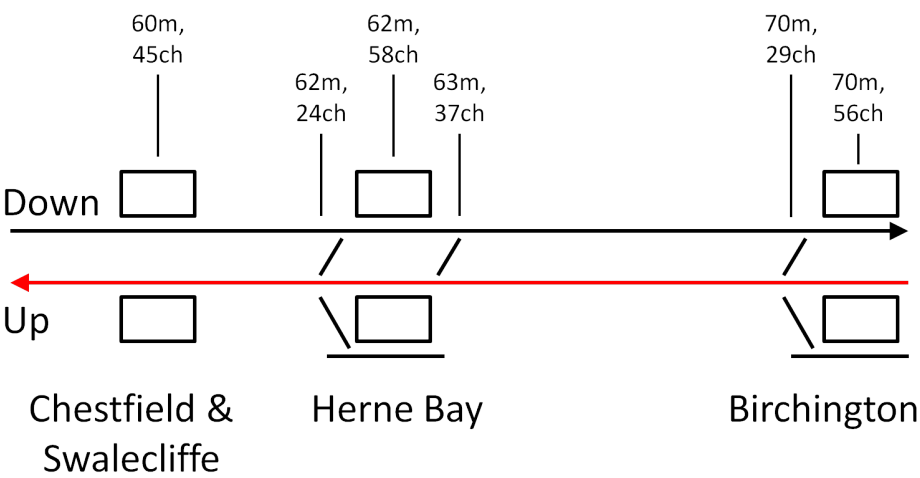
Figure 4.9: Trial route

Table 4.2: Trial runs

Run	mph	km/h	m/s
1	20	32.2	8.9
2	50	80.5	22.4
3	75	120.7	33.2

4.6 Results

4.6.1 Height above rail level

The height of the conductor rail above the running rail level was established by measuring from the tip of the fibreglass arm to the bogie. Geometrical conversions were then used to account for the heights of the bogie, the shoe bracket and the thickness of the shoe. Of course, the measurement is only accurate when the shoe is in contact with the conductor rail. At areas of impact, such as ramp ends, the shoe often loses contact with the conductor rail. In this case the measurement is not valid.

The height above rail level result for the 20 mph pass of the 16 km trial section is shown in Figure 4.10. Each data point on the graph corresponds to approximately 10 mm of track, mechanically smoothed over the shoe length of 150 mm. Although this is a dynamic measurement, the vast majority of the result exists within the envelope defined by the static tolerances of 72 mm to 86 mm. There are notable excursions beyond the lower end of this limit. These represent times where the shoe was either on a ramp, or there was no conductor rail and the shoe was suspended from the down-stop. In the latter case, the vertical distance from the running rail to the base of the shoe is 53 mm.
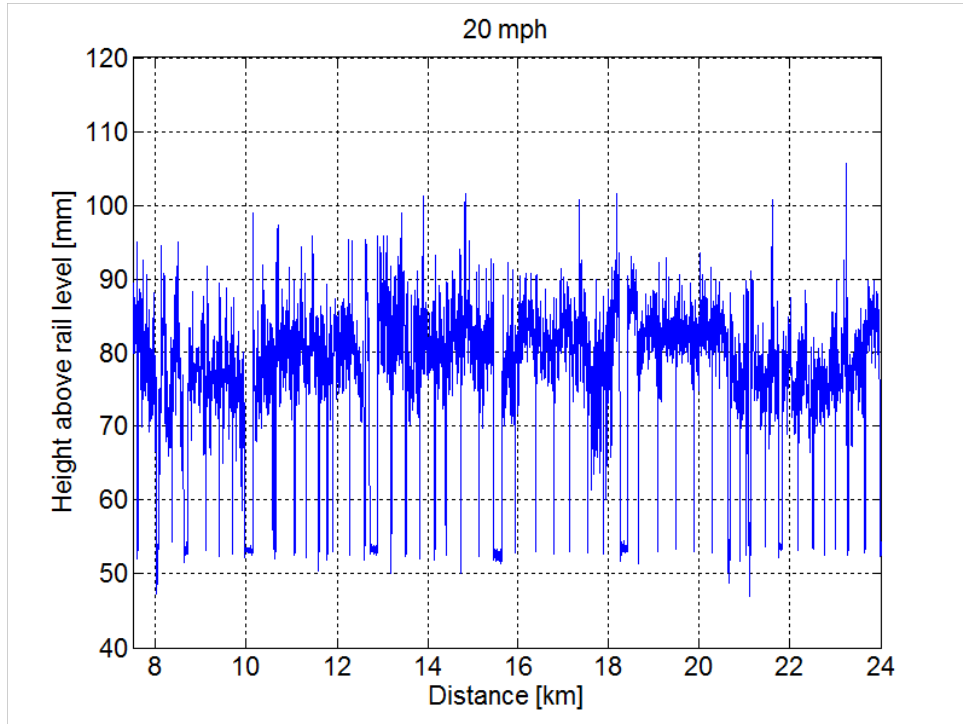
Figure 4.10: Conductor rail height above rail level for 20 mph trial run

The results for the 50 mph and 75 mph passes of the same trial section are shown in Figure 4.11 and Figure 4.12. The results show good consistency over the three trial runs with the majority of the result still remaining within the static tolerances. More extreme outlying points appear as the speed increases, but this is to be expected in the dynamic environment. Of note is the down-stop value. The down-stop bracket is mounted on the vehicle and so results taken when the shoe is on the bracket should be reasonably constant, however, as the speed increases, the results become more varied and, at high speeds, even seem to go below the down-stop limit. This is due to the construction of the down-stop system. The bracket itself is mounted on a flexible fibreglass beam, which spans between the two axleboxes. With greater speed, the dynamic effects put greater forces into the system. This causes a greater degree of flexing to be exhibited by the beam.

114
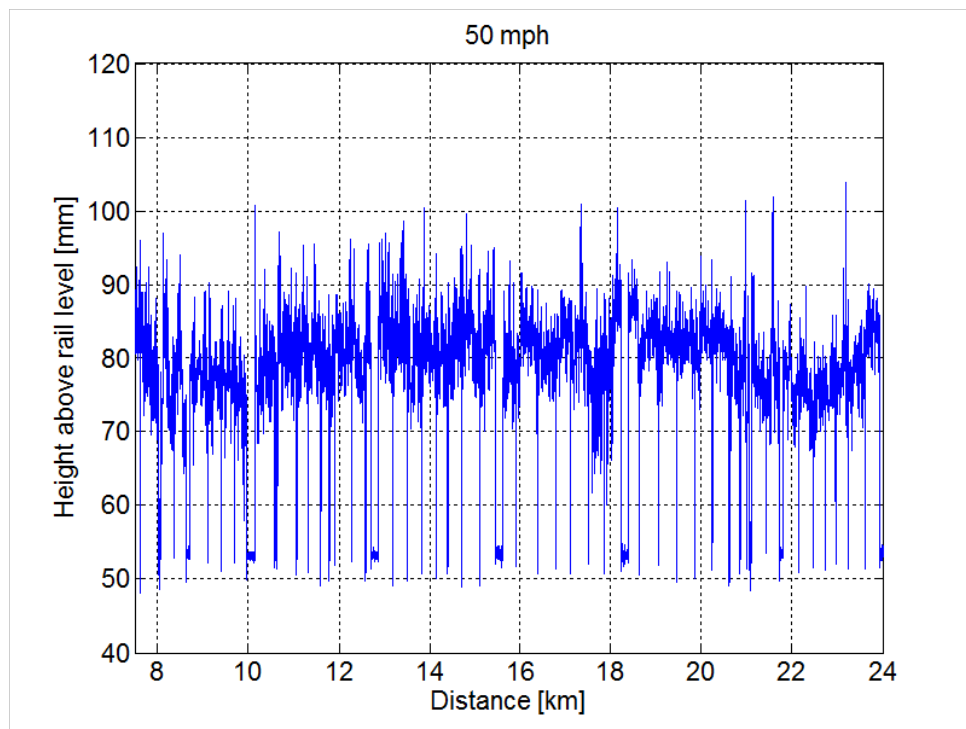
Figure 4.11: Conductor rail height above rail level for 50 mph trial run
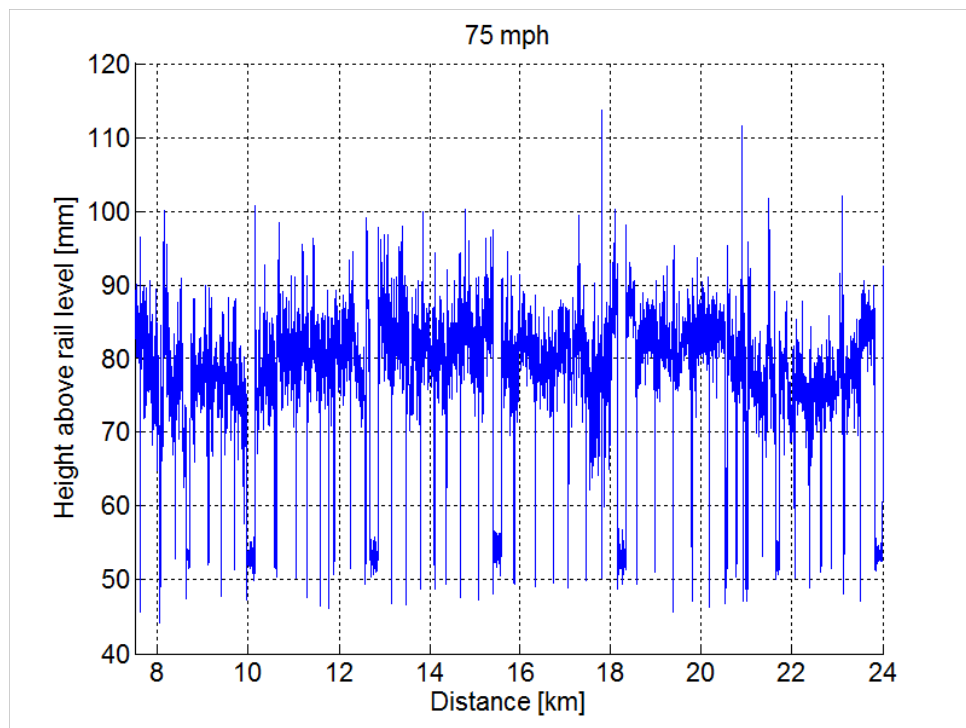
Figure 4.12: Conductor rail height above rail level for 75 mph trial run

To better quantify the behaviour of the conductor shoe, and assumed rail height, in relation to the defined tolerances, it is possible to consider the histogram of observed heights. This has the added advantage of allowing an easy comparison of the three different runs. Figure 4.13 is one such histogram, it was generated using 0.5 mm bins. The figure shows two peaks, one for normal operation and a second representing the time that the conductor shoe has spent hanging on the down-stop.
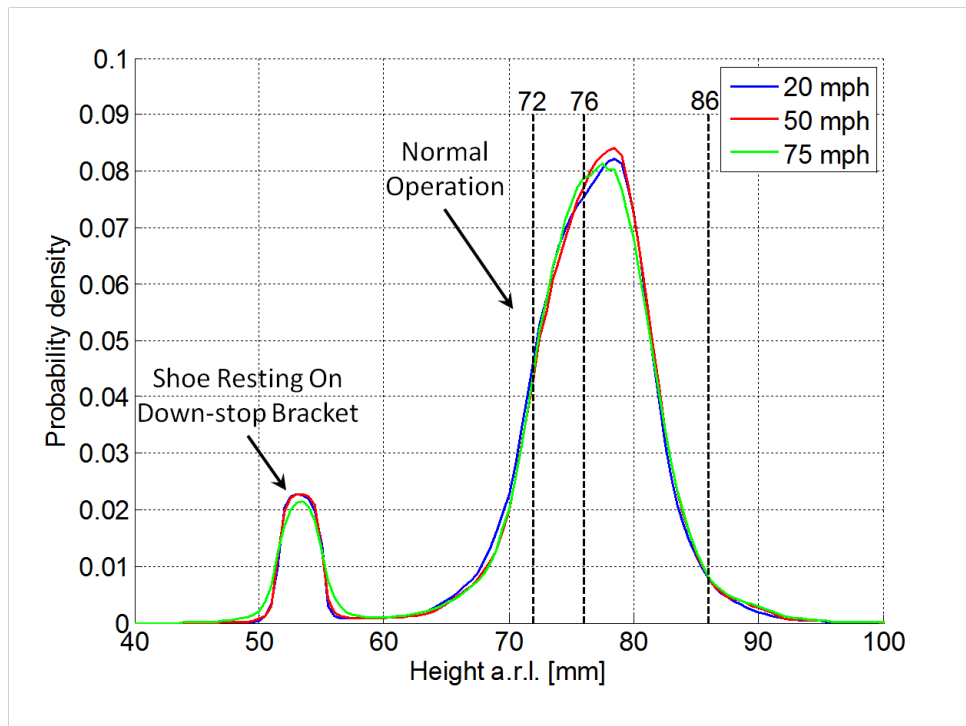
Figure 4.13: Height a.r.l probability distribution including downstop

As the focus is on using the position of the conductor shoe to identify the relationship between the running and conductor rail heights, it is acceptable to remove the data where the shoe is not in contact with the rail. Including this pre-processing step, before constructing the histogram produces Figure 4.14. The figure shows the consistency of the system over the three speeds. This is a clearer representation than the separate figures showing the raw data as the MATLAB plotting algorithms tend to highlight outlying data points.

The static tolerances are marked on the figure. While the modal average height is slightly above the nominal, the main peak is skewed to the bottom end of the tolerated range. This brings the mean value of 77 mm back towards the nominal 76 mm.
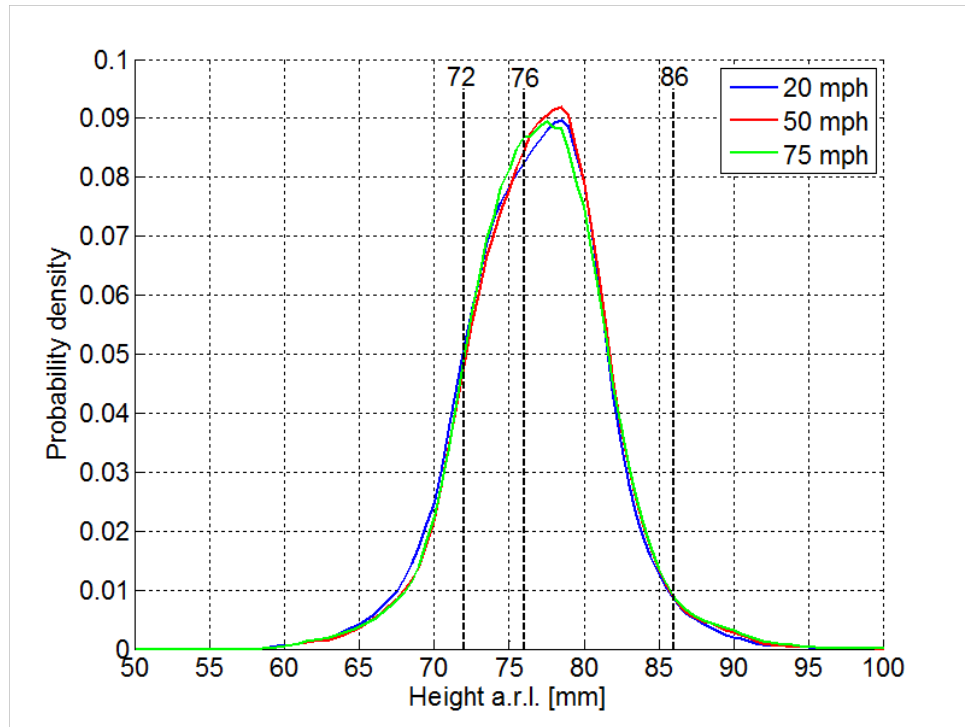
Figure 4.14: Height a.r.l probability distribution excluding downstop

While Figure 4.14 does indicate the probability that any given sample will fall outside the acceptable static tolerances, this information is more easily identified using a cumulative probability curve. Figure 4.15 illustrates the cumulative probability up to any particular height. From the figure, it is possible to determine that 14% of the conductor rail falls below the static standard's minimum permitted value of 72 mm. The conductor rail is less likely to be high, with only 3% being above the static maximum of 86 mm. The figure also confirms that the mean conductor rail height is 77 mm. Once again, the figure also indicates consistent behaviour from the three runs at different speeds.
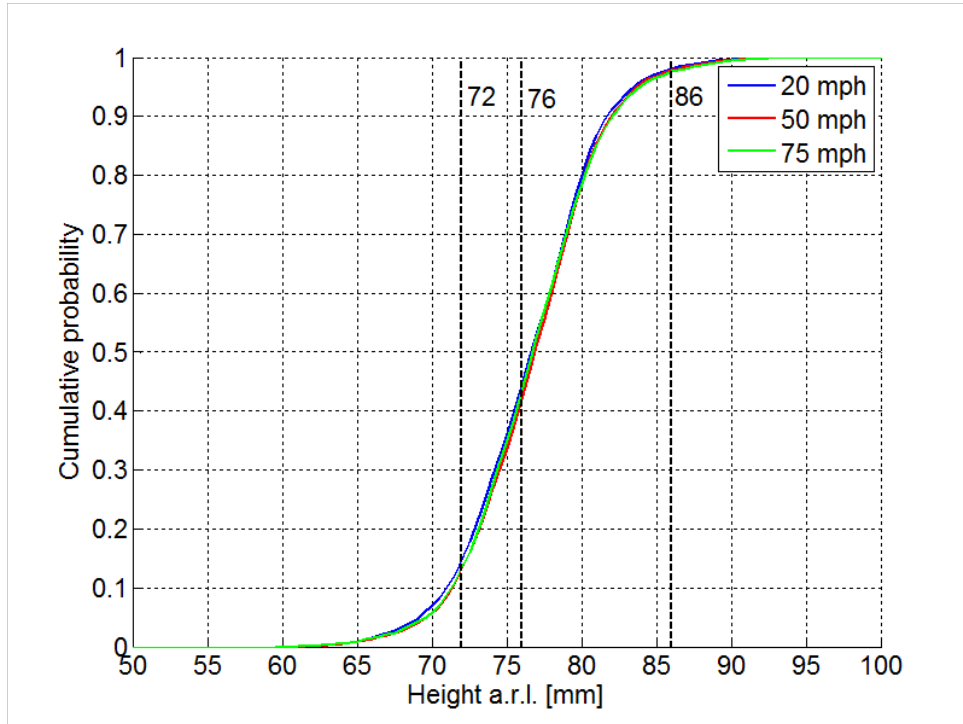
Figure 4.15: Cumulative height a.r.l probability distribution excluding downstop

To consider this consistency over the three different speeds, it is possible to compare the trial runs using results averaged over each of the 16 kilometres of the trial section. This is shown in Figure 4.16. The figure shows the largest discrepancy between speeds as approximately 2 mm. In the region of the test where this occurred there are numerous, brief excursions to the downstop position but the shoe does not remain on the down-stop for any substantial period of time. Such excursions imply the presence of up and down ramps where the dynamic effects of travelling at greater speed are maximised.

Standard deviations can be used to establish a metric for how closely the data obtained are clustered around the mean value. This can be used as another way to determine the consistency of performance of the shoegear when operating at the three speeds. Table 4.3 shows the mean values and standard deviations for each of the trial runs. The results include the data taken when the shoe was on the down-stop.
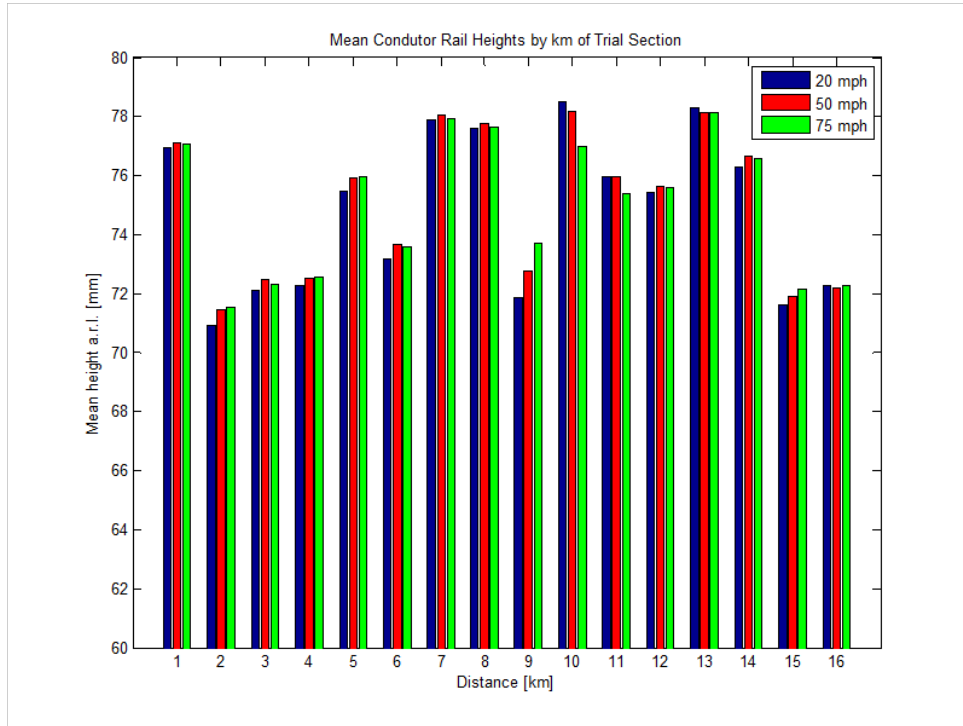
Figure 4.16: Mean height a.r.l for each track km and for each trial run

Table 4.3: Means and standard deviations of height a.r.l for each trail run

Speed [mph]	Mean [mm]	Standard deviation [mm]
20	74.8	7.9
50	74.9	8.0
75	75.0	8.0

As shown previously in Figure 4.13 and Figure 4.15, including the data taken when the shoe is on the down-stop means that a second, artificial, peak occurs in the distribution of rail height values. This clearly distorts the real measurement of conductor rail height. The positions of the conductor rail gaps are, however, consistent between the three runs and so it is still acceptable to use mean and standard deviation values for comparative purposes. Figure 4.17 shows the values for one standard deviation calculated for each kilometre of the trial section, and for all three speeds. The consistency is again very apparent with only a millimetre of discrepancy between the runs in the most extreme case.
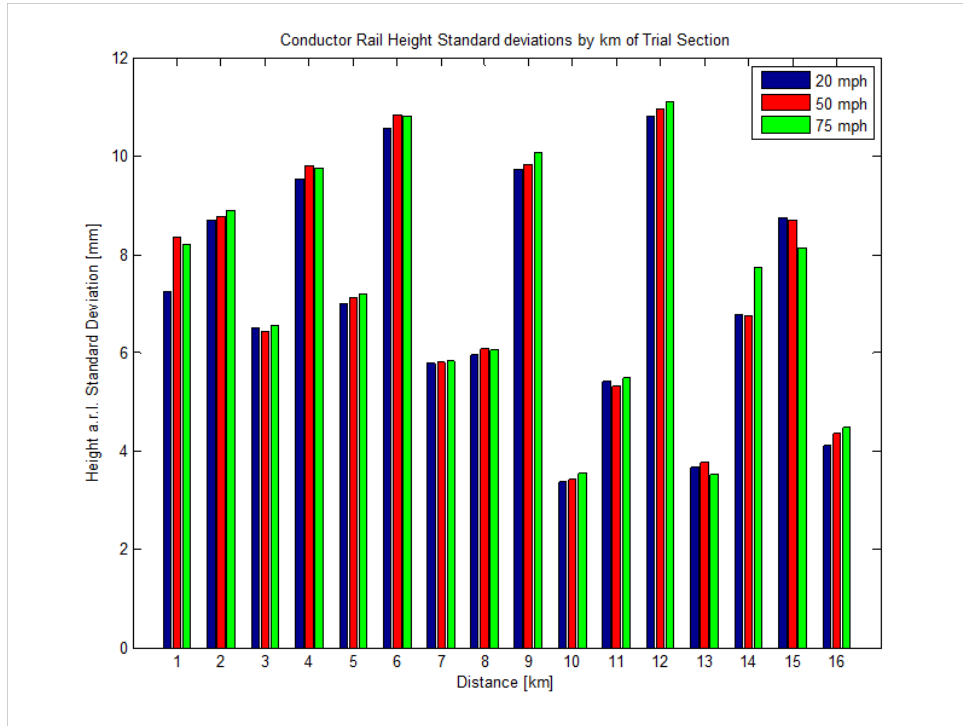
Figure 4.17: Standard deviation of height a.r.l for each track km and for each trial run

While both figures (Figure 4.16 and Figure 4.17) illustrate the high level of consistency between the three trial runs, it is not possible to use data at this resolution to explain the variation in standard deviation. Looking at the figures, there is a possible correlation whereby the greater the height a.r.l, the lower the standard deviation - for example between 6 and 10 km. Plotting the two factors against each other, however, reveals little correlation, as indicated by Figure 4.18. Increasing the distance resolution from 1 km to 100 m (Figure 4.19) suggests, however, that there may be a correlation, particularly at the upper bounds of the height a.r.l. where the standard deviation is typically below 5 mm. This could be due to the torsion spring being at one extreme of its operation. The other end of the figure indicates that the standard deviation is particularly low when the shoe is on the down-stop, which is clearly visible.
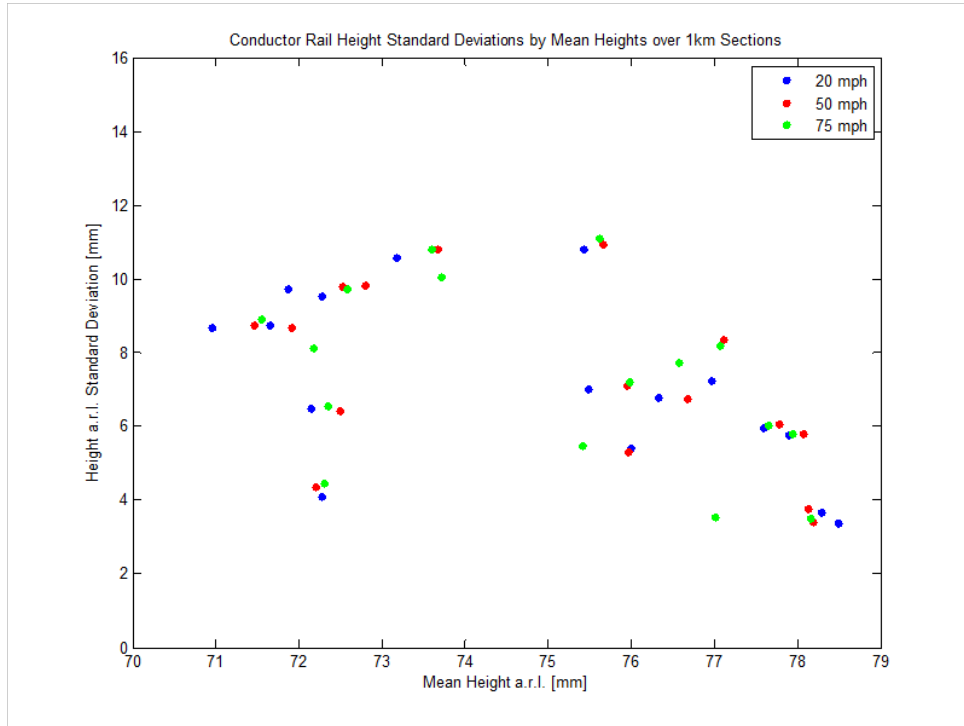
Figure 4.18: Correlation plot for mean and standard deviation over 1 km
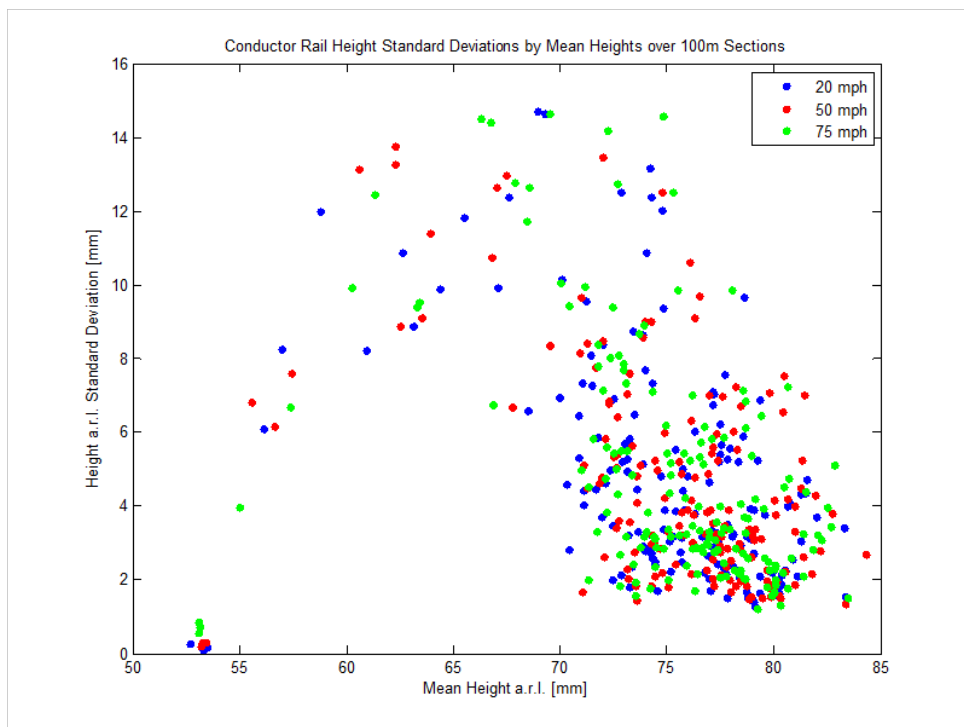
Figure 4.19: Correlation plot for mean and standard deviation over 100 m

121

4.6.2 Force

The eight strain gauges used in the instrumentation were arranged in two strain gauge bridges distributed to be representative of the whole of the fibreglass arm. The result from each bridge was acquired and recorded by the system three thousand times per second. By averaging the results from the two bridges, it is possible to obtain an estimate for the force within the fibreglass arm. This can be combined with the weight of the shoe and shoe bracket to give an estimate for the contact force between the shoe and the third rail under normal operation. In line with best practises, the sampling rate of 3 kHz was selected to be sufficiently high (in line with Nyquist theory) to allow consideration of the forces associated with the movement of the shoe and bracket, which have a combined mass of 10.1 kg and thus an inertial aspect to their behaviour. Futhermore, the upper limit of the sampling rate was based on the capacity of the communications link between the two sides of the instrumentation architecture. The selection was probably rather high, given that the system does not measure impact style dynamic effects, but the capacity was available and access to the vehicle was so limited that it was felt that every effort should be made to obtain as much potential information as possible.

The recorded forces for the 16 km trial section are shown in red in Figure 4.20. The blue line on the figure is the conductor shoe height a.r.l. and can be used to identify when the shoe and rail are in contact. At 20 mph, and with a sampling frequency of 3 kHz, each force point on the figure represents a distance of approximately 3 mm. Again it is worth noting that the software used to produce the graphs will enhance the appearance of extreme or outlying points.

The first thing to notice, is that the vast majority of the data recorded fall within the static tolerance range of 250 N ±50 N. The inclusion of the blue, height a.r.l. data in the figure is primarily to illustrate the correlation between the lower average forces and the time spent with the conductor shoe hanging on the down-stop. Although the system is not designed to

measure high impact, instantaneous forces, the peak values can be used indicatively and are ratiometrically consistent. It is possible to observe from the figure that the highest forces occur at the end of gapped sections, where the shoe regains contact with the rail, having spent time hanging on the down-stop. These are likely to be impact forces.
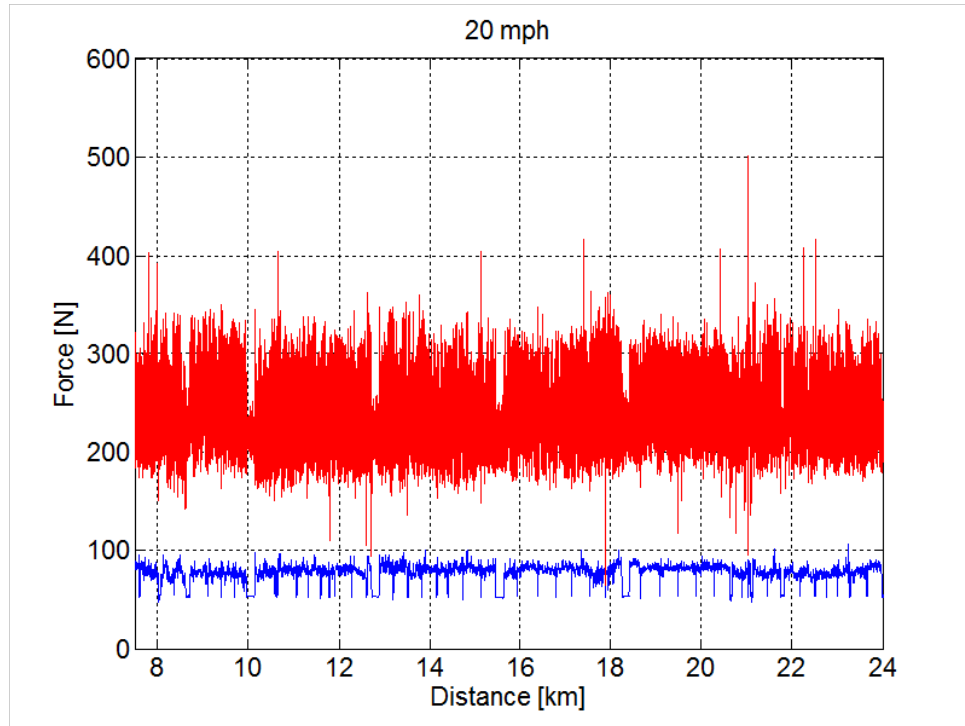
Figure 4.20: Contact force for 20 mph trial run, with overlaid displacement measurement (blue)

Figure 4.21 shows measurements of the same section of the track taken with the vehicle travelling at 75 mph. Again the force measurements are shown in red and the height a.r.l. is shown in blue. The figure shows a far greater dynamic component to the measurement. Peak forces, particularly at the ends of gapped sections, are greatly increased. The results during regular contact, however, are still largely consistent with the 20 mph test run and are, in the majority, once again within the static tolerances.
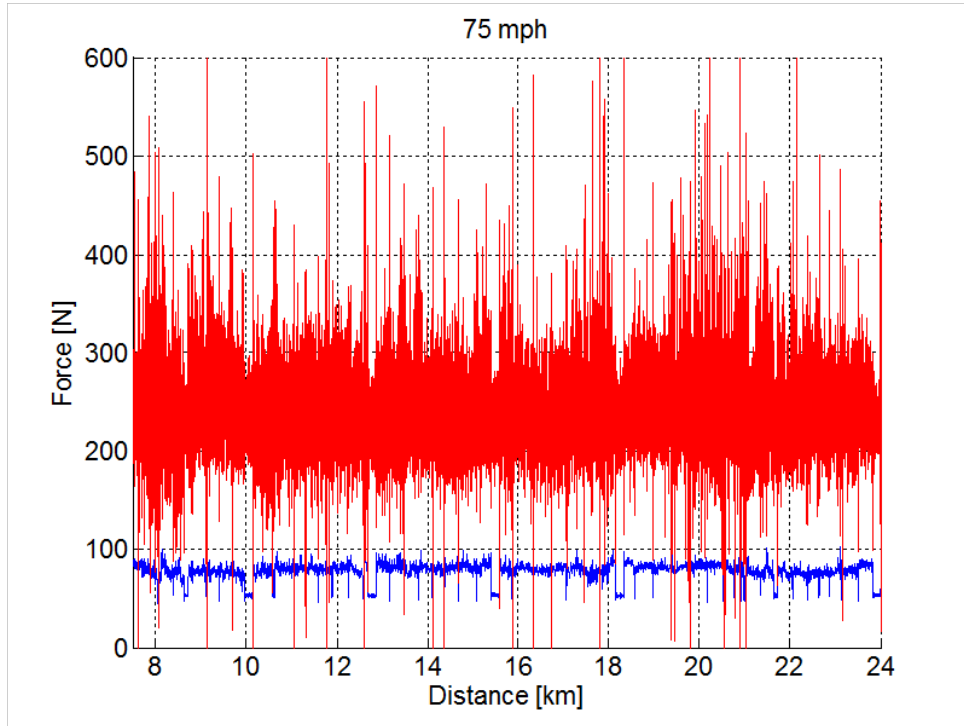
Figure 4.21: Contact force for 75 mph trial run, with overlaid displacement measurement (blue)

Closer consideration of a single gapped section taken from these two figures confirms that the contact force is generally consistent, with differences occurring due to higher frequency effects present in the data recorded at higher speed. Figure 4.22 and Figure 4.23 show the data taken around one 200 m gapped section. The higher frequency force variation when the conductor shoe is hanging on the down-stop is likely to be due to the resonances in the shoebeam which is a flexible structure that hangs between the axle boxes and has the shoe down-stop bracket mounted on it. This may also have an effect on the larger impact forces as the direction of shoe movement due to this resonance could be significant in the final impact forces recorded.
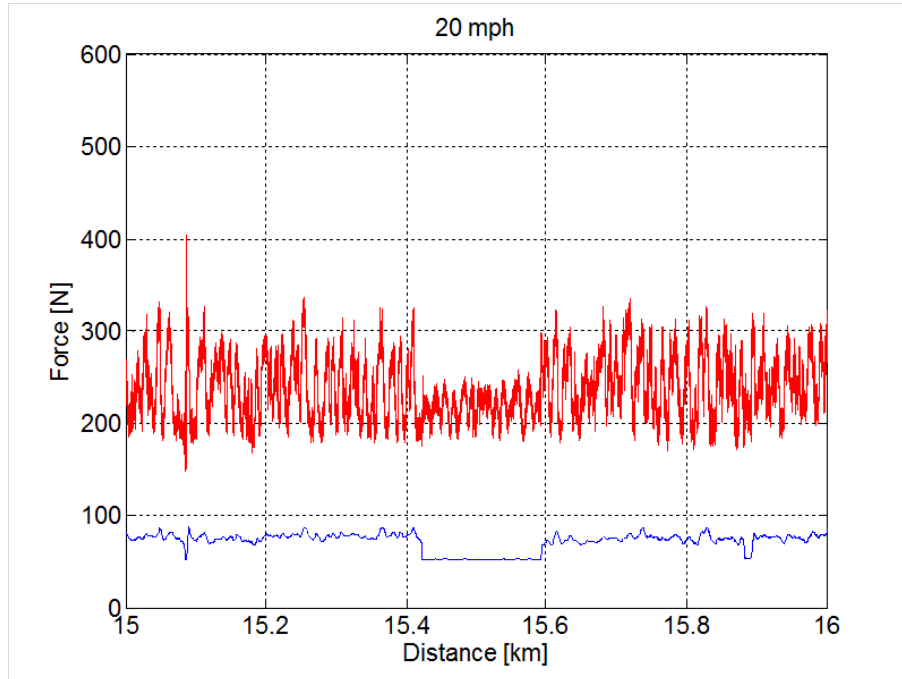
Figure 4.22: Contact force for 20 mph trial run shown at a gapped section, with overlaid displacement measurement (blue)
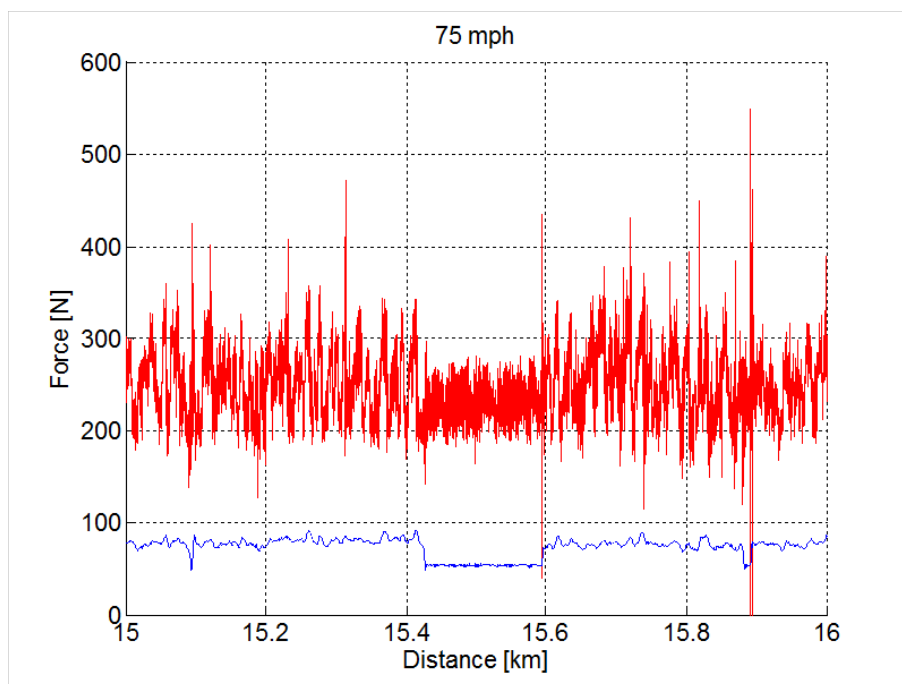
Figure 4.23: Contact force for 75 mph trial run shown at a gapped section, with overlaid displacement measurement (blue)

As with the height a.r.l. results it is possible to use a histogram to gain a better understanding of the contact forces over the whole trial. This also allows comparison of the three different speed runs to be made more easily. Figure 4.24 shows the probability density of the different forces taken over the 16 km trial section. The figure shows that, although not quite centred on the 250 N nominal, the mean force is reasonably close to this, and the majority of the data falls within the ±50 N tolerance of this value.
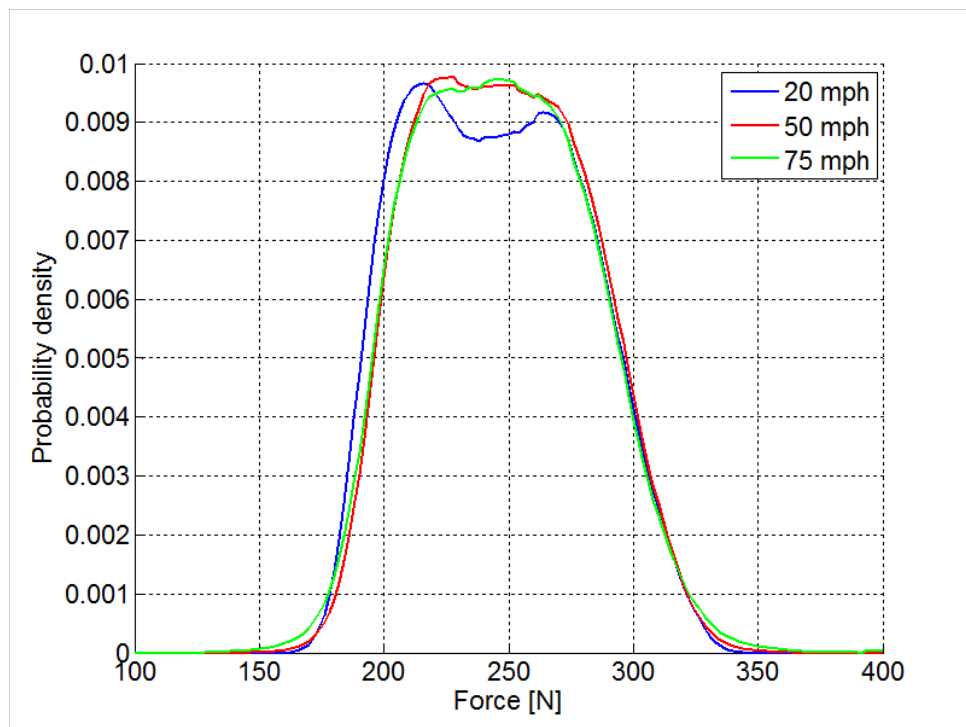
Figure 4.24: Contact force probabilty distribution

The 50 mph and 75 mph runs are reasonably consistent which suggests that the short-term, high magnitude forces do not significantly contribute to the result. There is a slight bias towards lower forces found in the 20 mph test. This may be significant as it speaks towards the mode of behaviour of the shoegear. The main arm of the shoegear is made of laminate fibreglass. Due to this construction, there is a small amount of flexibility in the arm. While reasonably easy to flex, in comparison to the main torsion spring, the arm itself can only adapt to small deflections. The higher frequency, smaller magnitude deflections can, therefore, be accommodated by the

flexibility of the arm, but larger deflections must still be accommodated by the torsion spring. Figure 4.22 and Figure 4.23 suggest that larger deflections, which tend to correlate with larger forces, occur at higher speeds. At lower speeds, the flexibility of the fibreglass arm may be able to accommodate some of the force.

To get a better understanding of how the observed forces relate to the nominal value of 250 N and tolerances of ±50 N, a cumulative probability plot is used Figure 4.25. The figure again shows good consistency across the three speeds, with the 20 mph curve revealing slightly reduced forces when compared to the other two. The mean force is approximately 240 N, slightly lower than the nominal value. The force only falls outside the ±50 N tolerance around 15% of the time (10% low and 5% high) but as with the height above rail level, these tolerances are for static conditions and the measurements are dynamic.
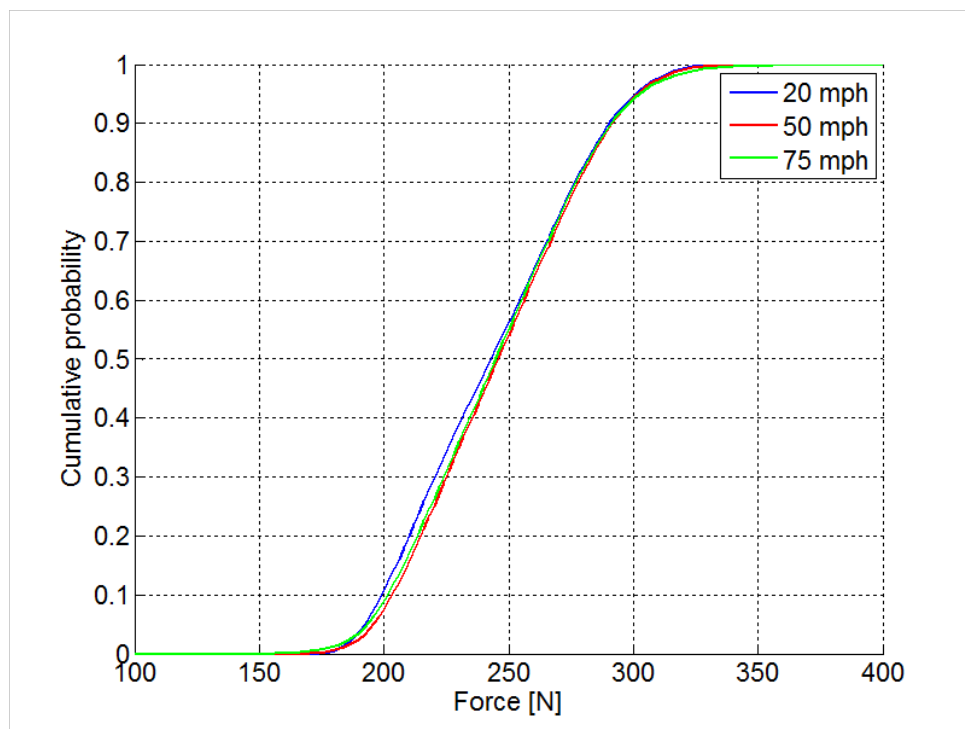
Figure 4.25: Contact force cumulative probabilty distribution

4.6.3 Height above rail level vs. Force

A simple model for contact force might include a primary component consisting of a force proportional to spring displacement. In this case, plotting contact force against displacement should reveal a straight line. This is not the case with the shoegear where laboratory tests have revealed significant hysteresis effects in the shoegear system. This is thought to be caused by the "slip-stick" nature of the torsion spring and the sheath around which it is wound. Due to this apparent hysteresis, direct comparisons of single measurements of height and force are not necessarily meaningful. An alternative way to visualise the relationship between the height a.r.l. and the force is to use a contour plot to consider multiple cumulative probability densities in parallel. An example of this kind of plot is shown in Figure 4.26.

The figure is produced by taking the data, in this case from the 50 mph run and dividing it into 1 mm wide bins based on the height a.r.l. The data is then used to generate cumulative probabilities of force which can be stacked together to generate a contour plot. Data is extracted on the 5% intervals to generate the final figure in a more readable format.

From the figure, it is possible to see that, despite the hysteresis, there is a general trend towards greater forces being recorded at greater conductor rail heights. This would correspond to the torsion spring being under greater strain when the shoegear is raised. The distribution of the contours seems most uniform around 80 mm. Below this, there is a concentration of forces towards the (static) lower limit of 200 N, and above towards the (static) upper limit of 300 N. At the very extremes of conductor rail height (below 67 mm and above 95 mm) the force distribution begins to spread out again. Using the 50% contour line, it is possible to see that a mean contact force of the nominal 250 N is obtained at 82.5 mm. Hence, the force is somewhat low for the nominal conductor rail height of 76 mm.
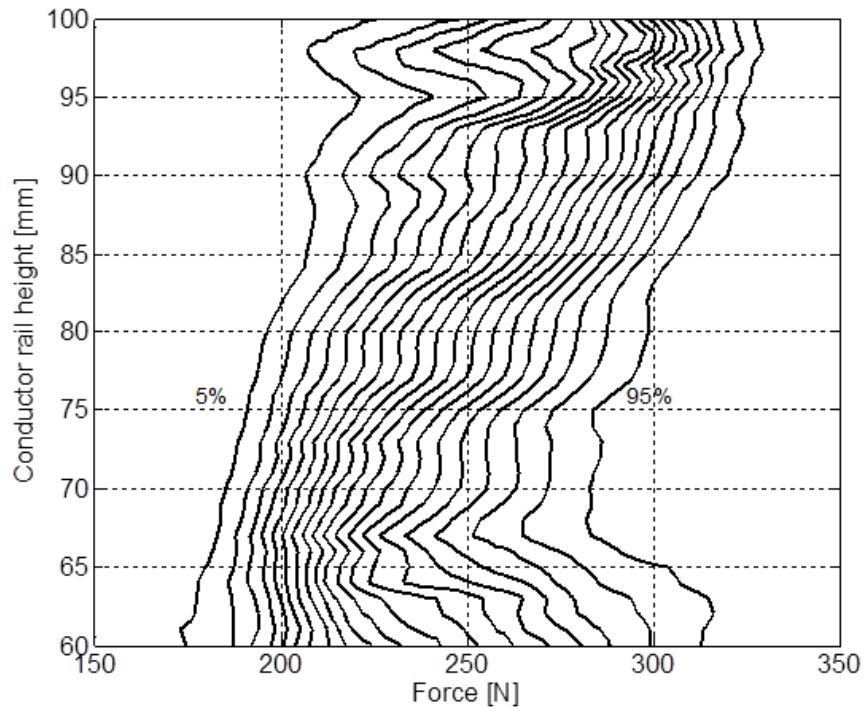
Figure 4.26: Contour plot of multiple cumulative probability densities of force

4.6.4 Torque

As mentioned previously the use of eight strain gauges, arranged in two separate strain gauge bridges introduces the possibility to take measurements of either the lateral or torsional flexibility of the fibreglass arm. Laboratory tests were undertaken to compare the merits of the two approaches. Lateral variation in the position of an application of force to the underside of the conductor shoe demonstrated limited observable torsion. This is thought to be due to the width of the shoe and the structure of the shoe mounting bracket. Torsional variation was much more clearly visible when the force was applied over the (substantially longer) length of the conductor shoe. Furthermore, in general operation the lateral variation of the position of the conductor rail relative to the running rails should be minimal, and the lateral variation of the bogie on the running rails limited. Given the also limited curvature of the top surface of conductor rail,

and the natural wear of the conductor shoe to conform to this, there should in theory be no substantial distribution of force laterally across the bottom of the shoe until extensive movement has occurred. Torsionally, however, there are sections of the railway where ramps are used to smooth the passage of the conductor shoe on and off the conductor rail. For the shoe to conform to these ramps, some element of pitching must be possible and this should be identifiable as torsional flexibility within the arm. Finally, tests conducted into the removal of ice from the conductor rail using conductor shoes had suggested some pitching motion of the shoe, where the toe became trapped in ice was occurring. While no ice was expected during the trials, an understanding of this pitching motion, if any, would support that previous work. It was therefore elected to configure the strain gauge bridges for torsional measurement as it was felt that this would yield more valuable information about the dynamic behaviour of the shoegear.

By configuring the strain gauges as a front and back pair (relative to the direction of travel) it is possible to use the difference between the two results to infer something about the pitch of the shoe. Torque is a measure of force multiplied by the distance through which that force is applied around a point. By referencing to the longitudinal centre of the shoegear it is possible to identify nominal values of torque for theoretical contacts at the front and rear of the shoe, and also for the shoe down-stop bracket. Applying the nominal 250 N through the tips of the shoe, and the observed mean down-stop force of approximately 200 N on the tip of the down-stop gives predicted values of ± 37.5 Nm. This is shown in Figure 4.27
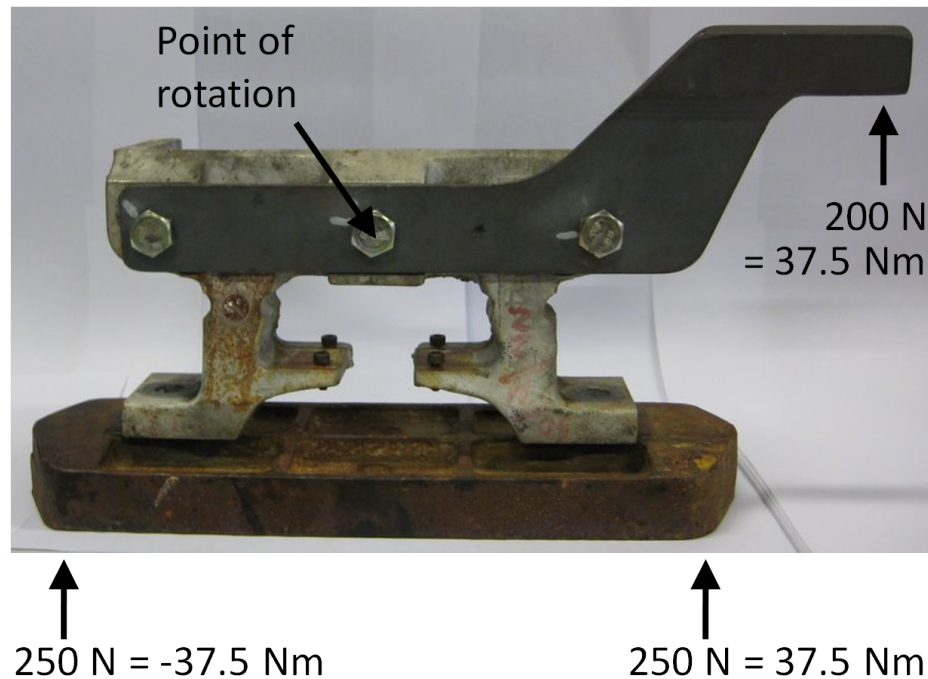
Figure 4.27: Contact points on a conductor shoe with associated torque values

The results for torque, as obtained using the difference between the two strain gauge bridges are shown in Figure 4.28. As predicted, the majority of the values fall within a range of ±37.5 Nm. Notable exceptions include the times when the shoe was hanging from the down-stop bracket. At these times the torque is slightly higher than expected, possibly due to the quality of the assumption of the force to be 200 N at this speed.

Also of interest is the tendency for the torque to be positive, the mean torque at 20 mph is 13 Nm. However, not including the time spent on the down-stop, the largest values of torque displayed are all negative. The negative values are larger because of the frictional forces which act along the bottom of the shoe. Assuming a coefficient of friction of 0.2, and a nominal contact force of 250 N being applied through 100 mm, the contribution of friction to the torque is approximately 5 Nm. One explanation for the slightly positive bias in the torque is that at impacts, such as joints, the shoe is thrown upwards by the impact rather than being pulled

downwards by the shoe dragging over the point of impact.
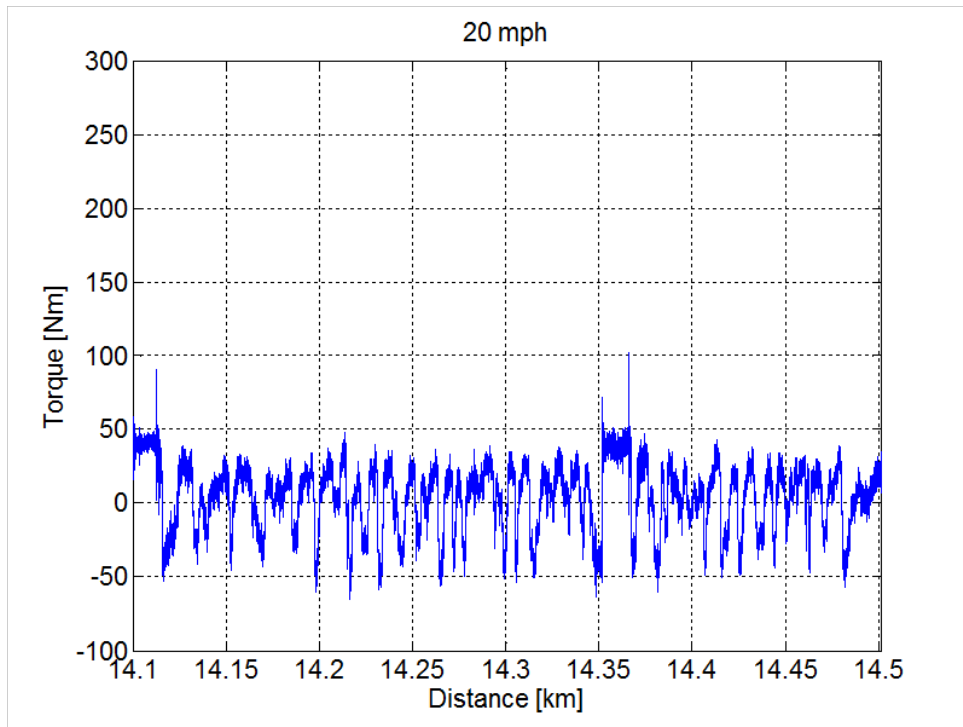
Figure 4.28: Torque in the fibreglass arm for the 20 mph trial run

Increasing the speed of the vehicle has two noticeable effects on the torque. These are shown in the comparison of Figure 4.28 and Figure 4.29. As the torque is inherently linked to the force, and the force showed increased high frequency components at higher speeds, one effect is that the higher frequency components of the torque measurement become more pronounced. This is particularly noticeable during gapped sections when the shoe transitions onto and off the down-stop. As with the force measurements there is some question of the accuracy of the system during these high impact transitions, but it is still possible to use the results comparatively.

The other noticeable effect is a decrease in the average torque. One factor already associated with negative torque is friction; however, friction should not alter with speed. That said, the dynamic contact behaviour of the shoe with the rail may be changing. As suggested previously, the small magnitude, higher frequency, components of the force may be accommodated by the

flexibility of the fibreglass arm while the lower frequency, larger magnitude, components are dependent on the torsion spring. If this is the case, the increase in speed may lead to more of the higher frequency components and thus change the behaviour of the shoe in relation to the rail. It is thought that a rocking motion may dominate the dynamics at higher speed while a stick/slip approach may prevail at lower speeds. This would suggest lower values for torque, and greater susceptibility to friction, when the speed is increased.
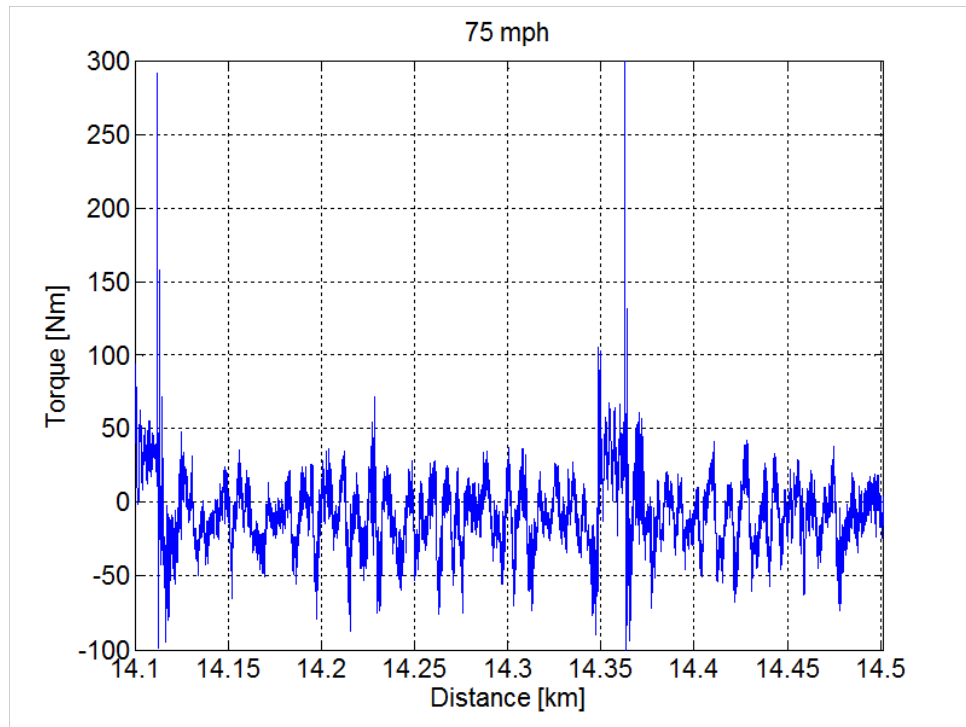
Figure 4.29: Torque in the fibreglass arm for the 75 mph trial run

This difference in torque with increasing speed can be shown particularly clearly using a cumulative probability curve, Figure 4.30. The figure suggests that the oscillation of torque around zero (flat consistent operation) occurs at approximately 50 mph, with the higher and lower speed runs skewed positively and negatively respectively. The figure also indicates the trend towards lower frequency, greater magnitude, negative values of torque with convergences of the three speeds occurring at -50 Nm and 38 Nm. This is reflected in the large negative spikes or torque seen in Figure 4.29.
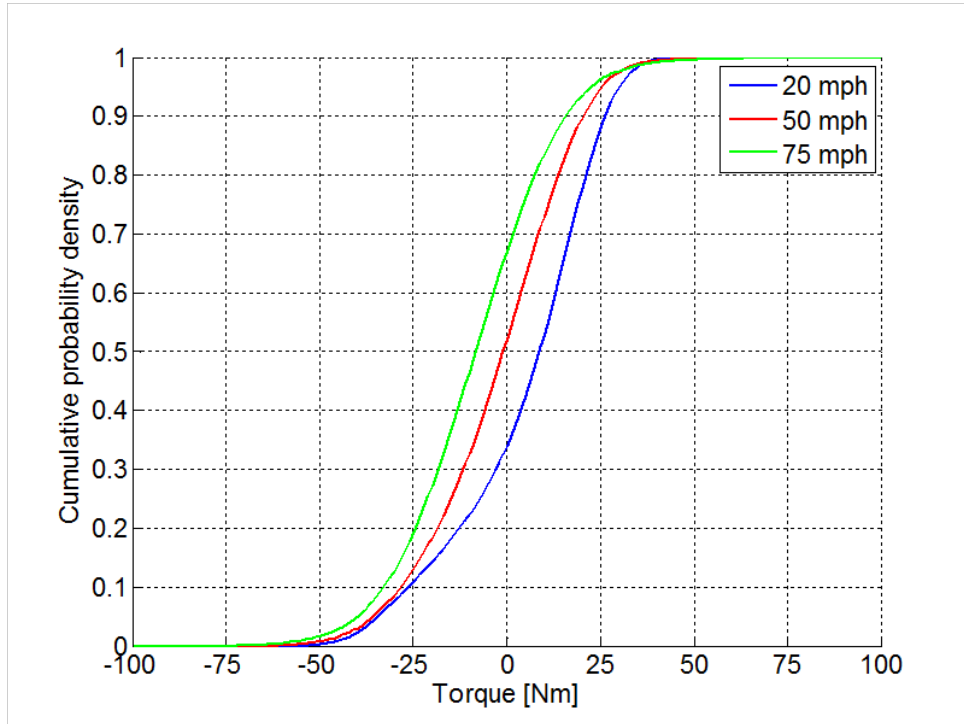
Figure 4.30: Cumulative probability densities for torque within the fibre-glass arm

4.6.5   Verification of results

Prior to installation, the instrumentation used in the field trials underwent thorough laboratory testing. This was in the form of static tests carried out on the bench, as well as dynamic tests performed using the spinning rail rig housed at the University of Birmingham. Despite this, it is still beneficial to be able to draw comparisons between the trial results and those obtained using other systems.

Atkins: South Eastern Trains Third Rail Height Survey 2004

In 2004, Atkins were commissioned to undertake a third rail height survey with the aim of identifying areas where the third rail was operating outside of its specification. The instrumentation in the survey consisted of drawstring displacement transducers and a voltage transducer. The system made no measurement of the contact force and was undertaken several years prior to the

work described above. However, there is still value in comparing the height of the conductor rail recorded by the two systems.

Figure Figure 4.31 is an example output taken from the Atkins report. It shows the heights of the conductor rail on both sides of the vehicle in blue and green. The result shown in red is representative of the conductor rail voltage. The section shown is part of the line between Birchington on Sea and Herne Bay. The system measures a down-stop height when there is no third rail present, this seems to be at 60 mm which is not representative of the down-stop calibration values observed during the trials described here. The report claims this to nominally be 57 mm, again not the value observed during these trials but this is a number that has been modified over the years. The report describes the section between 23400 m and 23850 m as "low areas", however, the sections outside this region can be considered normal.
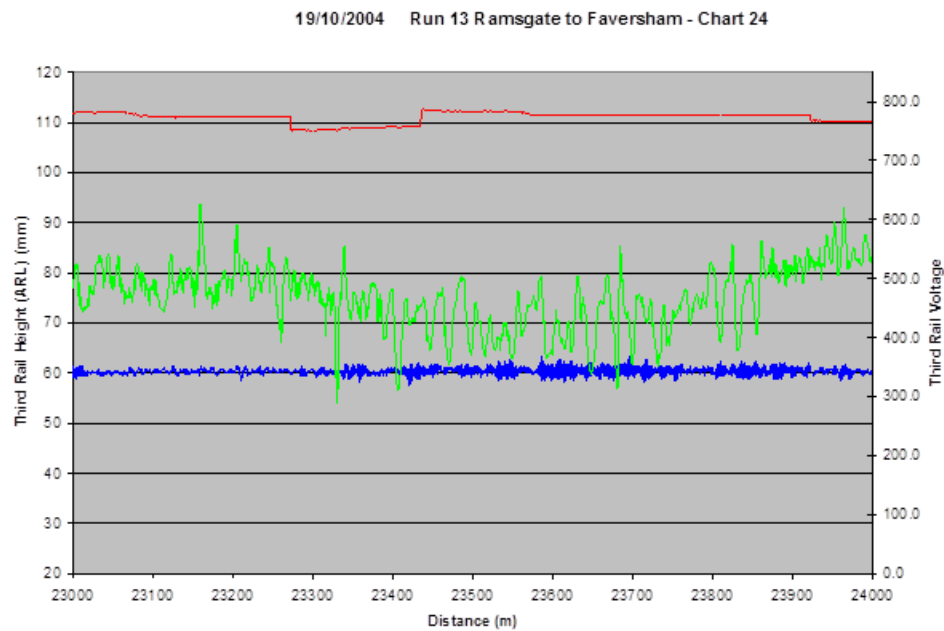
Figure 4.31: Conductor rail height a.r.l. measurement taken from Atkins report

Although the data displayed in Figure 4.31 was obtained at only 40 Hz, and so significant smoothing will have occurred, it is possible to draw comparisons between the results presented

earlier and those from Atkins. Outside of the "low areas" the drawstring measurement system reports the conductor rail height to consistently be between 73 mm and 83 mm (±5 mm). No speed is explicitly stated for the Atkins work; however the system assumes that the shoe is in contact with the rail so presumably the vehicle operated at a low speed. Comparing the Atkins results with those presented here for 20 mph, the results are broadly equivalent if the high frequency components of the author's data are removed or smoothed away. This is encouraging as it goes towards verification of the results.

Mutual agreement of multiple sensors

Another method that can be used to increase confidence in the recorded results is to be able to verify one measurement with another result obtained using different sensors. While not perfect as the sensors are not managed by independent systems, this method can be used to increase confidence in the transducer components of the system.

It has already been shown that the results from a pair of strain gauge bridges can, with appropriate processing, give results for both contact force and torque. During the discussion of the torque results, consideration was also given to the point on the shoe through which the force was applied (heel or toe). Taking this one stage further, it is possible to imagine a single point on the bottom of the shoe through which any given force is applied, and that this would be reflected in the recorded torque. Hence, combining the force and torque results, it should be possible to come to an approximation for this longitudinal "contact point". An example of this is shown in Figure 4.32.
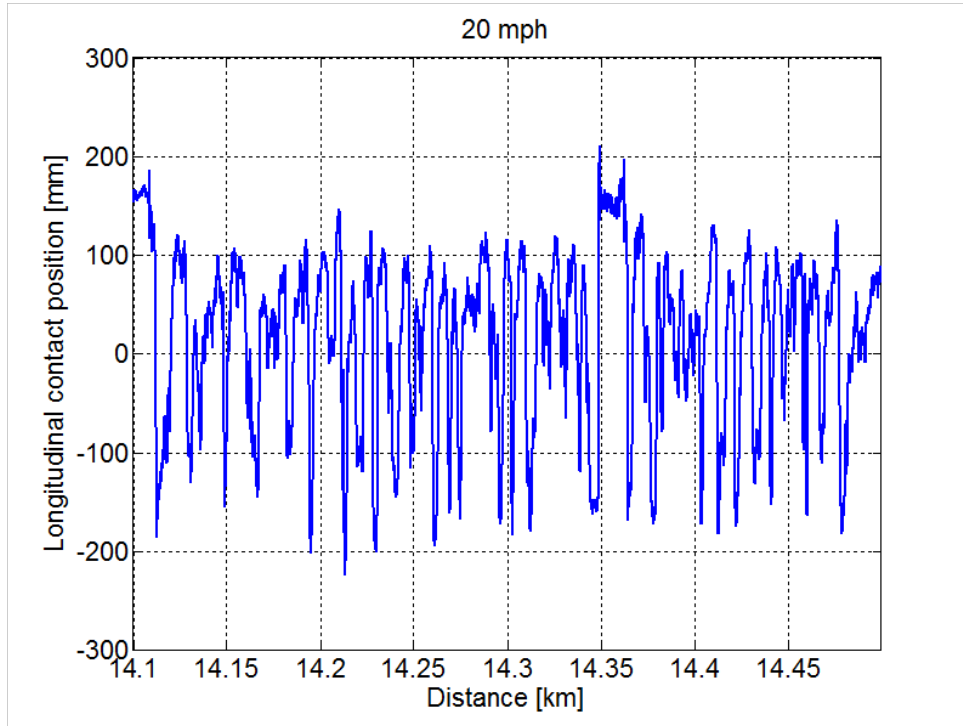
Figure 4.32: Estimated longitudinal contact point

The other main sensor involved in the instrumentation was the displacement transducer. The displacement transducer was sampled at 750 Hz which at a speed of 20 mph corresponds to approximately once every 10 mm. By considering the difference between each measurement point, it is possibly to identify the vertical displacement of the conductor shoe (relative to the bogie) corresponding to each 10 mm of track along which the vehicle moved. The combination of vertical and horizontal displacements can therefore be used to provide an approximation of the slope of the conductor rail relative to the running rails. This approximation is clearly very dependent on the consistency of vehicle speed (in the absence of a tacho signal) and is smoothed over the length of the shoe (300 mm).

The approximation for the longitudinal contact position is obtained from the strain gauge measurements. The approximation is derived from the torque results and is based on the extreme values of torque corresponding to the entire load being transferred through either tip of the

conductor shoe (300 mm apart), and the assumption of a linear relationship existing between the two extremes. Using this, it is possible to estimate a theoretical longitudinal contact point between the tips of the shoe based on the observed torque. It is worth remembering that this is a theoretical point of contact which would in practise not be possible due to the surface dimensions of the conductor shoe. As the approximation for the conductor rail slope is obtained from the displacement transducer. These measurements are therefore completely independent prior to the analogue to digital conversion, the microprocessor, and the signal processing chain. The two results can, with a little scaling, be combined to produce a graph such as that shown in Figure 4.33. The longitudinal contact position is shown in blue, and a scaled version of the slope in red. The section in the middle represents the shoe hanging on the down-stop and so it and the points of entry / exit to it are not comparable across the two technologies, however, the remainder of the figure shows significant coherence between the results taken from the different transducers.
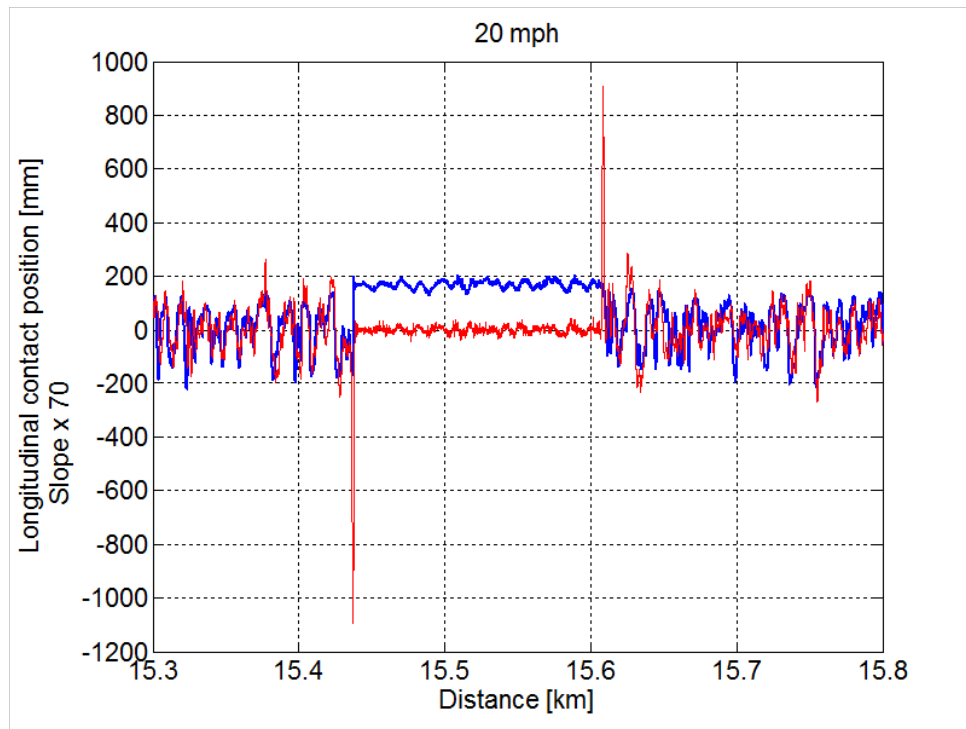
Figure 4.33: Estimated longitudinal contact point and conductor rail slope

4.7 Conclusions

This chapter has presented work carried out to instrument a Class 375 railway vehicle. The instrumentation recorded displacements of the shoegear and forces between the conductor shoe and the third rail. Similar to the system described in the first case study, the instrumentation was based on a general architecture of control and conditioning local to the transducer, with visualisation and recording carried out separately. This project, however, required the instrumentation to be far more ruggedised due to the vibrational nature of the operating environment. The systems developed were capable of operating at line speeds and were used for a single day of dedicated trials. The work described here has also been documented in a journal paper [107] for which the author of this thesis is the first author.

The displacement results generated by the system have been used to assess the quality of the third rail infrastructure throughout the section of line over which the trials were performed. This example of the vehicle acting as a measurement tool for the infrastructure is particularly important as it leads to a measurement of the infrastructure in it's regularly loaded condition. Dedicated measurements, such as those from the New Measurement Train, are not subject to the standard loading patterns generated by the vehicles that regularly use the network and as such may not be truly representative. Generally speaking, the infrastructure performed well with a few notable exceptions at specific points such as ramp ends which are beyond the scope of a general performance study anyway. The results from the displacement instrumentation have been compared to similar results from previous work and have been deemed acceptable given the timeframe between the two sets of measurements.

The instrumented vehicle has also been used as a device to monitor its own operation, specifically the operation of the shoegear components. The force measurements have been considered over both the whole trial distance and at specific areas of interest such as ramps and gapped sections.

This was the original focus of the work as very little was previously known about the dynamic performance of the shoegear. The results have shown that for the majority of the time and in general operation, performance is within the limits of the static specification. The equipment was not designed to measure impact forces at ramp ends so the results from these sections are only indicative. The indications are that the results are outside of the static specification, although this is not surprising. Furthermore, some areas of particular interest have been identified in these regions.

Investigation has also been undertaken into the relationship between conductor rail height and exhibited contact force. Rather than a direct contextualisation of one signal by the other, this has been an exercise in using the signals together to better understand the operation of the instrumentated system as a whole and has has lead to the identification of significant hysteresis within the shoegear system. Combination of this investigation with the results obtained from the force measurements has lead to the development of the theory that the flexibility of the fibreglass arm accommodates high frequency small magnitude displacements, while the torsion spring at the heart of the shoegear manages lower frequency, larger magnitude components of any displacement.

Beyond the mandated objectives of the project, work has also been undertaken to consider the torque within the fibreglass arm. Of interest in itself, this work is particularly valuable when considering the non-general case of conductor rail / shoegear interactions caused by impacts. The shock inputs "felt" by the shoegear at points of impact, such as ramp ends around gapped sections, are particularly clearly illustrated in the torque results. Despite using the same (albeit additional) sensors as the force measurements, the alternative analysis to establish torque allows a better understanding of conductor shoe behaviour in these critical zones which may be used to inform maintenance decisions for those critical zones in the future. Furthermore, given that

the force measurement is only indicative at these locations, the indicative nature of the torque measurement is equally valid.

The instrumentation has shown the shoegear to be reasonably consistent over the three trial speeds. This may allow it to be used in the development of future instrumentation systems. Some additional, higher frequency, information is included as the speed increases, and the shoe is more likely to become separated from the conductor rail. This has implications for the displacement measurement if it is considered to be a measurement of the conductor rail rather than the shoe. When considering the behaviour of the shoe, interesting comparisons can be made using the results from the different speeds over features such as ramp ends. One application of such comparisons is a project being undertaken with another TOC to instrument a pair of shoegear to identify sections of the infrastructure, particularly around gapped sections, which require specific inspection and maintenance. The system being developed for this project will be fixed to an in-service vehicle and so must operate at line speeds without decreasing in performance.

The work presented in this case study has shown that through a comparably minor extension of the sensor set it is possible to enhance a set of instrumentation such that it can become "self verifying"; i.e. data from one transducer can be used to give confidence to results obtained with a second unit of a different variety. This comparably minior modification, combined with the instrumentation architecture used in this project means that no substantial architectural changes are required. While this is beneficial in terms of effort, such a process is not without flaws as systematic errors would not be identified. It can, however, still provide an increase in the confidence associated with the data, compared to a single sensor measurement.

Despite the ruggedisation performed for this project, extending the system to operate during service conditions is not trivial and modifications to further increase system robustness are being considered. Some of these modifications are practical, such as implementing air curtains to help

maintain the laser displacement transducer. Others involve a more academic approach of going back to the extended dataset to investigate a possible reduction in sensors, or a change in the system architecture to reduce the number of single points of failure.

CHAPTER 5

CASE STUDY 3 - INERTIAL MEASUREMENT

5.1 Introduction

This case study describes a project in which the alignment of railway track was monitored from a regular service railway vehicle. The technologies used in the work are architecturally similar to those used in the first two case studies, but different in terms of transducer selection and installation / operating environment. The work extends instrumentation systems previously installed as part of the Enhanced Rail Contribution through Improved Reliability (ERCIR) project, which was a Department for Transport funded LINK project involving a variety of industrial partners, the University of Birmingham and Loughborough University. A clear evolution of this instrumentation is visible and directions in which such instrumentation systems might be extended further are discussed.

In addition to these objectives, this case study is included in this thesis partly to introduce an alternative variety of transducer that could potentially be interrogated in order to contextualise other data streams, partly to extend the concept introduced in the previous case study whereby multiple transducers can be used to measure the same physical effect, and also to introduce the concept of data volumes and the management required for the expanded data repositories that

modern instrumentation systems are starting to require.

5.2 Motivation

From the very first railways back in the early 19th century, through to the modern, high speed, lines of today; the quality and the alignment of the railway track used remain the most significant factors in ensuring vehicle safety and ride quality. With good track alignment, a metal wheel on a metal track can provide a safe and reliable interface, smooth enough that vehicle suspension can easily produce an acceptable ride quality. The challenge lies in maintaining the railway and hence maintaining that interface.

Many systems exist to maintain the railway track to the required standards. Faults or wear identified within the rails lead to rails being repaired using techniques such as grinding, or even removed and replaced. Many track faults, however, are not directly related to the quality of the rails, but instead to the substructure that the rails are mounted on. Remedial work for substructure faults may be as minor as simply redistributing some ballast, or as major as tamping or even re-laying whole sections of railway line.

Before it is possible to repair track faults, however, they must be located. This is the purview of track inspection technologies and techniques. In the UK, much of this track inspection is performed manually with Network Rail employees walking the lines to visually inspect them. There are many advantages to this manual approach. Track workers can be deployed almost continuously and in large numbers thereby giving good network coverage. They can be trained to identify a large number of faults; and they can become familiar with particular sections of the railway and so can identify changes in track condition. They can also use their initiative. This, however, is also one of the disadvantages of manual inspection. Writing in Modern Railways [11], Andrew McNaughton highlights the inconsistencies that can arise between track workers'

modern instrumentation systems are starting to require.

5.2 Motivation

From the very first railways back in the early 19th century, through to the modern, high speed, lines of today; the quality and the alignment of the railway track used remain the most significant factors in ensuring vehicle safety and ride quality. With good track alignment, a metal wheel on a metal track can provide a safe and reliable interface, smooth enough that vehicle suspension can easily produce an acceptable ride quality. The challenge lies in maintaining the railway and hence maintaining that interface.

Many systems exist to maintain the railway track to the required standards. Faults or wear identified within the rails lead to rails being repaired using techniques such as grinding, or even removed and replaced. Many track faults, however, are not directly related to the quality of the rails, but instead to the substructure that the rails are mounted on. Remedial work for substructure faults may be as minor as simply redistributing some ballast, or as major as tamping or even re-laying whole sections of railway line.

Before it is possible to repair track faults, however, they must be located. This is the purview of track inspection technologies and techniques. In the UK, much of this track inspection is performed manually with Network Rail employees walking the lines to visually inspect them. There are many advantages to this manual approach. Track workers can be deployed almost continuously and in large numbers thereby giving good network coverage. They can be trained to identify a large number of faults; and they can become familiar with particular sections of the railway and so can identify changes in track condition. They can also use their initiative. This, however, is also one of the disadvantages of manual inspection. Writing in Modern Railways [11], Andrew McNaughton highlights the inconsistencies that can arise between track workers'

interpretations of what they discover. He explains that with automated inspection techniques, variability is reduced and so the thresholds at which intervention occurs can be set higher.

Using dedicated vehicles to monitor the condition of the railway is not a new idea. Modern inspection vehicles have become reasonably complex, possessing extensive combinations of sensing techniques [19]. Inspection vehicles are, however, expensive and so tend to be low quantity assets which must be circulated around the network. This reduces inspection frequency. Also, track inspection vehicles do not possess the same dynamics as the vehicles that use the track regularly and so do not "stress" the track in the same way while measuring it. Furthermore, as described in [17], dedicated inspection vehicles tend to operate rather slowly which can have an impact on the timetables that can be operated during measurement periods.

There have been calls in the literature [45], [24], [61], particularly from the United States and Japan, for track measurement systems to be fitted to in-service vehicles. This would allow more regular inspection coverage while maintaining the dynamics of measurement and not disrupting the timetable.

The key to employing in-service measurements would be to develop a suitable suite of instrumentation. Initially this involves over-instrumenting a vehicle to allow a comparison of the sensors to be made and thus conclusions about their relative merits to be drawn. Previous work conducted by the University of Birmingham has investigated the appropriate sensor set for inertial measurements performed from the bogie of a vehicle. Further data, particularly around features such as Switches and Crossings (S&C) is now required to develop algorithms for automatic fault detection.

As part of the conductor shoe instrumentation described in the previous chapter, some elements of an inertial measurement system were fitted to a Class 375 railway vehicle. This instrumentation set was augmented to a full Inertial Measurement Unit (IMU) as part of this ongoing work

towards in-service inspection.

5.3 Instrumentation

5.3.1 Sensing technologies

The complexity of inertial measurement instrumentation is often related to the level of data that is required of the system. Some, very simple, systems barely qualify as inertial measurement, employing simple tilt switches to identify movement or vibration. Beyond a threshold inherent to the sensor, these systems are not quantitative. Other, more complex, systems employ displacement transducers, either mechanical or non-contact (such as those described in the previous chapter) to identify movement of one component of a system relative to another. These are quantitative sensors but are limited to relative measurements between two points rather than absolute displacement results, if the whole system is affected the transducer won't register anything. True inertial measurement systems make use of accelerometers, gyroscopes or a combination of both to track an object's behaviour through time.

The inertial measurement systems discussed in this chapter, make use of Micro ElectroMechancial Systems (MEMS) accelerometers. These are particularly suited to the high vibration and shock environments found on the bogie while providing sufficient precision for the measurements and the rest of the signal processing chain. For projects specifically concerned with inertial measurements, Vibrating Structure Gyroscopes (VSG) are used. These provide additional accuracy compared to more traditional counterparts, and are sufficiently accurate while being substantially more robust than Fibre Optic Gyroscopes (FOG).

5.3.2 Overview

This case study describes development work for, and outputs from, an inertial measurement system. The system is designed to be fitted to a railway vehicle which is then used as a tool of measurement to record vertical and lateral track irregularity. This is essentially a metric of track quality. Using the vehicle as a tool of measurement in this way means that the length between the points of the vehicle in contact with the track is significant. Mounting the system on bogie of the vehicle limits the contact separation to approximately 2.5 m which is the shortest possible solution without fitting equipment to the individual axles of the vehicle. The instrumentation architecture selected for this work therefore had to consider the issue of mounting transducers in the inhospitable bogie environment, while still using standard (PC based) recording equipment. This lead to the development of another split site architecture with transducers separated from data recording. As part of the evolution from the ERCIR project, the data acquisition components of the system were co-localised with the transducers themselves.

Previous versions

The work being described in this chapter is an extension of a project carried out for the Department for Transport under the ERCIR banner. In that project, an inertial measurement system was fitted to a metro vehicle and later to a diesel multiple unit. The first inertial measurement unit was typical of academically driven work. The system aimed to measure track quality from a railway vehicle. To do this, numerous measurements were taken. The instrumentation consisted of a full six degree inertial measurement unit (three axes of accelerometers and three of gyroscopes) attached to the bogie, combined with separate axlebox mounted accelerometers and displacement transducers linking the instrumentation sites together [105]. A block diagram of the instrumentation can be seen in Figure 5.1.
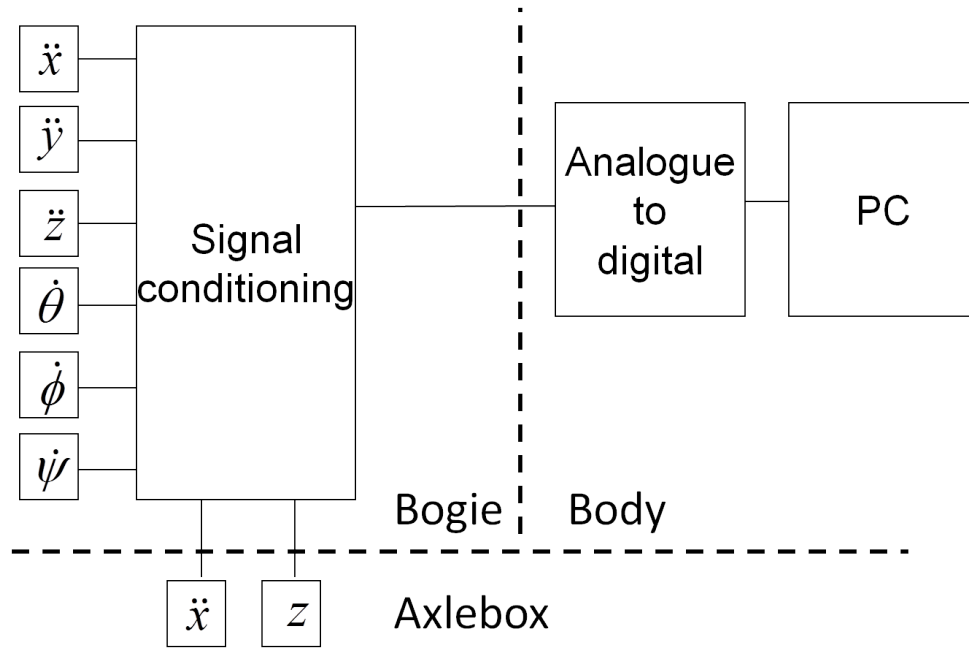
Figure 5.1: Version one of the inertial measurement system

The data from the instrumentation was analysed and the merits of the various transducers established [108]. Axlebox accelerometers were identified as the best transducers for identifying short wavelength effects [109] but were notably difficult to install and maintain. The accelerometers and gyroscopes attached to the bogie proved to be more suitable for general rail inspection.

A second set of instrumentation was then developed as part of the same project. This instrumentation focused on the bogie mounted equipment and did not include the axlebox or displacement transducers. The project was focused on vertical and lateral alignment but the inclusion of the longitudinal accelerometer and roll rate gyroscope involved little overheads and so they were included. This reduced instrumentation set is shown as the top two sections of Figure 5.1. As a testament to this academic approach, the data from the roll rate gyroscope has subsequently been used.

Both systems employed analogue signal conditioning at the transducer but then used multicore cabling to relay the signals to the vehicle body for analogue to digital conversion.

The extension to the work undertaken as part of the ERCIR project involved the installation of inertial measurement units as part of other instrumentation systems. Building on the work done in the earlier projects, the data from the inertial measurement units were used as part of a positioning and alignment system for the data as well as for inertial measurement and infrastructure condition monitoring in its own right.

One of the extensions involved fitting a system to the same Class 375 railway vehicle that was instrumented for conductor shoe / shoegear performance in the previous chapter. The inertial measurement unit was mounted in the box on the side of the bogie, where the microprocessor, signal conditioning and strain gauge amplifiers for the conductor shoe instrumentation were located. The inertial measurement transducers are highlighted in Figure 5.2.
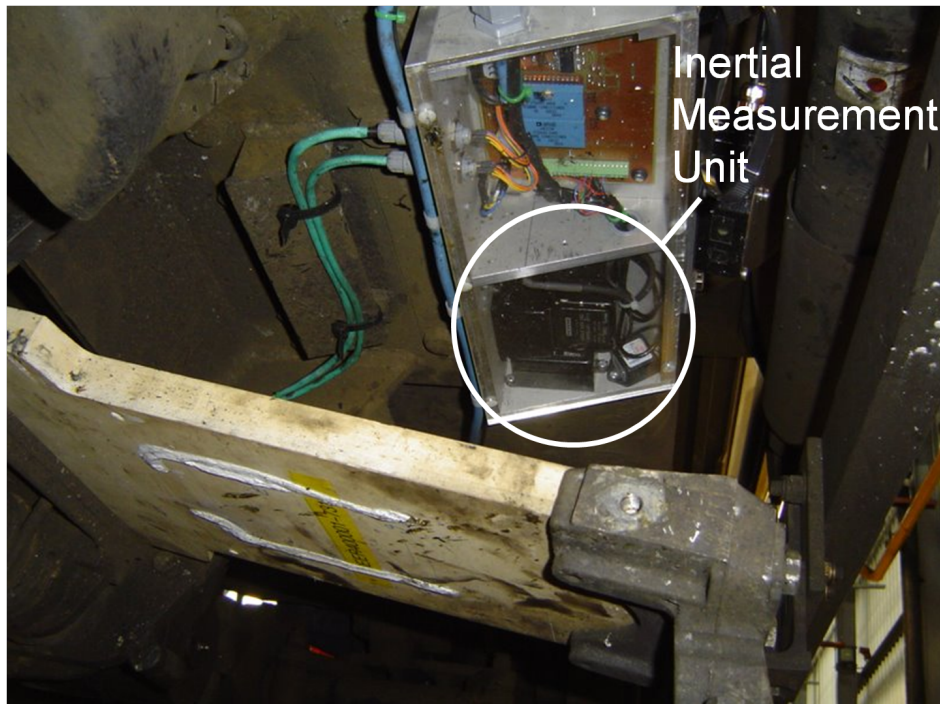
Figure 5.2: Inertial measurement unit fitted to Class 375

The system comprised three axes of acceleration measurement using a single tri-axis capacitive MEMS accelerometer and three axes of rotational measurement using three individual VSGs. Gyroscopes were used in all three axes because of the overlap between the two projects. The inertial measurement work required that the pitch and yaw be recorded to be processed as part of an investigation into vertical and lateral irregularity. The conductor shoe work required that the yaw be recorded to assist with vehicle positioning, and that the roll be recorded as part of the geometry calculations involved in identifying the conductor rail height relative to the running rails. This overlap is shown in Figure 5.3. The tri-axis accelerometer was selected because of its performance in the vertical and lateral directions, and precision of relative axes that a single package provides. The longitudinal result was recorded primarily as additional information. Subsequently, the results from the longitudinal accelerometer have been used to augment a poor quality positioning signal also obtained as part of the instrumentation.
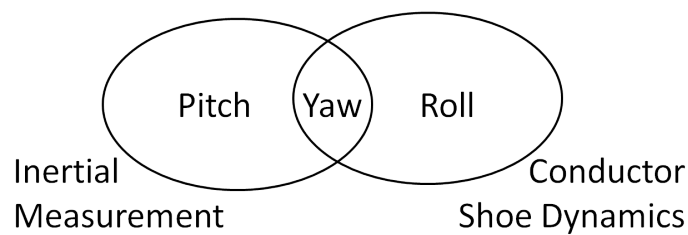
Figure 5.3: Gyroscope use in inertial measurement and conductor shoe dynamics work

The inertial transducers selected were all analogue and so signal conditioning and analogue to digital conversion was used prior to the results being stored on a computer. Unlike the work undertaken as part of the ERCIR project, in which a diesel multiple unit was instrumented, the Class 375 instrumentation was at particular risk from electrical interference and so the connections to the transducers were kept as short as possible. Instead of a multicore cable to the vehicle body, the analogue to digital conversion was performed using a microprocessor located on the bogie and the data transferred to the PC using a fibre optic link. A block diagram

of this, more advanced, system can be seen in Figure 5.4.
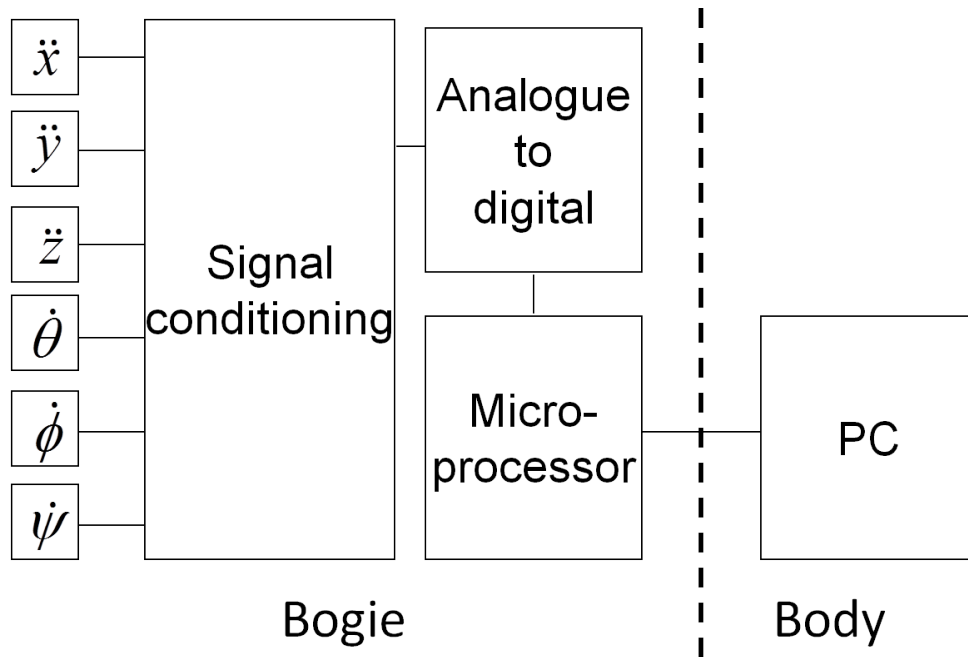
Figure 5.4: System overview diagram for the inertial measurement instrumentation

5.3.3 Electronics

The accelerometer selected for the instrumentation was a CXL10TG3 from Crossbow Technology. The TG10 is part of crossbow's tactical grade accelerometer range. It is a tri-axis, ± 10 g capacitive MEMS accelerometer with very low noise, good stability, and an operating bandwidth right down to DC. The tri-axis package ensures that the three axes are precisely aligned for minimal cross-pickup and also reduces the required footprint of the transducers. The 10 g range was selected based on predicted bogie movement in accordance with observations from previous instrumentation work.

The gyroscopes used were 299640-100 units from BAE Systems' VSG collection. They are single axis VSGs with a range of $\pm 50^{\circ}/s$ and very low noise. BAE use them in automated gunfire and missile systems as well as Unmanned Aerial Vehicles (UAV). The rotational rate of a railway

vehicle is largely based on vertical or lateral curvature and track alignment. A railway vehicle has significant mass, and therefore inertia, and is therefore unlikely to rotate particularly quickly around any of its axes. The $\pm 50°/s$ gyroscopes selected are slow by many standards, but fast enough for the predicted motion of the bogie, and therefore have greater resolution in this region.

Both of the above sensor types have analogue outputs and so need analogue to digital conversion. Using the same systems as those used for the conductor shoe sensors, the signals from the inertial measurement equipment were subjected to some basic anti-aliasing filters before being sampled by an LTC1859 analogue to digital converter at a rate of 750 Hz. This frequency was substantially lower than that used in the previous ERICR project but relates to the interrogation of acceleration and rotation sensors fitted to the bogie of a railway vehcile. This corresponds to a large inertia and so the acceleration and rotation of the bogie should be sufficiently small to be easily identified by such a sampling regime. Furthermore it meant that the signals could be interlaced with those from the conductor shoe instrumentation, thereby removing the need for additional electronics.

The analogue to digital converters were managed and interrogated by a PIC 18F2680 micropro-cessor. A single SPI bus was used to communicate with both controllers. Separate chip select lines were used to manage the bus. The microprocessor used an external crystal oscillator to control its operating frequency and hence to manage the timings for the analogue to digital conversions. Once samples had been taken, the microprocessor transmitted serial versions of the signals through a fibre optic media converter to the PC. The system used an RS232 protocol but implemented Transistor to Transistor Logic (TTL), rather than the associated voltage levels.

5.3.4 Firmware

As the conductor shoe and inertial measurement projects were always closely linked, the firmware modifications to combine them were reasonably straightforward. The interlacing between high and low speed signals was simply a case of modifying the lookup table that managed the analogue to digital converters, and implementing additional logic to control which of the two units should be activated. The framing messages, which were periodically inserted into the data to allow recovery from slipped bytes or data corruption, were selected to be unique while considering the signals from all of the different transducer types. Finally, the baud rate of the serial link to the PC was selected to accommodate the combined volume of data. This was, perhaps, the most challenging aspect of the modifications as the link speed was already close to the limit for reliable transmission. The selection of sampling rates (750 Hz and 3 kHz) was primarily based on the limitation of the fibre optic serial link. Likewise, the use of a dedicated driver chip rather than MOSFET control was a requirement of the data rate.

5.3.5 Software

The system used a two phase approach, whereby the data was stored onto the computer in a raw format and then converted to accessible data using a post-processing step. This was essentially dictated by the volume of data created by the addition of the inertial measurement channels. Table 5.1 shows the impact of the inertial measurement systems on the dataset. Prior to the inclusion of the inertial measurement unit, the computer was fast enough to decode the serial stream and, more significantly, write the formatted output to a file. With the inertial measurement data included, the system could only perform real time processing of a subset of the data. This was adequate for the real time visualisation which was only updated at 1 Hz.

Table 5.1: Data rates used in the conductor shoe dynamics and inertial mesaurement instrumentation stets

Conductor shoe dynamics		Inertial measurement	
Signal	Data rate [Hz]	Signal	Data rate [Hz]
Strain gauge bridge 1	3000	Longitudinal acceleration	750
Strain gauge bridge 2	3000	Lateral acceleration	750
Displacement	750	Vertical acceleration	750
		Roll rate	750
		Pitch rate	750
		Yaw rate	750
		Temperature	750
Total	6750	Total	5250

Data extraction from the raw data was carried out, along with that of the conductor shoe dyanamics data, using a second piece of software written in C#. Post processing of the inertial measurement data was performed using MATLAB. This post processing is described in detail in section 5.4.

5.3.6 Mechanical construction

The inclusion of the inertial measurement instrumentation into the conductor shoe dynamics setup had a smaller effect on the mechanical construction than might be expected. The signal conditioning electronics had to be expanded and an additional analogue to digital converter included, but beyond that there was very little change to the system. The strain gauge amplifiers, microprocessor and fibre optic media converters remained unchanged. The requirements for encasing electronics and mounting them on the bogie are no more or less stringent than those

for the conductor shoe dynamics transducers. The only difference, in fact, was the volume of equipment being installed. This was accommodated with a second chamber being added to the instrumentation, as shown in Figure 5.2.

The mass added to the system by the inclusion of the inertial measurement transducers is summarised in Table 5.2. The inertial measurement components added approximately 0.5 kg to the instrumentation which including the transducers totalled 4.5 kg. This difference was not ultimately significant in comparison to the 5.5 kg of arc shielding that the instrumentation was replacing.

Table 5.2: Inertial measurement unit transducer masses

Component	Mass [g]	Quantity	Total mass [g]
Accelerometer	110	1	110
Gyroscope	125	3	375
		Total	485

5.4 Data Processing

An ideal railway vehicle operating on an ideal section of perfect track would only move longitudinally along that track. In practice, railway track is not ideal. Flexibility in the substructure and manufacturing tolerances for the system inevitably lead to vertical movement. The need for changes in direction also lead to movement in a lateral direction. Furthermore, railway vehicles are sizeable objects with spacing between the wheels of a bogie and between the bogies of a car. As the front of a vehicle will reach any change in the alignment of the track first, rotations around the lateral and vertical axes are inevitable either on a bogie or vehicle body scale. Rotations around the longitudinal axis are also introduced through the alignment of the two running rails or dynamic effects within the vehicle. It is, therefore, possible to describe a

railway vehicle as having movement in three dimensions: longitudinal, lateral and vertical; with rotations around each of those three axes: roll, pitch and yaw. This is shown in the top section of Figure 5.5.

Measurement Axes

z

x

y

Yaw, Ψ, Psi

Roll, θ, Theta

Pitch, Φ, Phi

Travel along track, distance, s / time, t

Sensors

Vertical acceleration

Lateral acceleration

Yaw rate

Longitudinal acceleration

Roll rate

Pitch rate

Tacho
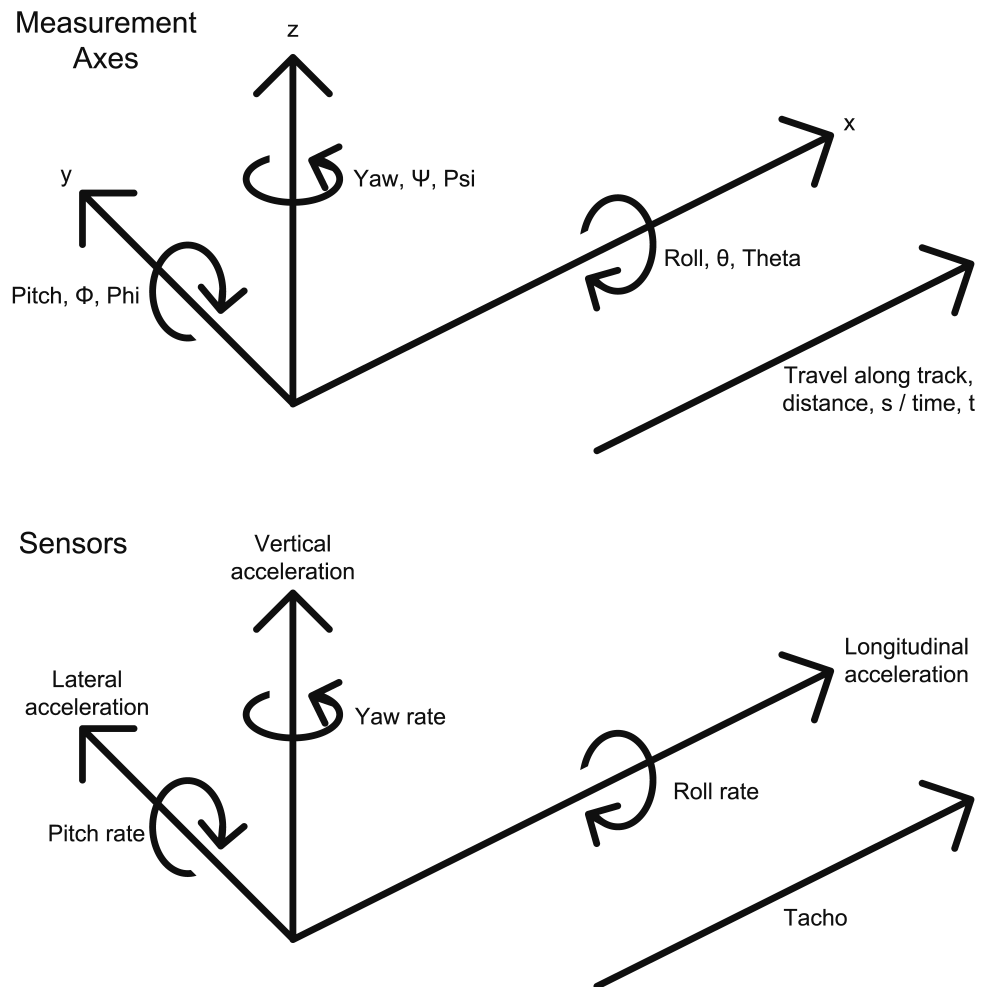
Figure 5.5: Inertial measurement dimensions and transducers

The inertial measurement unit fitted as part of the work described in this case study records accelerations in three orthogonal directions and rates of rotation around those three axes. It is therefore referred to as a six degree of freedom inertial measurement unit. While the sensors used correlate to the three dimensions of movement and the three directions of rotation around those axes, the sensors do not measure the physical characteristics directly. Instead of measuring movement in the x, y and z directions, accelerometers are used to measure accelerations in these

directions. Likewise, instead of measuring roll, pitch or yaw directly. Rate gyroscopes are used to identify the rates of those three rotations. This contrast is illustrated in the differences between the upper and lower sections of Figure 5.5.

Signal processing is used to identify the physical characteristics from the recorded data. The changes in this processed data correspond to the alignment of the track in a particular direction. It is standard practise in the railway industry to filter track alignment data over 35 m and 70 m wavelengths. This means that a filter is applied to the output of the processing which eliminates components of the signal with wavelengths longer than the 35 m or 70 m of interest. Everything shorter than the particular wavelength is retained. This is done, for example, so that track alignment features can be separated from substructure or geographical effects. Higher speed railways, such as those in Europe, occasionally also require alignment data to be filtered over 140 m. The railway industry is particularly interested in the vertical and lateral irregularity of the track as these have direct implications for ride quality and system wear.

5.4.1 Accelerometers

The literature contains many examples of accelerometers being used to quantify ride quality, for example [29], [44], [46], [51] and [56]. Particular consideration is also given to the location on the vehicle at which the accelerometer is mounted in relation to the wavelengths being considered [25]. Measurement of acceleration is ideally suited to ride quality assessment where the accelerations being felt by the passenger are the final output. As a mechanism for directly identifying track quality, however, measurement of acceleration is less suitable due to the processing required to identify displacement from acceleration. The relationships between acceleration, velocity and position are well understood and the processing chain works well in one direction. For example, vertical position recorded throughout a longitudinal journey with respect to time can

be differentiated to give vertical velocity and then again for vertical acceleration. The effects of noise in the original displacement measurement are minimised in this process as differentiation only magnifies high frequency error and a railway vehicle is sufficiently massive that this can be ignored. Conversion in the other direction, however, requires integration which magnifies low frequency error and so is particularly susceptible to low frequency experimental noise [110].

5.4.2 Gyroscopes

It is also possible to identify vertical irregularity using a combination of a pitch rate gyroscope and a speed signal, both sampled in the time domain. Lateral irregularity requires similar processing based on a yaw rate gyroscope. Both of these techniques have been used extensively by Dr. Paul Weston at the University of Birmingham and are used here under his guidance. The processing chains required are described in detail in [40] and [66] respectively. With consideration to the vertical case, illustrated by Figure 5.6:
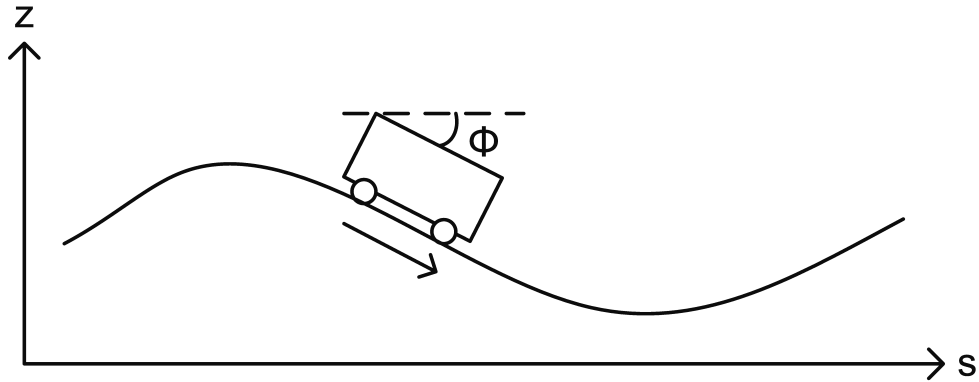
Figure 5.6: Dimensions and sensors for vertical irregularity

Vertical curvature (K) (the reciprocal of curve radius (R)) is defined as the rate of change of pitch ($\dot{\phi}$) with respect to the distance (s) travelled along the track. The pitch rate gyroscope reports the rate of change of pitch ($\dot{\phi}$) with respect to time (t). The vehicle tacho signal reports vehicle speed (v) directly; this can also be described as the rate of change of position along the

158

track (s) with respect to time (t). Hence:

$$\text{Vertical curvature } (K) = \frac{d\phi}{ds} \tag{5.1}$$

$$\text{Pitch rate } (\dot{\phi}) = \frac{d\phi}{dt} \tag{5.2}$$

$$\text{Tacho } (v) = \frac{ds}{dt} \tag{5.3}$$

$$\text{Verticle curvature } (k) = \frac{d\phi}{ds} = \frac{\frac{d\phi}{dt}}{\frac{ds}{dt}} = \frac{\dot{\phi}}{v} = \frac{\text{Pitch rate}}{\text{Tacho}} \tag{5.4}$$

As described in equation 2 in [40].

Equally, for a similar curve, a vertically sensing accelerometer mounted on a vehicle would report acceleration $(\ddot{z})$ with respect to time (t). The primary component of this acceleration would be due to centripetal acceleration, thus:

$$\ddot{z} = \frac{v^2}{R} = v^2 \times K \tag{5.5}$$

hence:

$$K = \frac{\ddot{z}}{v^2} \tag{5.6}$$

As described in equation 1 in [40].

Furthermore, equation 3 in [40] then states that:

$$\hat{z} = \int \int K \; ds \; ds \tag{5.7}$$

which implies that:

$$\hat{z} = \int \int \frac{\ddot{z}}{v^2} \tag{5.8}$$

and:

$$\hat{z} = \int \int \frac{\dot{\phi}}{v} = \int \int \frac{\text{Pitch rate}}{\text{Tacho}} \, ds \, ds \tag{5.9}$$

A low pass filter is then applied to present results over either 35 m or 70 m as per the standard alignment processing and presentation techniques described earlier.

The paper in which this processing was originally presented [40] describes a number of different methods for the estimation of vertical irregularity based on measurements from both accelerometers and gyroscopes. The different approaches vary in complexity and this is visible in the number of assumptions made. The processing chain presented here makes two key assumptions. Firstly that phi is small, i.e. $\phi \approx sin(\phi)$, and secondly that the bogie remains parallel to the track. Both of these assumptions relate to the fact that accelerometers not precisely aligned with their axis of measurement will also record components of the dimension into which they have been misaligned. For example, if the vehicle was to traverse a particularly steep section of track, or pitch significantly while on a flat section, the vertically sensing accelerometer would also record a component of the longitudinal acceleration. This is further complicated by the presence of gravity in vertical accelerations. Technically, the same effect may occur due to a rotation around the roll axis although this is not mentioned in the paper. Obviously, the same pitching effect would have a negative impact on calculations based on the output of a pitch rate sensor such as the gyroscopes being discussed. The entire concept of using the bogie as a tool of measurement for the track requires the bogie to follow that track with some degree of accuracy. In the railway environment, gradients are rarely extreme enough for the first of these assumptions to be significant. The validity of the second assumption would be dependent on the

160

condition of the track and the specification of the vehicle being used. Dr Weston has confirmed that in all the trials in which he has used the technique the assumption has remained valid.

Another consideration when performing processing based on inertial measurement systems is the stability of the speed of the vehicle. Due to the dynamic nature of the suspension of a bogie, accelerating and decelerating a vehicle increases the pitching effect which brings in the additional complexities described above. Discussions with Dr Weston have confirmed that, in his experience, this processing becomes more consistent the closer the vehicle is to operating at constant speed.

Both Equation 5.8 and Equation 5.9 use a double integration in the final stages of processing to generate the estimate for height (or irregularity). As mentioned previously, integration is a noisy process which amplifies low frequency noise. This double integration is, however, essential to the processing chain. In the case of the solution derived from the gyroscope, this is a double integration is of a term that involves a division by speed. In the case of the accelerometer processing the division is by speed squared. When the speed of the vehicle is low, low frequency noise in the speed signal is significant. This noise is amplified in the integration process. In the case of the accelerometer, the square of the noise is amplified. This is why gyroscope based solutions outperform accelerometer based ones at lower vehicle speeds.

5.5 Trials

As described in the previous chapter, the conductor shoe and inertial measurement instrumentation were used for a single day of trials which took place between Ramsgate and Faversham.

The train path that the trials took, involved travelling from Ramsgate to Margate and then repeatedly between Margate and Faversham, before returning to Ramsgate. This path included

three passes at constant speed between Birchington on Sea and Chestfield and Swalecliffe. The trial section is shown highlighted in red on Figure 4.9.

In addition to the highlighted trial route, Figure 4.9 also shows the switch and crossing work present in the trial section. There are three significant points to note from the diagram:

1. There are only five pieces of relevant switch and crossing work in the trial section.

2. All of the switch and crossing work encountered was in the normal configuration.

3. All of the switch and crossing work was encountered in the trailing direction.

5.6   Results

The dataset obtained from the instrumentation was processed using the techniques described in Section 5.4. A selection of the results for vertical an lateral irregularity, using data obtained from the pitch and yaw rate gyroscopes are show in the following sections.

5.6.1   Vertical alignment

Processing the data from the pitch rate gyro, including filtering over 70 m, produces an estimate for vertical irregularity. This processing has been applied to the data taken during all three trial runs.

Figure 5.7 shows the result obtained for vertical irregularity over a 4 km subsection of the 16 km trial route. Despite the instrumentation being configured for S&C work, and the lack of S&C in the route, some significant irregularity is shown to be present. The most obvious irregularity occurs just before 22.5 km. Other areas of poor alignment are around 21.5 km, 20.75 km and just after 23 km. The general alignment quality at the start of the section is noticeably worse than

at the end, hence, the feature just after 23 km would likely be more noticeable to a passenger than those around 20.75 km.
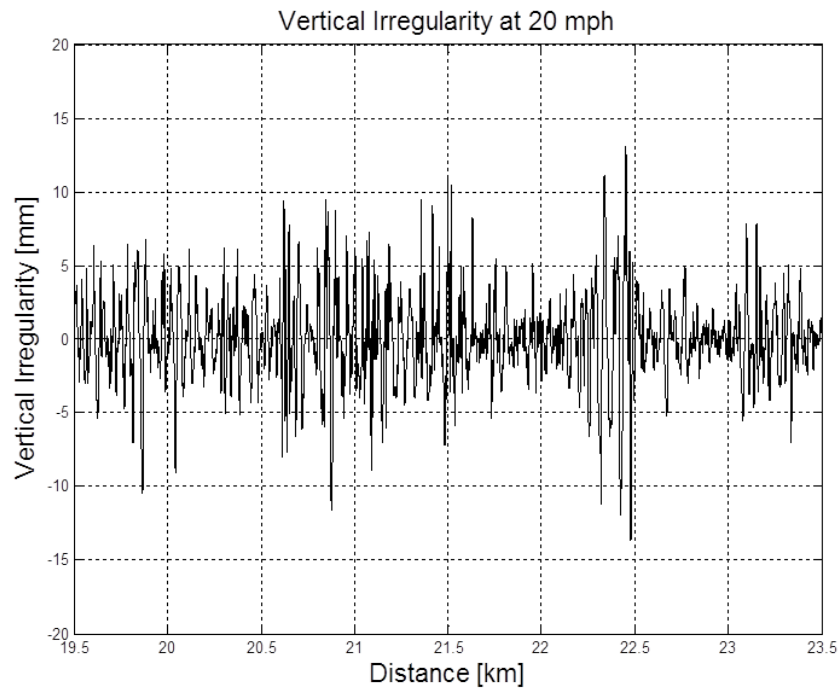
Figure 5.7: Vertical alignment recorded at 20 mph

The result for the same section of track, taken when the vehicle was travelling at 75 mph, is shown in Figure 5.8. From the figure, one can identify the same four areas as significant although the general background alignment is slightly worse so the earlier points could easily be missed. Generally, the result from the 20 mph and 75 mph test runs are reasonably consistent. The result from the 50 mph test displayed the same characteristics with the result falling between the other two.
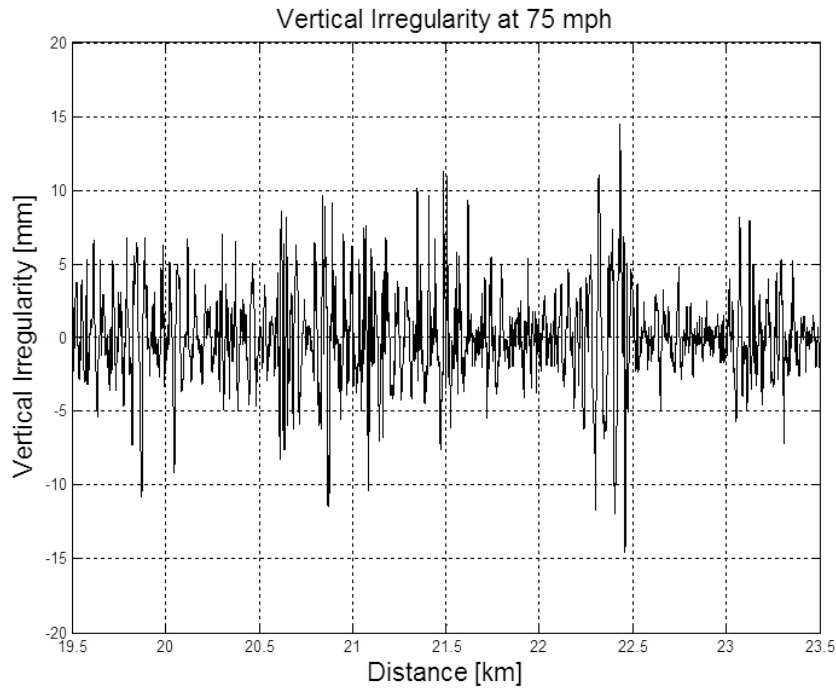
Figure 5.8: Vertical alignment recorded at 75 mph

To further illustrate the consistency between the results, the processed outputs from all three speeds have been plotted on the same graph, Figure 5.9. To separate the three lines on the figure, the data from the 50 mph run has been plotted unaltered while the 20 mph and 75 mph data have been plotted with +30 and -30 mm added respectively. The compressed scaling of the figure emphasises the consistency of the three runs as the differences between them, which are at most one or two millimetres, become less significant.

The largest irregularity in the vertical alignment data is between 13 mm and 14 mm. This is, however, a peak value with the largest consistent irregularity around 10 mm. This is within the extremes of the results reported in [40], as might be expected for a dataset that does not include any significant S&C. The general alignment, particularly at the start of the section presented in the figures, is moderately poor. It, along with some of the specifically poorer locations, will be investigated in Section 5.6.3.
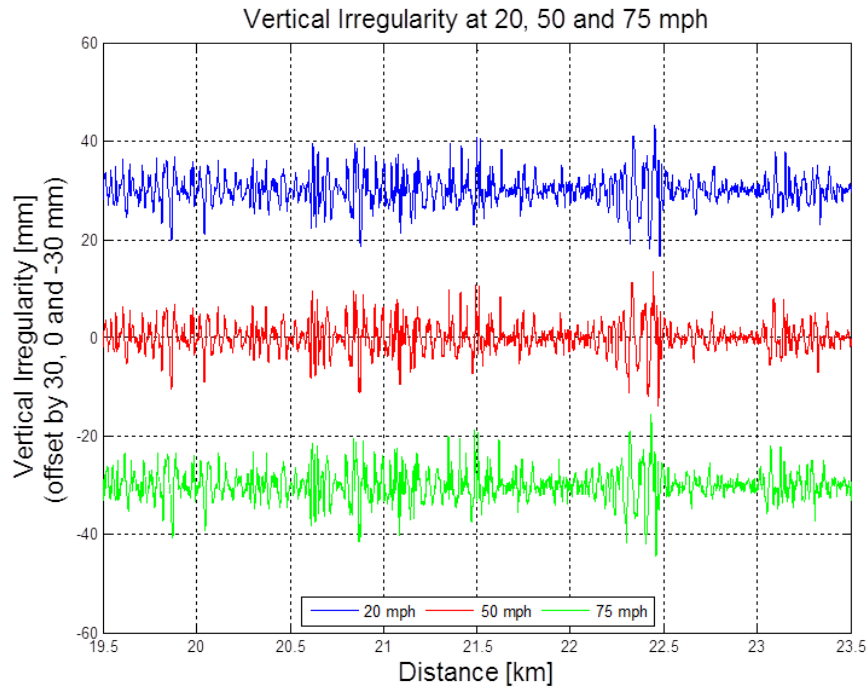
Figure 5.9: Vertical alignments recorded at 20 mph, 50 mph and 75 mph

5.6.2 Lateral alignment

The data from the yaw rate gyro, obtained in all three trial runs, have been processed as described in [66] to give estimates for the lateral alignment of the track. As the trial runs were not significantly affected by any S&C work, the lateral alignment results were expected to be particularly uninteresting. The 20 mph result from the same 4 km section as that presented in the vertical alignment result is shown in Figure 5.10.

The figure shows a good general alignment without any of the increased background level that was present at the start of the vertical alignment figure. The same four areas of lower quality alignment are even more clearly visible than in the vertical figure (Figure 5.7). The background levels shown in the figure are consistent with those found in [66]; the levels around features are slightly lower, however, no S&C were traversed in the most recent trial runs so the levels were expected to be much lower. The nature of the features will be investigated in Section 5.6.3.

165
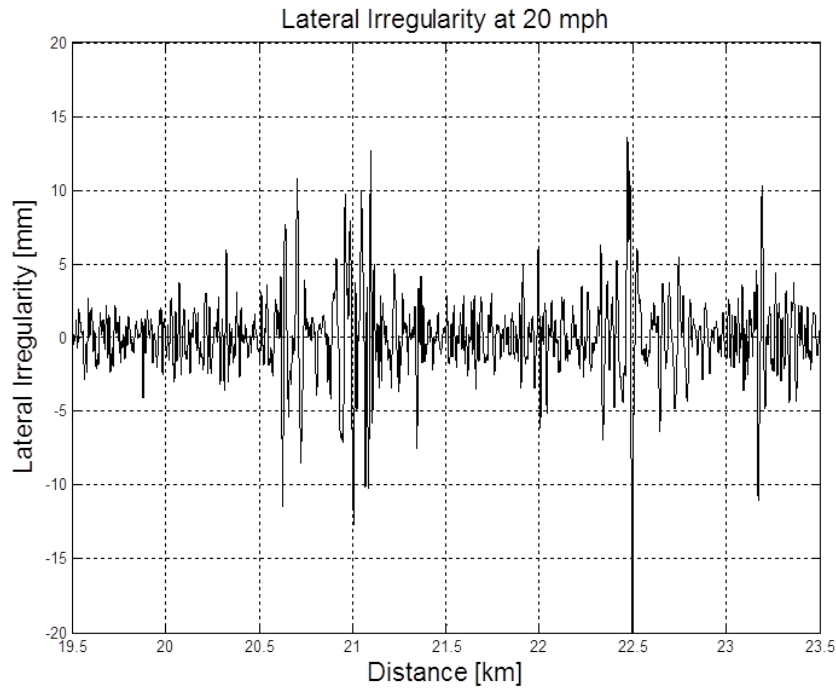
Figure 5.10: Lateral alignment recorded at 20 mph

The same figure, generated from the 75 mph data, is shown in Figure 5.11. The result is consistent with that identified at 20 mph. The general alignment is comparable with the 20 mph result but the peak alignment values are noticeably increased. This effect is more pronounced than that found in the vertical results taken at the same speeds. The 50 mph result, not shown, fitted predictably between the two presented results.

As with the vertical alignment results, plotting the three different speeds (offset) on the same graph reduces the space available for each in the vertical axis and so makes the differences harder to discern. Regardless, the plot of all three lateral results is presented in Figure 5.12. The figure highlights the consistency in the general alignment and clearly illustrates the differences in magnitudes of the results between the three speeds around features. The feature around 20.6 km is particularly noticeable as the shape of the vertical alignment curve is shown to differ at 75 mph compared to the other speeds. This is, however, likely to just be a dynamic effect.
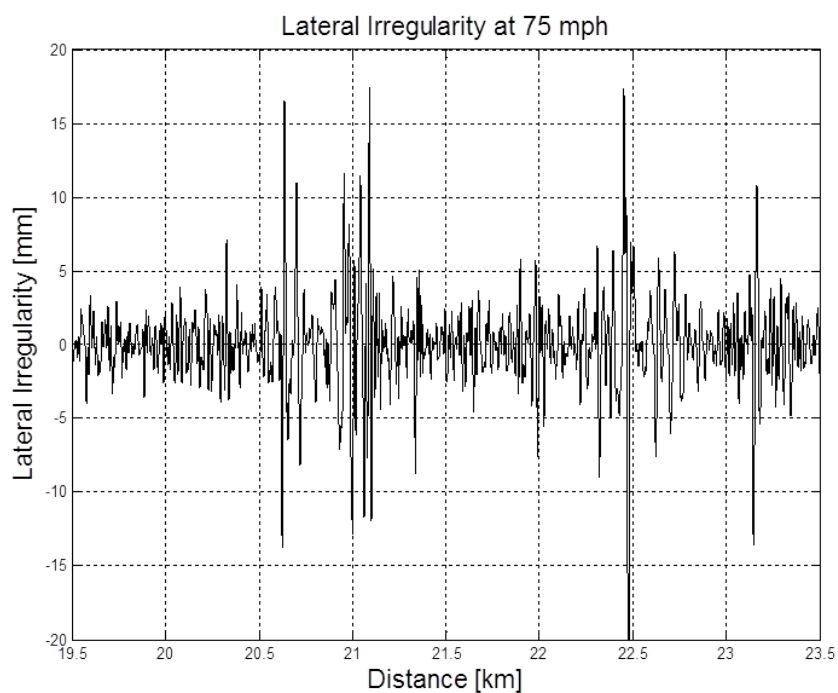
Figure 5.11: Lateral alignment recorded at 75 mph
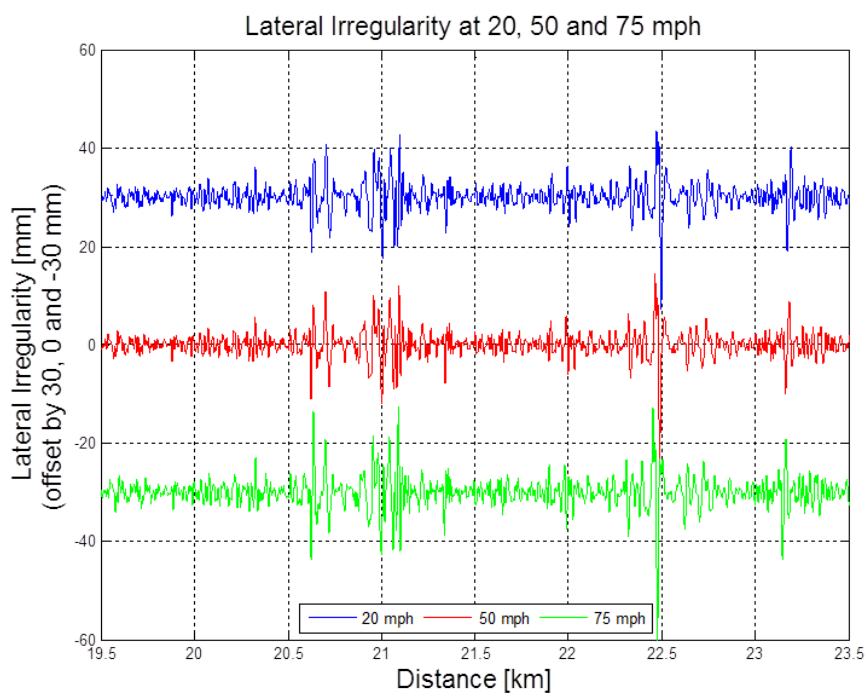
Figure 5.12: Lateral alignments recorded at 20 mph, 50 mph and 75 mph

5.6.3 Alignment with features

Most of the features that have been identified in the lateral alignment result can be shown to coincide with those in the vertical alignment result. Figure 5.13 shows the processed alignment data from both the pitch and yaw rate gyroscopes. As mentioned previously, the background level is higher in the vertical alignment data. Using a simple thresholding system at around +8 mm, the lateral alignment data reveals four features. These are located roughly around 20.75 km, 21.0 km, 22.5 km and 23.2 km. Using the same system on the vertical alignment data indicates the same features around 20.75 km, 22.5 km and 23.2 km, albeit fractionally earlier. This is possibly due to the vertical alignment data being obtained over the wheelbase of the bogie. The feature around 21 km is discernable but somewhat distorted by the data immediately following it. A further feature, located at 21.5 km is also visible.
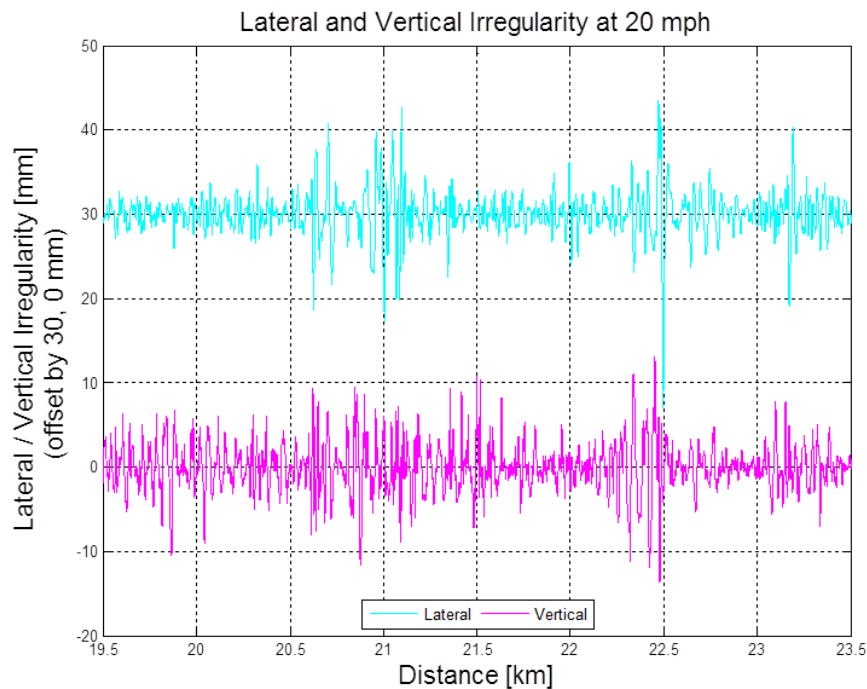
Figure 5.13: Lateral and vertical alignments recorded at 20 mph

Although all of the S&C in the trial route was encountered in the trailing direction and normal position, it is likely that some of the irregularity seen by the inertial measurement instrumentation could still be caused by the presence of S&C. Around S&C, the conductor rail is often disjointed or has to change sides which would appear as gapped sections to the conductor rail instrumentation.

Figure 5.14 shows the data obtained from the vertical alignment processing of the pitch rate gyro for all three of the test runs on the same figure as the conductor shoe displacement result obtained at the same time and discussed in the previous chapter. The features have been marked on the conductor rail plot and indexed A to E.
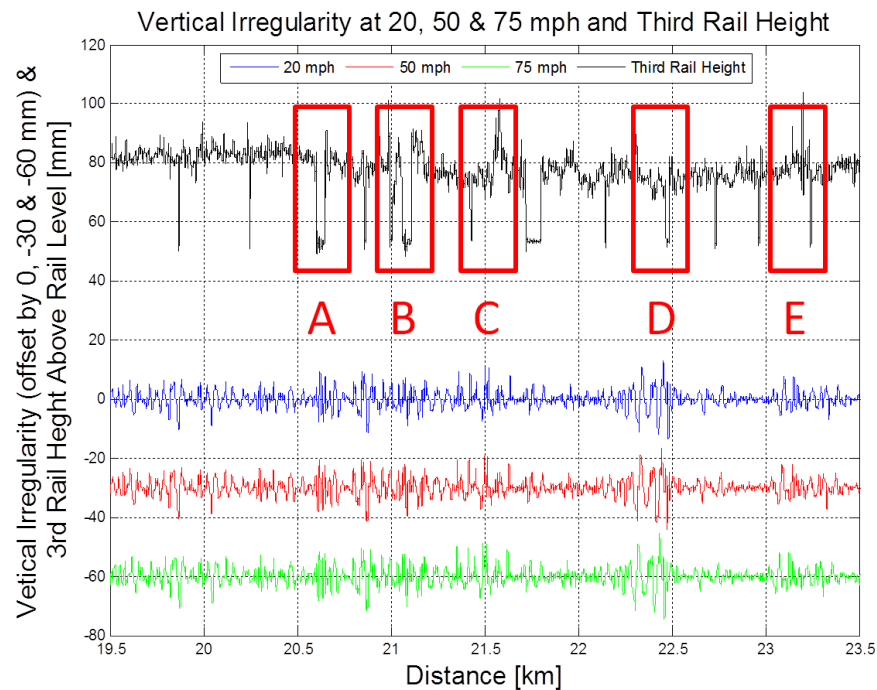
Figure 5.14: Conductor rail height a.r.l and vertical alignments recorded at 20 mph, 50 mph and 75 mph

Despite the additional feature, the background level of the vertical alignment result suggests that it is more suited to general rail condition analysis than feature detection. Features seem to be more discernable when using the lateral alignment result. Figure 5.15 shows the same

alignment between the conductor rail measurement and the features identified by the inertial measurement unit, this time using the lateral result. The same five features have been marked on the conductor rail plot despite only four of them being visible in the inertial measurement data.
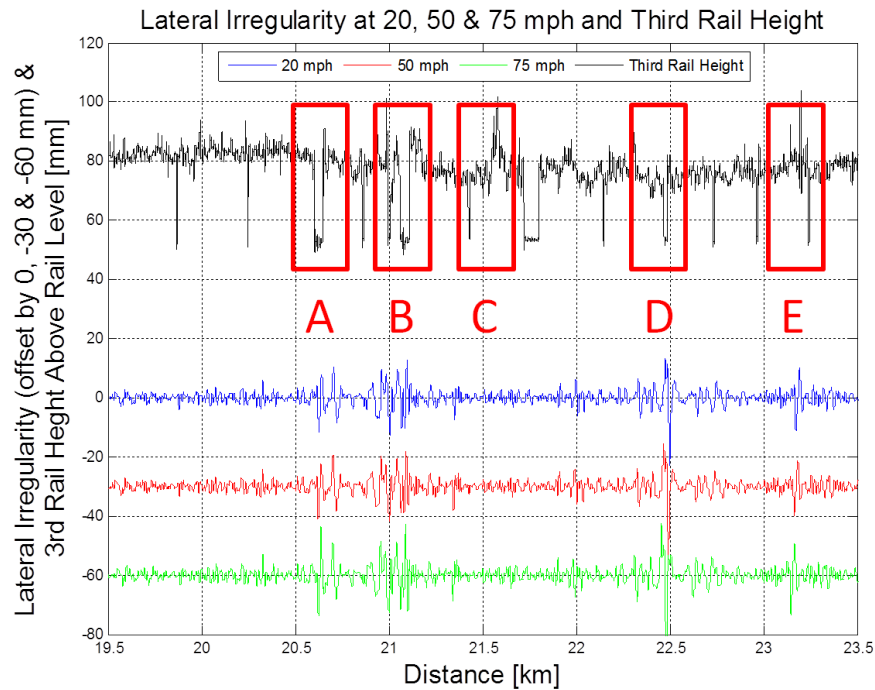
Figure 5.15: Conductor rail height a.r.l and lateral alignments recorded at 20 mph, 50 mph and 75 mph

The 4 km section of data presented in the figures was selected because it contained a number of features rather than because of the specific nature of any of these features. As it happens, only three sets of trailing, normal, S&C were encountered during the course of each of the trial runs and two of them are contained within the figures.

Features "A" and "B" correspond to the S&C around Herne Bay Station. Feature "A" is on the Birchington-on-sea side while feature "B" is closer to Chestfield and Swalecliffe. Herne bay is on a curve, which may have contributed to the high background levels in the vertical result. Aerial photographs of the S&C represented by features "A" and "B" are shown in Figure 5.16. These

images were obtained using Microsoft's Virtual Earth software. The second photograph in the figure corresponds to the junction on the Chestfield and Swalecliffe side of Herne Bay. This junction is actually a combination of a track-to-track crossover and the up siding joining the line. This may explain why the feature appears distributed in the alignment figures, particularly the vertical alignment figure.

Figure 5.16: Features "A" and "B"

Having located the Herne Bay S&C in the data, it is possible to look at the area around the station more closely. Figure 5.17 shows the same vertical irregularity data combined with the conductor rail height relative to the running rails shown previously. The figure has been zoomed in to consider only the area between the two sets of S&C which mark the station boundaries.

The first set of S&C occurs at approximately 20.6 km. Oscillations are visible in the vertical alignment and there is a gap, or a changing of sides, visible in the third rail data. Approximately 200 m further along the track, there is another short gapped section which coincides with an area of poor vertical alignment. This is likely to be in the centre of the station where track maintenance is particularly difficult. At approximately 21 km, there is a second, short, gapped section. This is likely to be the walking route visible towards the right of photograph "B" in Figure 5.16. Again the fact that this is contained within the envelope of the station, where track

171

maintenance is difficult, can be observed in the alignment data. At approximately 21.5 km, there is a large gapped section which corresponds to the S&C at the Chestfield and Swalecliffe side of the station. The vertical alignment data through this set of S&C is possibly slightly better than that found through the S&C at the other end of the station; however, it does not indicate that the region is particularly better or worse than the area around it. Although not visible on the photograph, there is a bridge shortly after the station. It is possible that the difficulties associated with track maintenance during stations and under bridges is masking the quality of the result through the S&C.
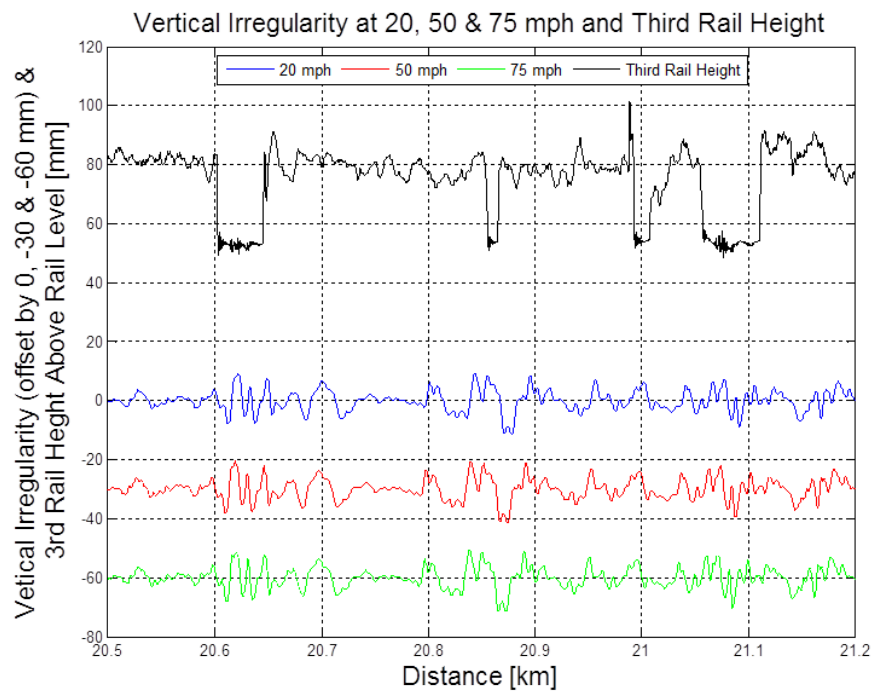
Figure 5.17: Close up view of conductor rail height a.r.l and vertical alignments recorded at 20 mph, 50 mph and 75 mph

A zoomed version of the lateral alignment data for the same section is shown in Figure 5.18. The inertial measurement results around the first set of S&C illustrate a characteristic "lurch" as the vehicle passes through the section. This is amplified with the increase in speed. The first short gapped section occurs in an area of rather good lateral alignment, suggesting that it is

172

not a significant feature and that the poor vertical alignment may be symptomatic of general track or substructure condition. Towards the end of the station, the general lateral alignment degrades. At the walking route just past the tip of the station platform, the figure shows a reasonably long wavelength excursion in only one direction, this is not characteristic of S&C. The final, longer, gapped section previously thought to be S&C shows a series of characteristic oscillations. As with the S&C at the start of the station, these are more significant with increased speed. Shortly after the S&C, the lateral alignment returns to a respectable background level. This suggests that, as with the gapped section in the centre of the station, any poor alignment seen in the vertical result is specific to the vertical direction and may possibly be related to the substructure.
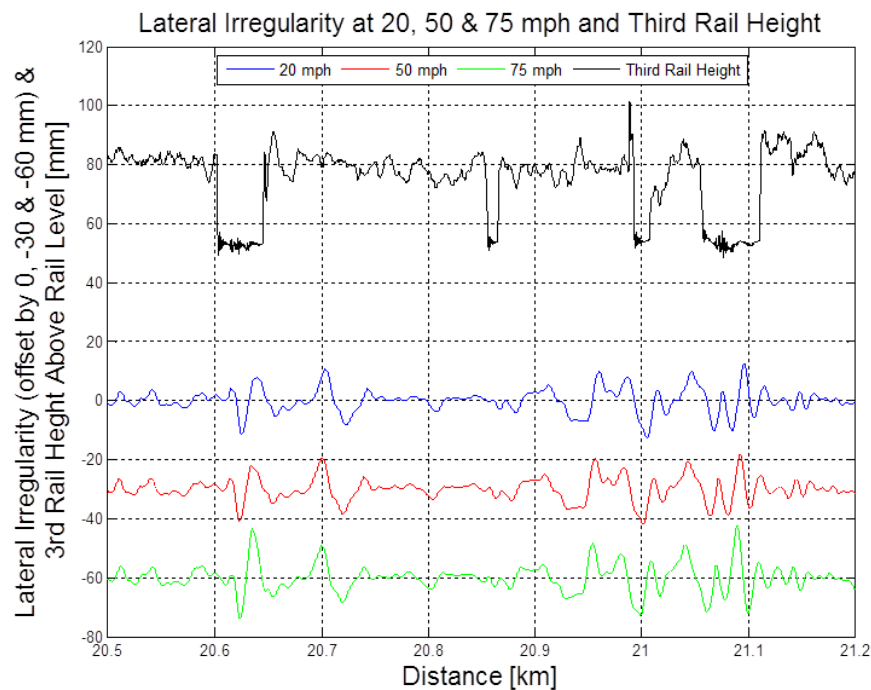
Figure 5.18: Close up view of conductor rail height a.r.l and lateral alignments recorded at 20 mph, 50 mph and 75 mph

Returning to Figure 5.14 and Figure 5.15; the irregularity around feature "C" is only visible in the vertical data. The conductor rail measurement indicates some abnormalities to general

running, but any effects are very short term. The section of line around feature "C" can be seen in the top section of Figure 5.19, a collection of photographs taken from Microsoft's Visual Earth software. The trial path was along the bottom track from right to left. No significant features are obvious although a short gap in the conductor rail may be visible towards the left of the picture. This could align with the dip in the conductor rail data; however, the result shown was not calibrated for vehicle roll and so it may be just as likely that it is an amplified artefact of the vertical alignment.
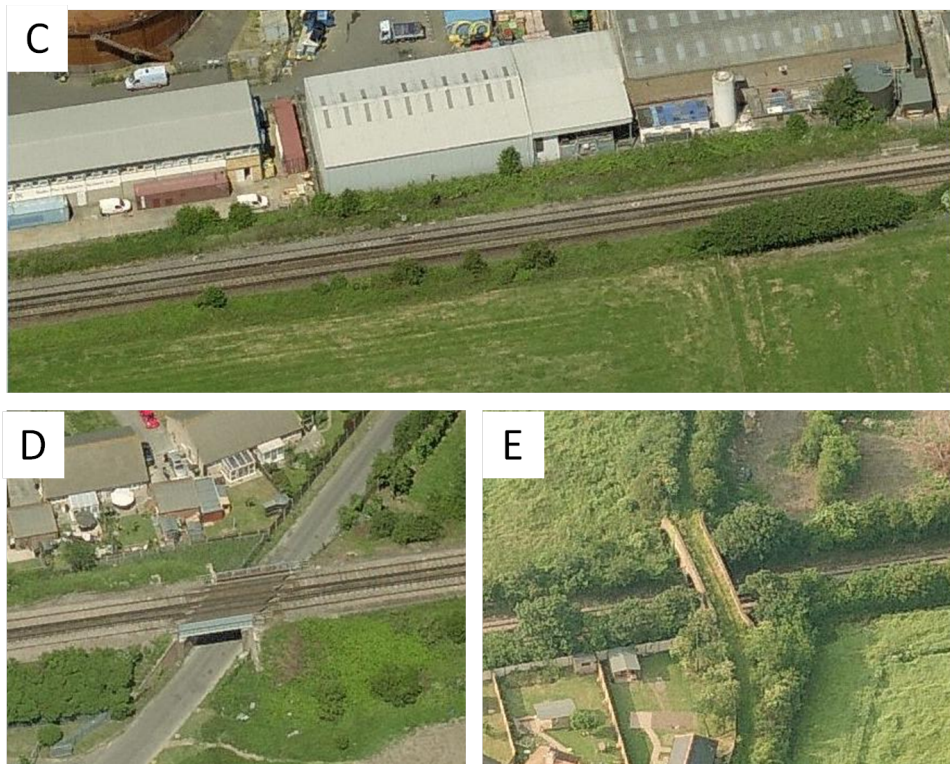
Figure 5.19: Features "C", "D" and "E"

The lateral alignment data indicates feature "D" to be quite an abrupt event while the vertical alignment data suggests that it had been building in magnitude for almost half a kilometre. The conductor shoe displacement data shows only a short gapped section. Feature "D" corresponds to a bridge allowing the railway to pass over a local road, it can be seen in Figure 5.19. The short gap in the third rail is likely to align with the length of the bridge. The railway that passes

over the bridge will, by necessity, be far stiffer than the general ballast based railway found in the area around it. Transitions from ballasted to more rigid railway need to be carefully managed to avoid large impacts as the vehicle transitions from the, movable, ballasted track to the comparatively immobile fixed structure. Lateral alignment to these fixed points in the railway is difficult and entry to the fixed section is often accompanied by a "jolt" such as the one visible in Figure 5.15. Vertical alignment is affected by the intentional changes in the stiffness of the track as it approaches the fixed structure. These changes are likely responsible for the increasingly "poor" alignment (as viewed over 70 m) present in the data displayed in Figure 5.14.

The feature designated "E" can be characterised by a very short but very significant drop in the conductor rail, a "jolt" in the lateral alignment which increases in magnitude with speed, and a minor increase in the vertical alignment result spread over a few hundred metres. Feature "E" represents the railway passing under a minor road bridge. The road in question appears almost disused but still features on local maps. As with the bridge associated with feature "D", track maintenance work around this kind of bridge is difficult. However, while realignment of track passing over roads is difficult due to the substructure, maintenance in this case would be difficult due to the practicalities of using heavy equipment at the site. The "jolt" in the lateral alignment is likely due to it being difficult to precisely align the track with the, almost fixed, section under the bridge. The stark, but relatively minor, change in vertical alignment quality around the bridge area would suggest a region particularly close to the bridge where maintenance was undertaken on a less regular basis. The drop in conductor rail is likely a very short gapped section under the bridge or the result of the dynamic system's interaction with the vertical alignment.

5.7 Conclusions

This chapter has presented work carried out to fit inertial measurement instrumentation to a Class 375 railway vehicle. The work described here has also been documented as a component of a broader journal paper considering the use of vehicle-based sensors for condition monitoring [106]. The author of this thesis is a named author on this paper.

The work presented here can be thought of as an extension of the previous ERCIR project, with this phase of the work illustrating an evolution in the instrumentation systems towards equivalent functionality being obtained using a reduced sensor set. This is the logical progression of research instrumentation, where initially a system must be over-instrumented to identify the key parameters before the instrumentation is refined to be more efficient.

The architecture of the instrumentation has also been upgraded to include local transducer management with separated data storage. This architecture reduces errors in the data, caused by noise in the wiring looms, by eliminating the majority of the cabling in the system and using digital, fibre optic, communications for the longer cable runs. Although previously thought of as a two-site system, this architecture could also be considered to be node based (albeit with only two nodes) as the electronics to marshal the transducers are now sufficiently complex and entirely electrically isolated from the data recording system. This will become significant in the next chapter.

The work presented in the case study has made use of gyroscopes in a role more traditionally filled by accelerometers. While not novel in itself, this extends the concept demonstrated in the previous case study where multiple, independent, transducer types can be used in place of each other, for redundancy, or as a verification tool. In this case, the gyroscopes were selected for increased resilience to low frequency noise in the tacho signal however higher frequency

measurements may be better suited to the accelerometer technology. This will be investigated further, in the next chapter, through the concept of complementary filtering.

Some switch and crossing work has been identified although the configuration and limited number of test runs have meant that only basic characteristics have been considered. Making use of the limited dataset, other features have been identified in the trial sections. Bridges, both over and under the track, have been inspected and their impact on track maintenance work considered. Based on these initial observations, the University of Birmingham has subsequently undertaken further studies into track geometry changes around these key features. The results have also allowed inspection of the general track condition as per that previously demonstrated in the ERCIR project. Laterally the track was generally in good repair, however, vertical alignment results suggested some areas that were in poor condition. Since the trials were conducted, Network Rail has performed extensive track maintenance and has actually re-laid much of the trial route.

Performing the three trial runs, at different speeds, and immediately following one another has allowed comparisons of the data and instrumentation performance at three different speeds. Although differences have been detected they have been mainly dynamic effects and the general consistency across the three speeds is high. The speed of the vehicle is a parameter in the processing required to analyse the data. However, as long as the vehicle speed was accurately known and reasonably consistent over sections, it should be possible to perform the alignment processing on a vehicle with a varying speed. There are limitations on the use of inertial (accelerometer and gyroscope) based instrumentation systems operating at low speeds, and there would be minimum distances where consistent speed would be required for the processing. This would put limitations on the use of such vehicles within stopping patterns, particularly for metro style services, but generally the railway operates with consistent line speeds and as such

these systems should function.

This leads to the possibility of instrumentation systems attached to in-service vehicles either recording vast quantities of network data or performing real time analysis. These kinds of long timescale, in-service measurements are essential if automatic inspection systems are going to be more heavily used as they allow the possibility of monitoring features through time and identifying evolving faults. Doing this would inherently involve the generation, recording and processing of much more data. An alternative approach may be to move the data processing to a real-time on board system. Adding long term database support and processing to existing regular inspections is an alternative approach; but the interval of inspection may not identify rapidly evolving faults and so such systems are unlikely to replace more regular manual inspections.

CHAPTER 6

CASE STUDY 4 - DISTRIBUTED EMBEDDED SYSTEMS

6.1   Introduction

The three case studies presented previously have described instrumentation fitted to railway vehicles. Each of these sets of instrumentation was installed for a particular experiment and used a range of sensor systems appropriate to that task. Although some evolution of the systems has been displayed, all three systems have been architecturally very similar and there been little crossover between them in terms of transducer selection. A summary of the sensing technologies used in each case study is shown in Table 6.1. At the core of each of these instrumentation systems lies a microprocessor, some form of analogue to digital conversion and a PC for recording data. These are arranged in the form of a local transducer management node, with a remote data acquisition component. This architecture is adequate for the simple systems described to date, but is not extensible to larger systems due to bottlenecks in the core processing and communications systems.

Table 6.1: Sensing technologies used in each of the three case studies to date

Case study	Technologies used
1 - Energy monitoring	Current, voltage, (tacho)
2 - Conductor shoe monitoring	Displacement, strain gauge
3 - Inertial measurement	Accelerometer, gyroscope

This, final, case study presents an example instrumentation framework which can be used to acquire signals from transducers located individually or in clusters rather than all in one location. The framework is equally applicable to networks of disparate transducers as well as matched sensor arrays, and is equally well suited to both infrastructure and rolling stock based applications. The framework is showcased through a case study involving a piece of work undertaken to develop an advanced instrumentation system, using sensors arrayed at multiple locations, that is capable of measuring multiple aspects of a railway vehicle in order to use the wealth of data obtained to give context to some of the more key signals. The work described in this chapter involved the development of a distributed system of embedded processors, each responsible for servicing a selection of transducers at different locations on the vehicle, and each communicating with the others to maintain synchronisation and a global data repository. In some ways, this is the natural evolution of the previous systems, where instead of a single sensor node communicating with a host, multiple sensor nodes communicate over a backbone network which also links to a data repository.

This backbone based instrumentation framework uses standard communication protocols with a custom presentation layer to allow an almost arbitrary number of sensor nodes to connect and share data on a common bus. Using a similar mechanism, an arbitrary number of data extraction nodes can also be fitted to the bus. These acquire the signals presented by the acquisition nodes and either transmit a required subset, or record the signals for later analysis.

The increased complexity of the instrumentation system allows a greater understanding of the operation of the vehicle to be obtained. For example, in the first case study which related to energy monitoring the data processing was curtailed by a lack of information relating to the impact of the driver and the environment on the operation of the vehicle. With the instrumentation architecture presented in this chapter, it is possible to obtain this information and combine it with the energy monitoring data during the analysis phase.

The instrumentation described in this case study was developed as an extension of that described in the first (energy monitoring) case study. Following on from the success of the single day trial, instrumentation was fitted to a vehicle remaining in passenger service. The vehicle was then operated, in-service, for an extended period with the instrumentation providing a rich dataset of energy use profiles for different journeys including realistic stopping patterns. In addition to extending the energy monitoring work from the first case study, the new instrumentation incorporated the inertial components demonstrated in case study three, and some additional inputs highlighted as relevant through analysis of the existing datasets.

6.2 Motivation

The motivation for the work described in the energy monitoring case study was largely concerned with the development of specifications for on train energy metering systems and what the recorded data could be used for. The work demonstrated that by sampling at higher data rates and at more locations than those proposed in the draft standards it is possible to obtain additional understanding of vehicle behaviour [111]. The work described here extends this concept to include a far greater range of signals thereby allowing a better understanding of the operation of the vehicle. The additional duration of the trials also provides a richer dataset to be used with existing analysis techniques.

Energy efficiency is now a significant concern for all Train Operating Companies. In addition to proposed schemes for compulsory energy metering [79] [112] and modifications to track access agreements to include charges based on energy usage [83], appearing to be environmentally conscious is becoming more important with customers. Merseyrail is aware of these issues and is moving towards a more energy efficient railway. In the short term, this is being fulfilled through modifications to operational procedures and, in particular, driving style. Prescribed outputs of this work include an analysis of different driving styles for their energy efficiency and recommendations to improve this over all of the Merseyrail routes.

Considering the longer term; the Class 507 and Class 508 fleets operated by Merseyrail date from between 1978 and 1980 [37] and are due to be retired in the foreseeable future. Energy usage data from these studies can be used to identify energy saving strategies that can be applied either to the infrastructure or on vehicles. It may be possible to take the understanding of energy use in the network gained in these studies into consideration when selecting the rolling stock that will succeed the existing fleets.

This work also extends that presented in the third case study relating to inertial measurements. The system developed is designed to operate unattended on an in-service vehicle and so systems have been included to allow alignment and positioning of the data during post processing. These systems have been expanded to include a full Inertial Measurement Unit, with tacho and GPS systems, fitted to the bogie. The long term data obtained by this inertial measurement unit over repeated routes can be used in the consideration of track quality and degradation, particularly around switch and crossing work.

6.3 Instrumentation

The instrumentation developed for this work uses a new architecture based on modular network. It has been designed to operate for an extended period of up to one year with minimal user interaction. User involvement should be limited to periodically replacing the storage media as the system approaches capacity. To ensure this level of independent operation, the system makes use of embedded microprocessors which are resilient to power-outages and can accommodate regular power cycling without notice. All storage is in the form of solid state media as this is more resilient to the level of vibration found on the vehicle and less likely to sustain damage or data corruption in the event of a power failure.

6.3.1 Overview

Although this instrumentation is described as an energy monitoring system, it is far greater in complexity than its predecessor which was the subject of the first case study. This system includes a far greater range of sensors and inputs than that described previously. These can be used to assist in the interpretation and contextualisation of the energy results. In addition, the system has the added complexity of local data storage and a requirements for a year of continuous (in-service) operation.

The energy monitoring system described in the first case study consisted of a traction current transducer, an auxiliary current transducer and a voltage transducer. A tacho signal was also recorded. These signals provided information on the energy consumption of the vehicle, and allowed some initial hypothesis about it's inner workings to be developed. The system was used to estimate driver behaviour, but did not directly monitor any driver or traction system performances. The instrumentation described here uses a far more complex structure capable of directly integrating an extended range of signals from around the vehicle. These signals can

be used to contextualise the energy data.

Initially, the same signals (traction current, auxiliary current and line voltage) were recorded. This time, however, the signals were integrated into a managed data stream which also included time and date stamp information. This allowed indexing to be performed to extract specific information from the expanded dataset. In addition to the traditional energy monitoring signals, the tacho signal was more closely integrated into the data stream so that alignment between vehicle performance and the energy consumption could be considered directly. Observations from the original energy monitoring case study suggested particular traction system behaviour which could not be verified from the energy signals alone. Logical (digital) signals controlling the camshaft based traction system were recorded digitally as a mechanism for confirming traction system performance and therefore explaining the effects seen in the energy data. Similar signals describing the driver's demands to the vehicle (including the type of braking) were recorded in order to contextualise the combined speed and energy data being used in establishing energy efficient driving styles. A signal corresponding to the weight of the vehicle was also recorded as in input to vehicle system performance. Another thing that was hypothesised to have a significant effect on the energy consumption seen in the first energy case study was the gradient of the track on which the vehicle was operating. There are two ways to identify this, either know the position of the vehicle and have a gradient map of the network, or to install an inertial measurement unit capable of recording vehicle pitch etc. The extended instrumentation system described here adopted both of these approaches fitting high grade GPS equipment as well as a full (6-degree of freedom) inertial measurement system. The full list of signals recorded by this system are summarised in Table 6.2. The two ends of the vehicle operate independently; quantities in the table are for both halves. Where a single instance of a signal is recorded that component is only fitted to one end of the vehicle.

Table 6.2: Signals recorded by the enhanced energy monitoring instrumentation

Quantity	Signal	Justification
2	Traction current	Traction energy
2	Auxiliary current	Auxiliary energy
2	Voltage	Energy calculations, one only per vehicle end
2	AETW	Vehicle acceleration moderated by this weight signal
2	Temperature	Temperature compensation in post processing
2	Acceleration: lateral (body)	Redundancy
2	Rotation: yaw (body)	Redundancy
2	RTCC	Time, date and sample request signal
2	Tacho	Vehicle speed and position
4	Driver direction request	Forward, neutral, reverse. Duplicated at both ends
10	Driver power demand	Four levels of demand plus coasting. Both ends
8	Cam shaft operation	Locating signal plus direction and shift requests
2	Braking mode indicator	Resistive or physical braking
1	Acceleration: longitudinal	Positioning
1	Acceleration: lateral	Positioning
1	Acceleration: vertical	Positioning
1	Rotation: roll	Positioning
1	Rotation: pitch	Positioning
1	Rotation: yaw	Positioning, particularly curvature alignment
1	GPS: position	Positioning
1	GPS: speed	Positioning
1	GPS: heading	Positioning
1	GPS: time / date	Alignment
1	GPS: validity indicator	Allows GPS data to be discarded as appropriate

The signals described in Table 6.2 are each available only at specific locations on the vehicle. In many cases access to these areas is limited and cable routes are in areas of significant electrical interference. Not withstanding the extensive processor load, it was therefore not possible for the system to consist of a single master processor with sensors located around the vehicle. An alternative architecture was required.

At the core of the instrumentation framework lies a communications bus. Every piece of electronics that attaches to this bus exists as a module. Modules can be used to put data onto the bus, or to take it off. The key concept of the framework is that the modules have standard interfaces and the communications over the bus use a standard schema to present data. In this way, custom modules (nodes) can be developed at key instrumentation sites around the vehicle, to be local to the sensors that they are managing, and to have the appropriate level of processing capability and custom I/O for their own individual instrumentation requirements. Recording modules can be located anywhere on the network, and through the same communications interface and schema can selectively extract signals of particular interest and/or from particular sensors. The system is therefore extensible as long as there is remaining capacity on the data bus at it's core.

The developed system comprises four clusters of electronics (maximum per vehicle end) each of which is responsible for servicing a selection of the signals described in Table 6.2. Each cluster attaches through a standard presentation system to the instrumentation framework developed specifically for distributed sensor networks. The backbone of the framework consists of a power and communications bus which provides a highly resilient mechanism for transferring the obtained signals around the vehicle. The clusters, the signals located at each cluster, and their approximate locations on the vehicle are indicated in Figure 6.1. The dotted line linking to the laptop indicates where the laptop is connected for commissioning and testing processes. The

laptop is not required for normal system operation but can be used for live, real-time visualisation.
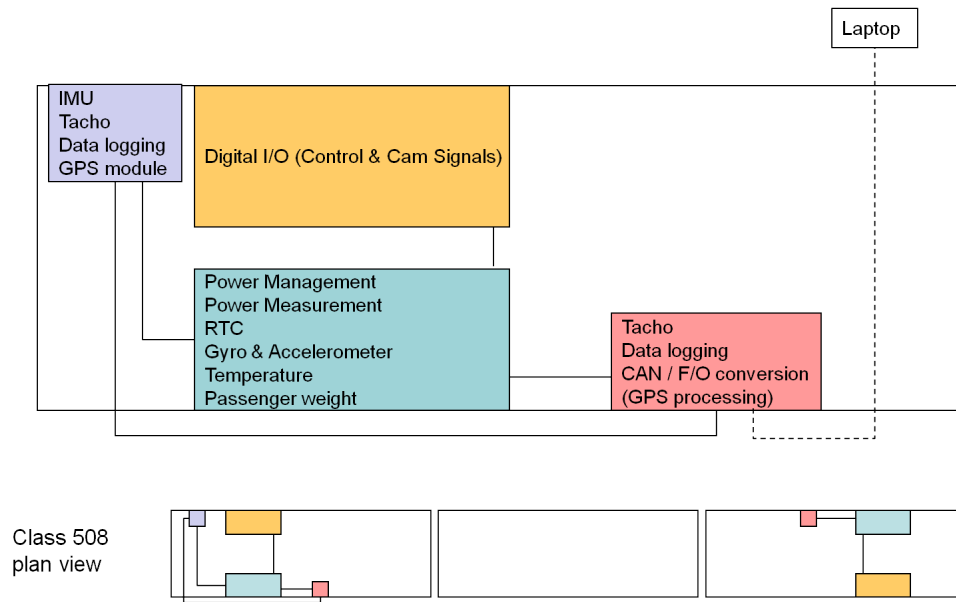
Figure 6.1: Instrumentation sites on the Class 508 vehicle

The instrumentation framework operates through the sensor nodes each packaging the data recorded from their individual sensors according to a standard presentation schema. The nodes then transmit this packaged data via the framework backbone. A second series of reporting nodes then also connect to the backbone. These are capable of deciphering the schema and unpacking the recorded data. Multiple types of reporting node are used to select designated subsets of the signals within the network based on the presentation schema, and report them to real-time visualisation systems, or to record the data being transmitted over the framework backbone in its entirety for later analysis.

The nodal structure of the instrumentation framework means that the system can inherently operate in a modular fashion. Providing there is "address space" in the schema and capacity on the bus, new nodes can be developed and incorporated into existing system without affecting the overall system reliability or operation. Furthermore, the nodal nature means that the majority

of failure modes from which a node may suffer will not materially affect the rest of the system. Single points of failure do still exist, notably in the power supply and data bus lines, but most can be eliminated through redundancy. For example, there is no reason why the data recording node within the system could not be duplicated or paired with a transmission and visualisation node.

6.3.2 Technologies

Each cluster of electronics consists of one or more printed circuit boards, each containing an embedded microprocessor. The data bus selected to link the different embedded processors together is the CAN bus. CAN stands for Controller Area Network which is a mechanism by which small systems can communicate over a bus without a master controller. CAN was originally developed between 1983 and 1986 by Bosch [113]. The technology was primarily championed by Phillips and, to some extent, by Intel. It is heavily used in the automotive and aerospace industries. In 1991, Bosch released a standard for version 2.0, the physical and datalink components of the protocol are described in ISO 11898 [114].

The CAN bus was selected as the system backbone for a number of reasons:

Firstly, the system does not require a separate clock and makes use of only a two wire interface. This level of simplicity is beneficial in a system which aims to reduce the amount of cable being installed around the vehicle. The two wire interface, combined with the power for the instrumentation cluster, can be installed in a single four core cable. This cable is twisted and a shield is added due to the nature of the environment, but this is still a huge reduction in cabling when compared to a single processor with a distributed network of sensors.

The CAN system is capable of operating at distances of up to 500 m at a data rate of 125 kbps. Faster data rates are available over shorter distances. The instrumentation being described here

uses a data rate of 100 kbps and operates over less than 20 m, and so is easily supported by the CAN protocol.

The CAN architecture uses a message based system where, among other things, each message has a message type identifier, a payload size indicator and a data payload. This fits ideally with the instrumentation system which has a number of different signal types originating from different transducers. It is possible to assign a CAN message type to each signal type, thus making data extraction significantly easier. Furthermore, these message type indicators can be prioritised by the CAN system. Higher priority messages will be transmitted in preference to lower priority ones which will retry in the event of both being put onto the bus at the same time. In the instrumentation system, the timely arrival of the timestamp messages is critical and so these have the highest priority. The signals related to the energy signals are also considered important and so have a reasonably high priority. The temperature signals, for example, are comparatively unimportant, as temperature varies only at a low frequency, and so are assigned a low priority.

The most significant reason for the selection of the CAN system, however, is its resilience to interference. The system operates using a two wire, balanced, signal which provides significant resilience to all but common mode interference. The CAN module interpreting the balanced signals then rejects any common mode interference that has been picked up in the wires. To improve the performance further, the two wires are twisted together to share exposure to electromagnetic interference, thus ensuring that it appears as common mode rather than on a single wire. This level of resistance is of particular concern inside the Class 508 vehicle where data bus cables are required to pass through regions of the train where high current switching occurs. This is well known as a source of interference.

The most logical alternative to the CAN bus would have been Ethernet. Ethernet is a multi-layer

protocol capable of transferring huge quantities of data and with native support for handshaking and data re-communication in the event of interference. Ethernet can easily operate over distances of several hundred meters and much further if media converters (to and from fibre optic communications) are used. There are, however, several reasons why Ethernet was not selected for the backbone in this instance. Firstly the hardware required to manage the Ethernet interfaces is substantially more power intensive than the CAN equivalent. Although trains are large, additional equipment fitted to them (when in-service) must either draw very small amounts of power, or have extensive systems to ensure that they de-activate in the event of the vehicle switching to backup power. Using Ethernet as the system backbone would have used a notable component of the system's power budget. The Ethernet interface is also more processor intensive which would have affected the capability of the core processors in the nodes to the extent that different processors would likely have needed to be used. An additional concern was the physical configuration of the backbone. Ethernet networks tend to form star configurations with the switching fabric at the centre. This architecture is not tremendously compatible with large scale distributed networks where cabling capacity is a limitation. While a star structure could have been used for the work being described here, as a structure it is generally speaking a less compatible with broad scale railway based instrumentation networks.

As the work presented in this chapter is an extension of the work presented in the first case study, a certain level of crossover between the technologies used is to be expected. The energy monitoring core of this set of instrumentation is still based around the DV 1200/SP2 voltage transducer and the LTC 600-SF and HTR 100-SB current transducers, all from LEM. In this case, however, one of each is fitted at each end of the train, and the systems used to monitor the transducers also monitor a range of other signals.

The Class 508 being used for the work has a system which monitors the pressure in the vehicle's

air suspension to form an approximation for the combined weight of the vehicle body and passengers. This information is then supplied to the acceleration and braking systems through the Air-to-Electric Transducer - Weighing (AETW) and Air-to-Electric Transducer - Braking (AETB) systems respectively. These systems modify the vehicle's acceleration and braking routines based on the combined weight of the passengers on the vehicle. The instrumentation developed for the trials takes a copy of the output of the AETW module and passes it to an isolation amplifier. The amplifier selected is the AD202JY from Analogue Devices which provides up to 1000 V of common mode isolation.

In addition to the main energy monitoring components and the AETW interface, the first cluster of electronics also contains thermometry, a longitudinal accelerometer, a yaw rate gyroscope and the system's real time clock module. The thermometry is included to allow compensations to be made to the calibrations of the other transducers (where required) as the ambient temperature is likely to vary significantly over the course of the year long trial. The accelerometer and gyroscope are lower quality, body mounted, versions of those installed on the bogie. They are included primarily as redundancy in the event of the equipment on the bogie failing in the harsh environmental conditions, but also serve as a secondary data alignment system and allow comparisons of the effectiveness of low cost, easily fitted components, with more expensive ones which are more complex to install / maintain. All of these analogue signals are serviced by an LTC1859 analogue to digital converter linked to the main microprocessor via an SPI interface.

A real time clock is used as the primary point of time alignment for the whole system and its outputs are used in the generation of the addressing key for the storage system. The clock chip also presents a high frequency, high accuracy pulse which is used as a trigger for high speed data acquisition.

The second cluster of instrumentation is located in the Wheel Slide Protection (WSP) cabinet.

The majority of the system's outputs are located in the second cluster along with the interface to the WSP unit. The WSP system provides a varying analogue voltage which is representative of vehicle speed and is presented at a series of test pins on the front of the unit. Clearly the instrumentation must not interfere with the WSP system and so an isolation amplifier is used. The unit selected is the same isolation amplifier as that previously described for use in the AETW connection. While the first cluster of instrumentation makes use of a full, external, 16-bit analogue to digital converter, the second cluster only has this single WSP signal to sample. It also requires this signal to be in a format which can be easily transmitted to the fourth instrumentation cluster. For these reasons, a voltage to frequency converter is used. The AD654 is a device from Analogue Devices which outputs a variable frequency waveform dependent on an analogue input signal. The frequency of this pulse is identified by counting the number of transitions between low and high in a fixed time period. The pulse is digital and so can be transmitted to the fourth instrumentation cluster over a fibre optic link using a simple, transistor driven, media converter.

The third cluster consists of digital inputs. The control system for a Class 508 is essentially a large logic circuit operating at 110 V logic levels. As the digital inputs located in cluster three are a combination of the demands that the driver is making of the vehicle and the way that the vehicle is interpreting them, it is imperative that the instrumentation does not interfere with the original signals. HCPL 3700 optical isolators from Agilent Technologies are used to provide isolation between the train and instrumentation systems. The isolators are rated to 2500 V of isolation and have a very low switching threshold of only 2.5 mA.

The final cluster of electronics used in the instrumentation is fitted to the bogie and is only fitted to one end of the train. This cluster is concerned only with the position of the vehicle and so measuring at one end is sufficient. The cluster is operated almost independently from the rest

of the system, with the only connections being a power supply and fibre optic relay of the GPS and tacho signals. The positioning cluster is a more advanced version of the instrumentation described in the third case study. An inertial measurement unit is augmented with a GPS unit and a tacho signal. The inertial measurement unit uses the same CXL10TG3 tri-axis accelerometer described in the previous work, but uses CRS09-02 gyroscopes from Silicon Sensing rather than the BAE Systems VSGs used previously. The CRS09-02 gyroscopes are capable of $100°/s$ compared to the $50°/s$ of the VSGs. In a railway application where the expected rate of rotation is very small (due to bogie inertial effects) the CRS09-02 units are therefore arguably less suitable however they also have a much lower power requirement than the BAE units which was critical in their application into this larger instrumenation system. The GPS module used is the GPS-622R from RF Solutions which is a high sensitivity, 65 channel, unit that outputs standard National Marine Electronics Association (NMEA) strings. The GPS strings are relayed to the second cluster of electronics using a fibre optic serial interface. In exchange, the tacho signal is relayed back by the same mechanism. Fibre optics are used to minimise electrical connections between the bogie and the body of the vehicle.

The system is designed to operate autonomously, and so records all of the signals from the various sensors to removable storage media; the capacity of the storage is a significant issue. The distributed nature of the system means that the combined processing power of all the nodes is significantly more than is required for the signals being acquired, at the rates that they are being stored. The limiting factor of the system is, therefore, the throughput of the storage device. It is possible to utilise the additional processing capacity within the nodes to perform "oversampling" of the signals. This means that the, faster changing, signals are sampled at an increased rate and then averaged within the local processor before being put onto the bus to be recorded. This is not necessarily done to remedy situations whereby the data rate is lower than the required operating rate for the acquisition of a signal (in line with Nyquist theory), but is

often used to reduce noise or interference in the sampled signal. Essentially, oversampling has the effect of applying a low pass filter to the data to smooth any high frequency effects within the signals thus making the recorded data less susceptible to outlying data points. The sample and data rates for the different types of transducers are shown in Table 6.3.

Table 6.3: Sampling and data rates used throughout the system

Signal(s)	Sample Rate [Hz]	Data Rate [Hz]
Traction current	8192	256
Auxiliary current	8192	256
Voltage	8192	256
Traction energy	8192	1
Auxiliary energy	8192	1
AETW	1024	1
Temperature	1024	1
Body inertial measurement	1024	16
RTCC	1	1
Tacho	8	8
Driver's demands	16	16
Cam shaft position	16	16
GPS	1	1
Bogie inertial measurement	8192	256

The precise selection of sampling and data recording rates shown in Table 6.3 was based on a combination of requirements in relation to the rate of changes in the signal, the spare processing capacity available in the node, and the relationship to other sampling rates within the same node in terms of firmware compatibility. Generally speaking, the system prioritised the signals

relating to energy as this was the primary focus of the work with the other signals being obtained to provide context for the energy related data. A data rate of 256 Hz for the raw current and voltage data was selected to be sufficiently high to identify any expected variations found in the camshaft controlled Class 508 being instrumented. This data rate also exceeds the 200 Hz used by the Naples group in their energy efficiency work [115], although this is not entirely comparable as they were instrumenting an AC system. More immediately, the 256 Hz was selected as the fastest data rate at which the monitoring nodes within the system could reliably operate while still retaining headroom within the instrumentation framework for additional signals, and while remaining compatible with the selected microprocessor. The sampling rate for the energy related signals was selected purely based on the remaining processor overhead. The selection of 8192 Hz allows the selected processor to reliably sample up to 8 channels and perform the down sampling to 256 Hz while simultaneously maintaining communication with the rest of the instrumentation framework. This sampling rate is also supported by the selected analogue to digital conversion unit. The traction and auxiliary energy data rates are significantly lower, but were only included to provide confidence in the current and voltage outputs and their alignment. Subsequently they have proved useful as a means of generating an easily manipulable energy summary over some of the longer term analysis.

With the energy signals fixing the primary sampling rates for the system, the other channels handled by the same processor then have to fit around the established schedule. Largely this involves powers of two so that internal divisions can be performed more efficiently. The vehicle weight and temperature measurements are unlikely to change at any great frequency, hence a 1 Hz data recording rate is both adequate and subsequently manageable. The sampling rate of 1024 Hz was selected for use in the vehicle weight measurement primarily because little was known of the performance of the transducer but also because it happened to fit into the schedule dictated by the 8192 Hz used for the energy measurements. Temperature was also

sampled at this frequency, primarily because having equivalently performing channels simplified the firmware used on the processor. The body mounted inertial measurements were included mainly to determine simple things such as the direction of travel of the vehicle. This too is unlikely to change quickly. The data from the body mounted gyroscope is naturally smoothed by the length of the vehicle body, and so this too has a limited rate of change. It was also also only included to be indicative of direction in the case of a primary inertial measurement system failure. A data rate of 16 Hz was selected as a compromise between required performance and nodal capacity.

The RTCC inherently only changes once per second and as the same chip was used to drive the sampling timings for the whole system there seemed little need to oversample it. Likewise, the GPS signal is pushed over a serial interface once per second by the GPS module. As the data arrives as NMEA strings, the processing for its interpretation is somewhat complex and so oversampling was avoided for the sake of processing headroom. Initially the driver's demands were configured for 1 Hz interrogation, however, a site visit revealed the operation of the camshaft to be significantly faster. The sampling in this node was therefore updated to 16 Hz based on the observed rate of change, the driver's demands were at the time driven from the same clock and so these were also modified.

The rates used in the bogie inertial measurement systems were selected based on minimal requirements established from [40] / [66] and [108]. They also take into account the maximum operational rates of the nodes as the bogie inertial measurement unit consists of 6 of the 8 channels that can be managed at these rates. The rate of 8 Hz was selected for the tacho signal based on laboratory testing. A tacho module (part of the wheel slide protection system) takes square pulses from a toothed wheel transducer and produces a continuously variable voltage output. This was interrogated by the node. Laboratory based experiments into changing the

frequency of the square wave, within likely parameters based on the number of teeth, wheel size and standard vehicle acceleration of less than 1.1 g, identified a maximum response which could be observed with a sample rate of 8 Hz.

The system has three output points, two of which are repeated at both ends of the train. These first two outputs monitor the communication bus which links the first three clusters of electronics. The contents of the bus are monitored and everything being transmitted is sent to both output points. The first output is a fibre optic serial interface to a laptop computer which can be installed in the vehicle body as required. This is similar to the systems used in the previous case studies and is intended to be used for commissioning, testing and live monitoring during dedicated trials.

The second output relates to the autonomous nature of the system. The instrumentation is designed to function for up to a year with minimal user intervention. To achieve this, the system records the data to removable storage media in an efficient binary format to maximise capacity. The user interaction should be limited to changing these storage media. The selected technology is SD card. Using 16 GB SD cards, it is possible for the system to operate for up to 42 days without loss of data. After this time, the system overwrites the oldest stored information. The vehicle will return to the depot for routine maintenance every 19-21 days; the storage media are exchanged when this happens.

The third output is related to the cluster of instrumentation located on the bogie. The bogie must, as far as possible be kept electrically isolated from the body of the vehicle. This means that the installed instrumentation system should not provide any additional electrical paths if possibly avoidable. For this reason, the instrumentation located on the body is not included on the main data bus. Furthermore, the bogie mounted components of the instrumentation are only present at one end of the vehicle and therefore adding bus capacity to support them would be

inefficient. Instead, the bogie mounted instrumentation has its own local storage which is then aligned to the rest of the vehicle's data outputs using the tacho, GPS, longitudinal accelerometer and yaw rate gyro signals which are also stored by the instrumentation attached to the vehicle body. The data rate from the bogie mounted transducers is such that the storage requirements are the same on the bogie as on each end of the train and so the same SD card based storage solution can be used and the cards exchanged at the same interval.

6.3.3 Electronics

As described in the overview of the instrumentation, the core of the system is made up of a number of independent embedded microprocessors linked together with a data bus. There are six microprocessors at the end of the train including the bogie mounted instrumentation cluster and five at the other. To simplify hardware design, and to maximise firmware code re-usability, the same type of microprocessor was used throughout the system. The microprocessor selected was the PIC24HJ128GP502, described previously in chapter three. The PIC24HJ128GP502 has a number of peripheral components, more than are required for any single circuit board and enough to cover the combined requirements of the system. There are two serial interfaces for communication between the clusters of electronics housed in the WSP case and mounted on the bogie, and for the link to the vehicle body. There is an I^2C interface to communicate with the RTCCs and a digital port expander used in the third cluster. There are two SPI interfaces to communicate with the analogue to digital converters and the SD cards. Finally there is a CAN interface to the main system data bus. The peripheral buses used by each circuit board are summarised in Table 6.4.

Table 6.4: Peripheral buses used throughout the instrumentation

Cluster	PCB	Purpose	Serial	I^{2}C	SPI	CAN
1	1	Energy	-	1	2	1
2	2	Tacho	2	-	-	1
2	3	Logging	-	-	1	1
2	4	Laptop	1	-	-	1
3	5	Digital I/O	-	1	-	1
4	6	IMU	1	1	2	-

Reusing a single processor type means that some instances are only lightly used while others operate almost at the capacity of the processor. The most heavily loaded processor is the one running the IMU mounted on the bogie. This module must service an 8 channel analogue to digital converter, a real time clock, a GPS serial stream, a fibre optic media converter, a tacho input, and removable storage in the form of an SD card. The PIC 24HJ128GP502 can be operated at up to 80 MHz which is sufficient to manage all of these components. Many of the input / output pins on the device can be allocated to different peripherals which simplifies the hardware development, and many of the pins are tolerant to 5 V inputs despite the device operating on 3.3 V. Although the PIC24HJ128GP502 has a built in peripheral to allow it to use a CAN bus, an external driver chip is still required.

The hardware described in this chapter is being installed onto a passenger carrying vehicle for a period of up to a year. Before this can happen, the system must undergo an approvals process to ensure that it is safe. Part of the approvals process is concerned with ensuring ElectroMagnetic Compatibility (EMC) in line with BS EN 50155 [88] and BS EN 51021-3-2 [89]. There are two parts to the EMC testing; firstly the system has to show that it is resilient to electromagnetic interference and secondly it has to show that it will not emit interference above threshold levels.

The four instrumentation clusters are housed in four metal boxes which provide electromagnetic shielding. The largest EMC risks occur at the links between the instrumentation clusters. To combat any interference induced over the 24 V power links between the clusters, each printed circuit board has a section dedicated to filtering and EMC protection. Diodes and capacitors are used between the 24 V and 0 V lines to clamp any large voltages to an acceptable level, and also for transient voltage suppression between both lines and the shield which acts as a ground plane for the system. The power lines then pass through a common mode choke before being smoothed by a combination of resistor / capacitor and inductor / capacitor circuits. A diode is also used to further protect the power supply. A schematic of this is shown in Figure 6.2. A similar protection system is used for the CAN data lines. The differential pair is passed through a common mode choke and a combination of capacitors and bi-directional diodes provide transient voltage suppression for the system.
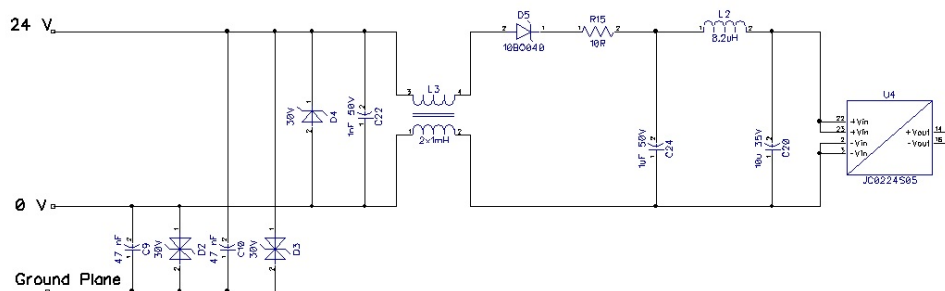
Figure 6.2: Standardised EMC protection system for power circuitry

The other points where the system is particularly sensitive to emissions are those where physical connections are made with the existing train systems. There are three types of these connections: the AETW interface, the tacho interface and the interfaces to the driver's demands and the camshaft system. Isolation is used to ensure that any interference in the instrumentation system cannot be transmitted to these connections. Connecting wires are also twisted, shielded and kept as short as possible.

As mentioned previously, the core of the AETW and tacho isolation systems are based around the AD202JY isolation amplifier from Analogue Devices. The unit itself provides transformer based isolation with a separate power supply available on the instrument side which can be used to power any active filtering. In addition to the isolation amplifier and the wires being made as short as possible, extra components are used to protect the AETW module in the event of wires becoming disconnected. A schematic for the AETW interface is shown in Figure 6.3. Even with this level of protection, it was felt that the risks associated with direct interfacing to the equivalent braking system, the AETB module, were too high and so no connections were made to it.
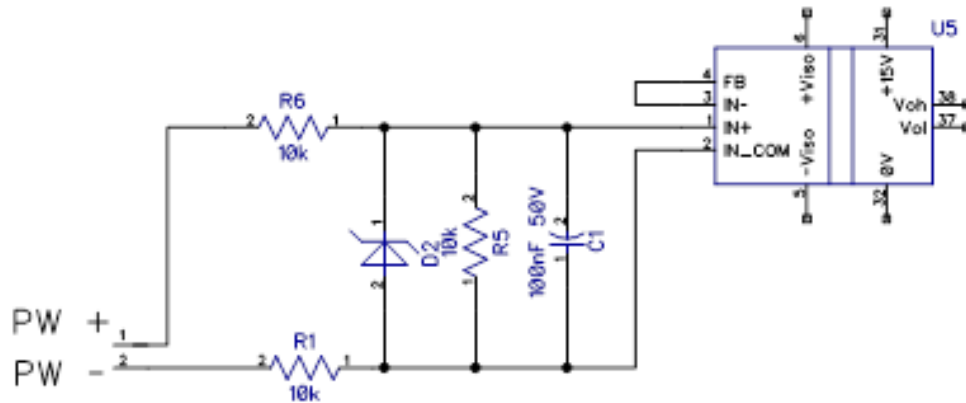
Figure 6.3: Vehicle weight transducer interface

As the driver's demands and camshaft control systems are essentially binary in nature, isolation is provided optically through the use of HCPL3700 units from Agilent. This prevents anything being transmitted from the instrumentation system onto the control lines. Even without this, the instrumentation operates at 3.3 V while the control system uses 110 V logic levels so it is unlikely that any effects would be significant.

The 110 V logic is actually rather challenging because while it is nominally 110 V and 0 V, in practice anything below 30 V is considered to be "off" while anything above 60 V is considered "on". The levels representing "on" can actually be as high as 150 V and spikes of up to 600 V

have been reported.

A network of Metal Electrode Leadless Face (MELF) resistors, a diode, and a Zener diode is used to manage the switching voltages. This is shown in Figure 6.4. The 1N4148 diode is simply present to protect the opto-isolator. The Zener diode was selected to modify the switching voltage of the opto-isolator. By reducing the presented voltage by 12 V, the predicted signal threshold voltages of 30 V and 60 V become more manageable within the supported current range of the opto-coupler. MELF resistors were used because of their capacity to accommodate transient voltage spikes. The values of the resistors were selected such that the combined series resistance and the modified threshold voltages aligned with the characteristics of the opto-coupler.
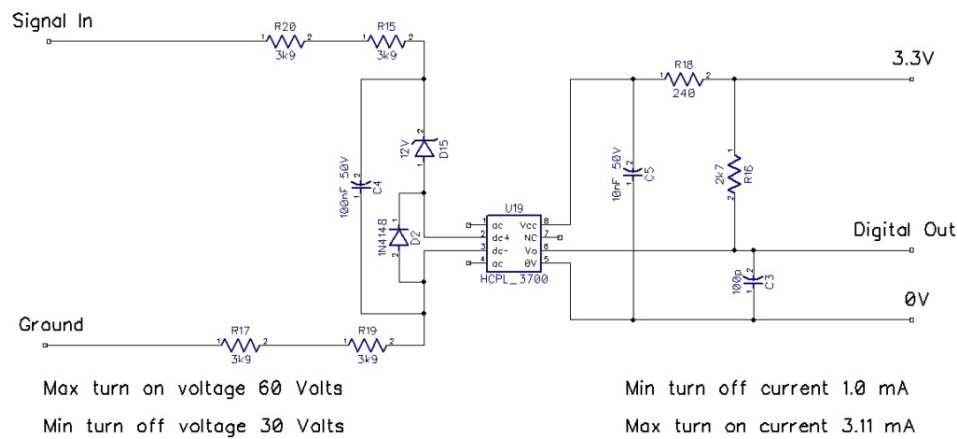
Figure 6.4: 110 V logic interface for driver's demands and camshaft control signals

6.3.4 Firmware

The instrumentation described in this case study is distributed over a whole vehicle. At each end, there are either three or four clusters of instrumentation. Each of these clusters in turn makes use of between one and three microprocessors. As described previously, this division is based more on functionality than capacity. In total, there are five microprocessors at one end of

the vehicle and six at the other. Although some components could be reused, this meant that six different types of firmware have had to be developed.

Three of the six different types of firmware are concerned with the acquisition of data from external analogue or digital sources and transferring that data onto the CAN bus running between the different instrumentation clusters. A further two types of firmware focus on transferring the contents of the CAN bus to other media, for storage or dynamic viewing. The final type of firmware combines aspects of the others in as much as it collects data from various inputs and then records it to storage media directly rather than involving the CAN bus.

Many of the components of these firmware are similar to those used in previous instrumentation projects. The firmware for the SPI interface to the LTC1859 analogue to digital converter, for instance, is similar to that used in the energy monitoring work described in the first case study. In this project, a similar interface is used in the first and fourth clusters of instrumentation, i.e. those concerned with the energy monitoring and inertial measurement systems. While the underlying firmware is similar in its use of interrupts to control timing, and SPI communication, there are differences in the way that the defined functions are used. As shown in Table 6.3, the current and voltage channels are sampled more frequently than the AETW, temperature and inertial systems. To achieve this, the interrupt routine responsible for managing the sample requests was modified to sample all of the high frequency channels and one of the low frequency channels or a dummy channel, each time the interrupt was called. A simplified implementation of this is shown in Figure 6.5. The remaining components of the interrupt routine, concerned with oversampling and averaging the running sums for each channel, remained unaffected.
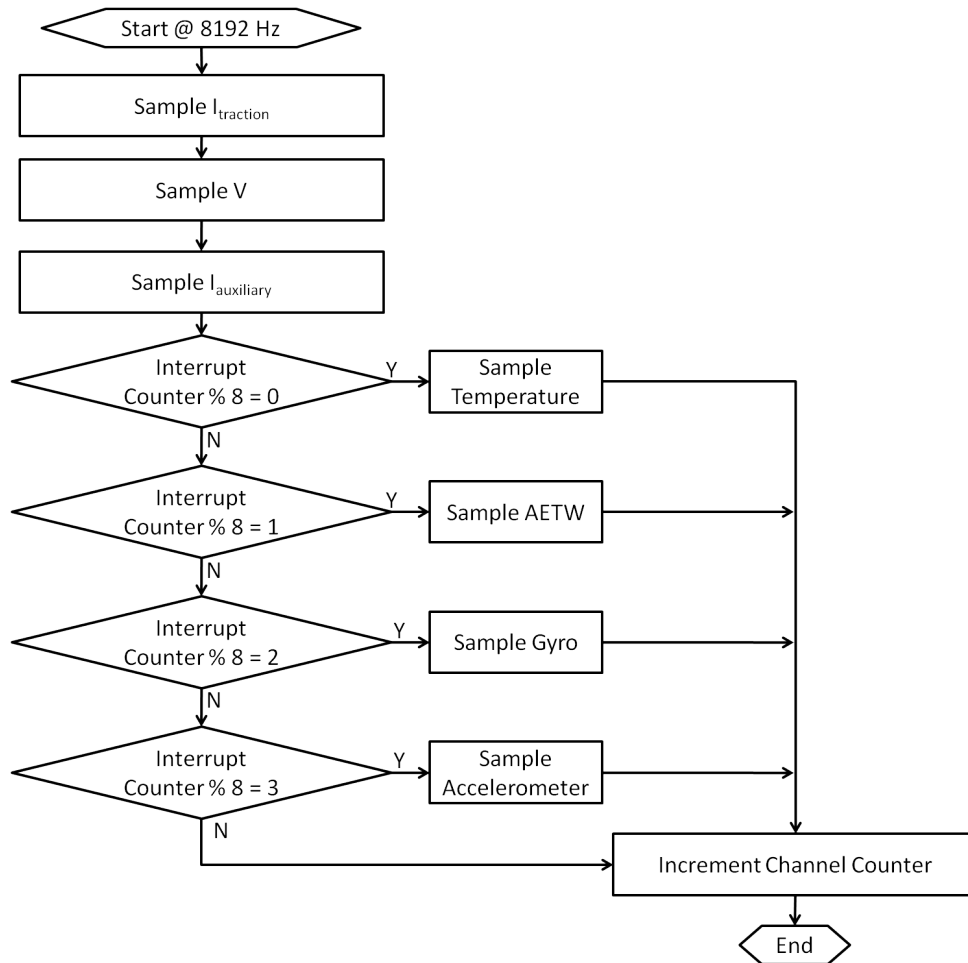
Figure 6.5: Flow diagram of interrupt routine

The two most major advances in the firmware used by this system compared to those discussed previously are the inclusion of the CAN bus and SD card systems.

The PIC 24HJ128GP502 CAN peripheral supports 32 DMA based input/output buffers. Of these, buffer zero was configured as an extremely high priority output buffer and buffers 16 through to 31 were configured as input buffers. First In First Out (FIFO) queues were applied to all of the used buffers, buffers 1 to 15 were not used. The peripheral has the ability to pre-filter data arriving at the input buffers, discarding data that does not meet a particular pattern or directing data that does. This feature was essentially disabled in favour of retaining all incoming data for inspection. In normal operation, only two of the sets of firmware need to read CAN

data, while only three sets need to transmit. The structure described above was, however, used in all cases so that CAN based debugging facilities could be developed. These included power up and status reporting and the ability to set system clocks remotely using CAN messages. A computer based CAN bus interface was used to access the CAN bus.

The CAN bus was operated in standard mode, rather than extended, partially due to the standard mode being suitable for the complexity of the system and partially due to fragility in the Microchip ECAN peripheral. With the extended mode disabled, a standard 8 word data frame could be developed. This is shown in Figure 6.6.

	SID [13:2]		0x0	0x0
	0x0000			
0x00		0x0	DLC [3:0]	
Data byte 1 [15:8]		Data byte 0 [7:0]		
Data byte 3 [15:8]		Data byte 2 [7:0]		
Data byte 5 [15:8]		Data byte 4 [7:0]		
Data byte 7 [15:8]		Data byte 6 [7:0]		
	0x0000			

Figure 6.6: Standard mode CAN message data frame

The first word of the data frame is largely concerned with the 11 bit Standard Identifier (SID). In a CAN system, these have several purposes. Firstly they provide a description to the user of the contents of the message. In this system, this corresponds to reference to the cluster of electronics that the message originated in and the type of sensor or input that it is associated with. The second function is related to message priority. Certain bits of the SID are used to judge the importance of messages during conflict / retry events with higher priority messages being re-transmitted first. In this system, timestamps and energy related signals were given priority over lower frequency inputs such as those from the AETW system and temperature

sensors. The final use of the SID is related to the filters that can be applied to messages being received to discard them or direct them to specific input buffers. In this system, message filtering was largely disabled and so this was not considered.

As the system was operating in standard mode, the second word is not used. This word is used to manage the identifier if the system is operating in extended mode and is, therefore, a required part of the standard CAN message structure. The third word contains the 4 bit Data Length Code (DLC). This is used to indicate how many of the following words will contain the data payload. Reducing the data payload allows fewer words to be transmitted per message, thereby saving bandwidth. The overhead component of the message is, however, fixed. In this system each SID may be associated with several signals to obtain the greatest possible efficiency from the message structure.

The data frame allows for up to 8 bytes of data arrayed as four 16 bit words. The LTC1859 analogue to digital converter, used in the first instrumentation cluster to return the current and voltage signals, is a 16 bit device which returns high and low bytes. These are ideally suited to packing into the words of the CAN message. Not all of the messages sent made use of the full capacity of the data frame. In these cases, the number of data bytes could be reduced. For example, in the first cluster of electronics, one message was dedicated to the voltage, traction current and auxiliary current measurements. Each of these consists of 16 bits, i.e. one word. The final word is unused and so does not need to be transmitted. Conversely, the locally computed energy measurements are generated by multiplying two 16 bit numbers together and so could require up to 32 bits each. The message containing both the traction and auxiliary energy makes use of all four words.

Table 6.5 shows the CAN message structure used along with the frequency that each type of message was transmitted. The sources of each type of message are shown in Table 6.6. Data

cells that are empty in Table 6.5 are unused and so for these messages the DLC is reduced so that these bytes are not transmitted. As discussed, the system makes use of the majority of the capacity of each message identifier that it uses. The major limitations on this are concerned with the microprocessor that each signal is sourced from and the data rate that the signals are being recorded. In cases where more than one message is transmitted by a cluster at a given frequency, it would not have been possible to compress the information into fewer messages. In most cases, each signal is allocated a number of bytes within a message. The only exceptions to this are the GPS messages where a single bit is used to indicate North vs. South or East vs. West. These messages could have been given their own byte, but this would have been inefficient, particularly as it would have involved an additional word rather than just an additional byte, and so the information was packed into the top (unused) bits of the GPS position messages. The other major inefficiency is in the reporting message where an odd number of bytes is used. This is acceptable because the reporting messages are only used during system power up and debugging.

SID 0x004 is used for the low frequency signals being measured at instrumentation cluster one. Table 6.5 shows that two bytes are reserved for AETB messages. These would be the outputs of the air pressure transducer connected to the braking system; however, it was decided not to connect to this transducer for safety reasons. This decision was made after the system was developed and so the capacity was retained despite the additional overhead. The hardware to support this transducer was also retained.

This message structure defines the communications between the system nodes that will occur over the backbone data bus. Essentially, the limitations and specification of this message structure constitute the communications schema used within the insturmentation. The schema has been defined to prioritise different types of messages such that different data rates, or signal

sources are packaged into different message descriptors. This was done in a customised format for this instrumentation system, but could be extended to allow the addition of different node types. In the first instance, additional nodes to use the spare sections in the schma, but ultimately this may be modified so that the data descriptions were more general to allow a broader transducer range. In this case, messages would likely be multi-part, or would use some bits to indicate a unique transducer ID. It is likely that the different levels in the schema would still correspond to sampling rates or signal priorities. Some of the other spare capacity in the schema could be used for the event based ouptus of processing nodes which, as prevoiusly discussed, would (non destructively) extract data from the bus in accordance with the schema and then transmit the results of the processing back onto it, or directly to the user.

Table 6.5: CAN messages used throughout the instrumentation network

SID	Name	Bytes	Description								Frequency [Hz]
			Data 1	Data 0	Data 3	Data 2	Data 5	Data 4	Data 7	Data 6	
0x000	Control		Variable size control message								
0x001	RTCC_Time	6	Year	Month	Day	Hour	Minute	Second			1
0x002	Power_1	6	Voltage		I traction		I aux				256
0x003	Inertia	4	Gyroscope		Accelerometer						16
0x004	WeightTemp	6	AETW		AETB		Temperature				1
0x005	Power_2	8	Energy traction				Energy aux				1
0x006											
0x007	Tacho	2	Count for 1/8 sec								8
0x008	GPS_Time	6	Year	Month	Day	Hour	Minute	Second			1
0x009	GSP_North	6	North (top bit=N/S)				Speed				1
0x00A	GPS_East	6	East (top bit=E/W)				Heading				1
0x00B											
0x00C	Digital	2	16 digital inputs								16
0x00D											
0x00E											
0x00F	Reporting	3	Board ID	RCON		State					

Table 6.6: Instrumentation cluster sources for individual CAN messages

| SID	Name	Message Source
0x000	Control	External
0x001	RTCC_Time	Cluster 1
0x002	Power_1	Cluster 1
0x003	Inertia	Cluster 1
0x004	WeightTemp	Cluster 1
0x005	Power_2	Cluster 1
0x006		
0x007	Tacho	Cluster 2
0x008	GPS_Time	Cluster 2
0x009	GSP_North	Cluster 2
0x00A	GPS_East	Cluster 2
0x00B		
0x00C	Digital	Cluster 3
0x00D		
0x00E		
0x00F	Reporting	All

The other significant difference between the instrumentation being described in this case study and the previous ones is that in this case the data is being recorded to removable solid state storage media in the form of SD cards.

Secure Digital (SD) cards measure 32 x 24 x 2.1 mm and can be used for a variety of digital storage applications. There are three main generations of SD card; SD, SDHC and SDXC. Regular SD cards have a maximum capacity of 4 GB, this corresponds to approximately 10 days in this application. SDHC cards operate up to 32 GB, while SDXC provide greater capacity, up to a theoretical 2 TB. The instrumentation being described in this case study uses 16 GB SDHC cards which give a system runtime of 42 days. The vehicles return to the depot every 19 to 21 days for routine maintenance.

PIC microprocessors, such as the PIC 24HJ128GP502, do not provide native support for SD cards, however, it is possible to communicate with them using commands issued manually over an SPI interface. Microchip have wrapped up some of these commands into a series of SD card utility functions. These functions, however, are designed to be used to access a FAT, or FAT32 file system operating on the SD cards. Such systems, when operated by a microprocessor, are not fast enough to handle the data throughput generated by this system. Instead, a custom library of functions has been developed. This library allows the microprocessor to access the raw memory sectors of the SD card without the file system overheads.

The SD card association was founded by three organisations; Matsushita Electric Industrial Co., Ltd. (Panasonic), SanDisk Corporation and Toshiba Corporation. The group was founded in 2000 and in 2006 produced version 2.0 of the physical layer specification component of the SD card specifications [116]. Figure 6.7 is taken from this document. It describes the initialisation process for SD and SDHC cards. The system uses a series of 48 bit command messages which must be issued in accordance with the flow chart and at specific timings. The timings are

described in maximum and minimum terms in the specification but can vary depending on card manufacturer and capacity. This system was developed using genuine 16 GB SanDisk class 2 media.
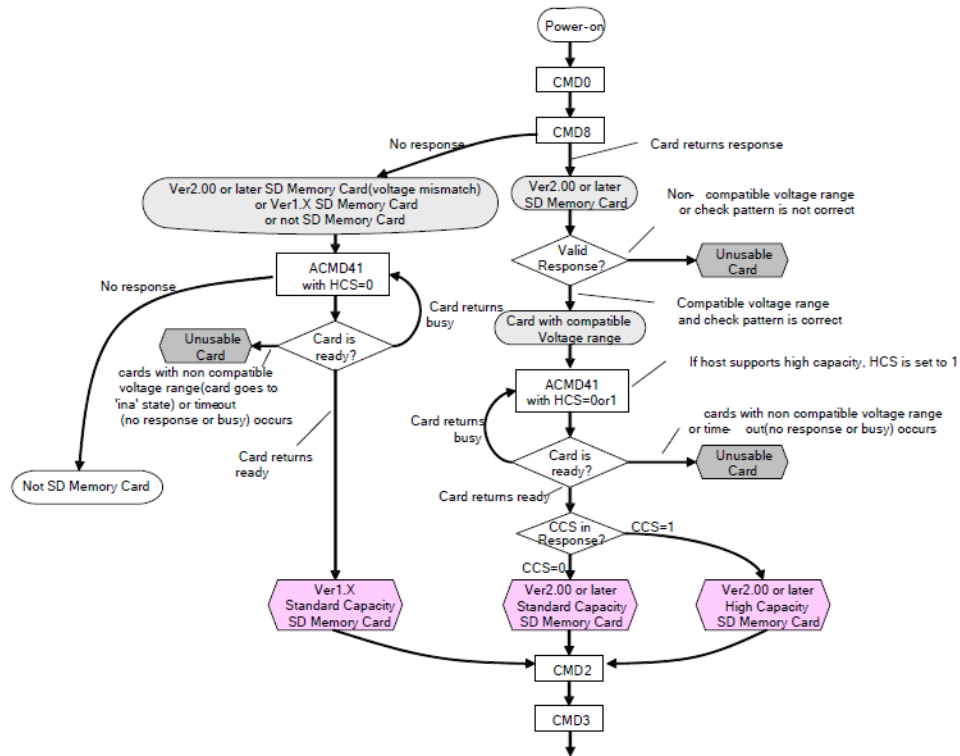
Figure 6.7: SD Card initialisation flow chart

To initialise the card, the system must first send Command 0 to prepare the SD card. The second command required is Command 8 which checks the condition of the card and prepares it for use with the extended command set. The third command required is the Application specific Command 41 which is used to identify the host capacity support of the card. This command is actually a combination of Command 55 and Command 41, were Command 55 is used to prepare the card for an application specific / extended mode command. Command ACMD41 is commonly repeated until the card has completed its initialisation routine or is deemed incompatible. The time allocated to this repetition is one of the significant factors in the range of capacities of SD card supported by a system. Finally, once the card has completed

initialisation, the system must issue Command 2 followed by Command 3. These commands ask the card to identify itself and its address information to the system.

Once the SD card has been initialised, the system issues Command 24 to inform the SD card that it wishes to write a block of data to the system. The command also takes an address which informs the SD card which location the data should be written to. In SDHC cards this is a block address where a block consists of 512 bytes. Upon confirmation, the system then issues a start block token and proceeds to transfer data one block at a time from its local buffers. The SD card accepts the continuing data stream as a burst mode transfer. Once the transfer is complete, the system issues two bytes to form a cyclic redundancy check and waits for confirmation from the SD card that the data has been stored. This can take up to 1 ms.

By not using a FAT based file system, the overheads associated with file management are negated and the raw capacity of the SD card can be realised. This system uses 16 GB SD cards with a custom file system which provides sufficient capacity for operation of up to 42 days. The custom file system is based on raw binary transfers to specific locations on the SD card. The locations are selected based on the timestamps associated with the data. The timestamp obtained from the RTCC is in the format of YYYY/MM/DD/hh:mm:ss. This is converted into a number of seconds, assuming there to be 31 days per month. This number of seconds is then multiplied by the 4 blocks of 512 bytes required per second before a modulus based on the number of blocks available on the card is applied to ensure that the address wraps around to the beginning once it goes past the end of the card. This scheme is not 100 % efficient as it assumes all months to be 31 days long and also allows for 24 hour operation. The scheme does, however, allow location on the card of any specified timestamp and the ability to avoid overwriting of previously recorded data without the need for a separate index to be maintained.

One type of node that uses these two technologies (CAN and SD) more than any other is the

logging node located in instrumentation cluster two. This node is connected to the CAN bus but, other than power up reporting and debugging, does not put any data onto it. Instead the node monitors the data bus and records all relevant, i.e. not debugging, data to an SD card.

Although the node does not need to interpret the data, some internal processing is still required. The data coming from the CAN bus is a continuous stream of variable frequency message types. The SD card data format requires that data be divided into 1 second intervals and that it is transferred to the SD card in 512 byte sections. To achieve this, the microprocessor uses a series of internal data buffers. The initial CAN receive buffers are unloaded to a series of FIFO buffers as part of the CAN interface. The data in these FIFO buffers are then decoded and placed into one of three 512 byte ring buffers. If a buffer is filled, the system advances to use the next buffer. Should a timestamp message be detected, the remainder of the buffer is cleared and the system advances to the next buffer. This ensures that timestamps always align with the start of one of the buffers. When a buffer becomes full, it is put into a queue to have its contents transferred to the SD card. If the start of the buffer contains a timestamp, the SD card address is also updated. This is summarised in Figure 6.8.
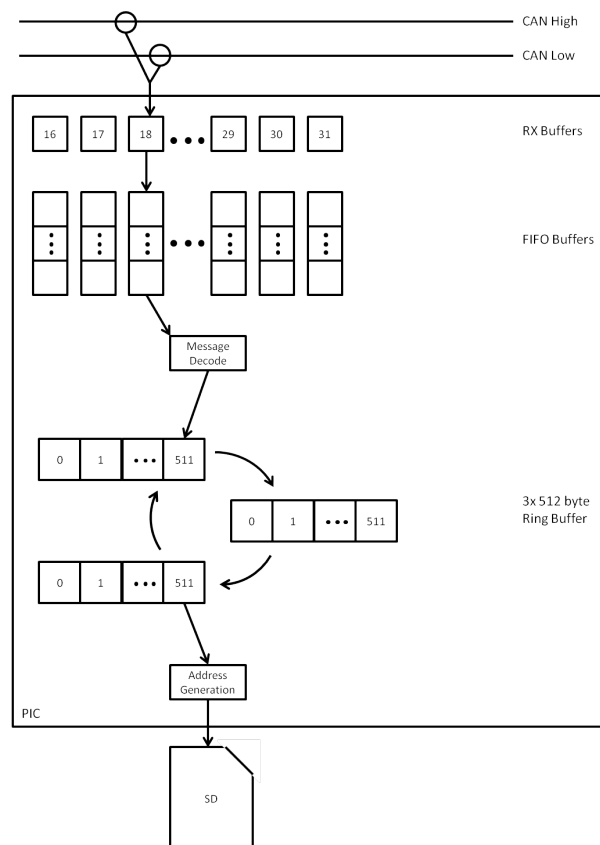
Figure 6.8: Microprocessor internal buffers for SD card storage system

6.3.5 Computer interface

In each of the three previous case studies, the system has had a single point of output to a computer which has recorded the data. In the cases of the energy and conductor shoe monitoring projects, the system had some form of real time visual indication of output. In the case of the inertial measurement work, the data rates and processing required were prohibitive of any real time visualisation solution and so the system was centred on storage and offline data processing. In the work being described in this case study, both situations apply. The data rate was too high, and the variety of signals was too broad, to allow a full real time visualisation. Furthermore, the data acquisition equipment was distributed in three locations over the length of the vehicle. It was, therefore, necessary to store the data for processing later. Some visualisation was, however, beneficial as the system was installed for an extended period and so some commissioning testing was required.

Data visualisation was achieved by taking a fibre optic clone of the CAN bus contents and transmitting it, along with a framing message, to a PC mounted in the vehicle body. A circuit board responsible for this activity was located in each of the second instrumentation clusters, i.e. at both ends of the vehicle. The results from the inertial measurement unit attached to the bogie were not relayed in this manner. The data being received by the computer located in the vehicle body was stored in raw form to a data file on the computer's hard disk and also processed in real time for display. The raw data recording was a direct copy of the serial stream, however, the system was not fast enough to decode and present all of the data in real time. Instead, the system sampled the serial stream every 100 ms to record the latest value of each of the signals. These were then used in the visualisation. While this method discards huge volumes of data, the visualisation was only ever intended to be a commissioning tool and so the reduced output quality was acceptable. An example of the real time visualisation software output is shown in

Figure 6.9.
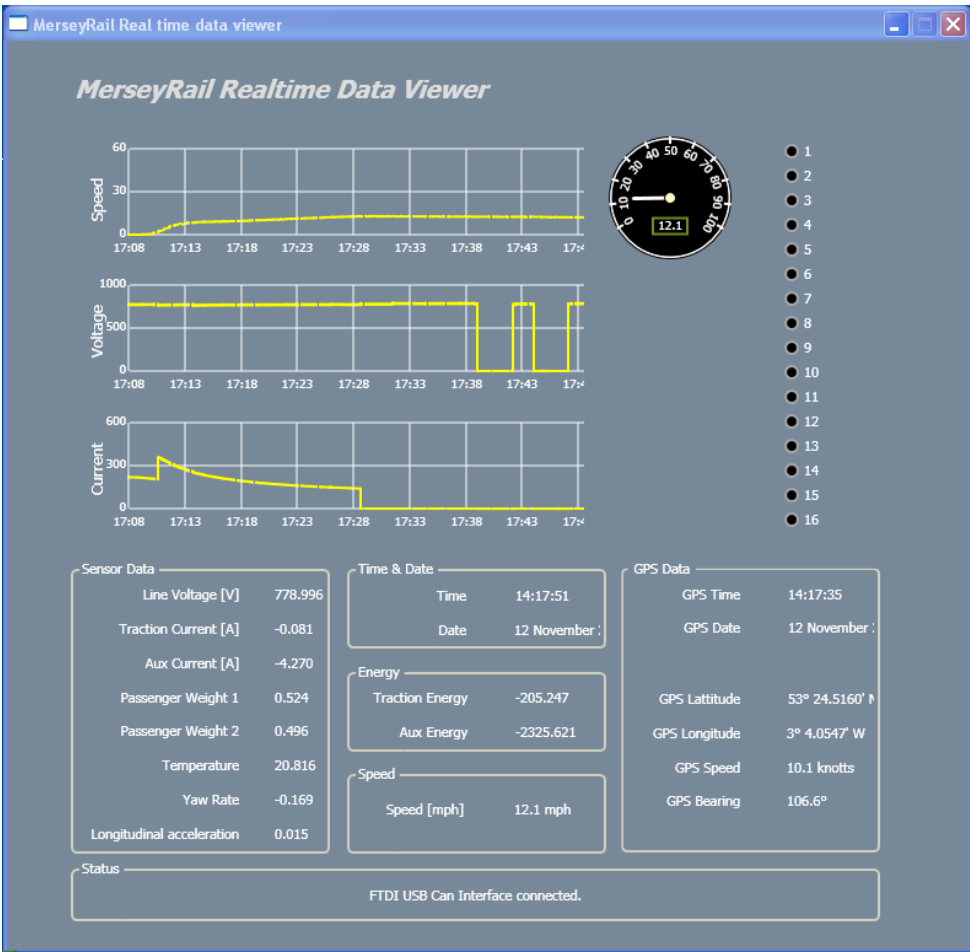
Figure 6.9: Real time visualisation software for energy monitoring instrumentation

As explained previously, in the work being described in this case study, the real time visualisation was only used for testing and commissioning of the system. The main data outputs were in the form of three 16 GB SD cards taken from the vehicle each time it returned to the depot for routine maintenance inspections.

While possible, reading the SD cards via a microprocessor communicating with a PC would be inefficient and slow. A better solution is to use an SD card reader attached directly to a PC. Unfortunately, Windows based computers have standard methods of dealing with SD cards which mask the raw data that the system has stored on them. A solution to this is to use a

microprocessor to erase certain key memory addresses on the cards that Windows uses as part of its file system. Once this is complete, Windows no longer knows how to interpret the card, thus leaving the raw contents accessible. In this way, it is possible to download the entire binary image of a, class 2, 16 GB SD card in approximately 40 minutes.
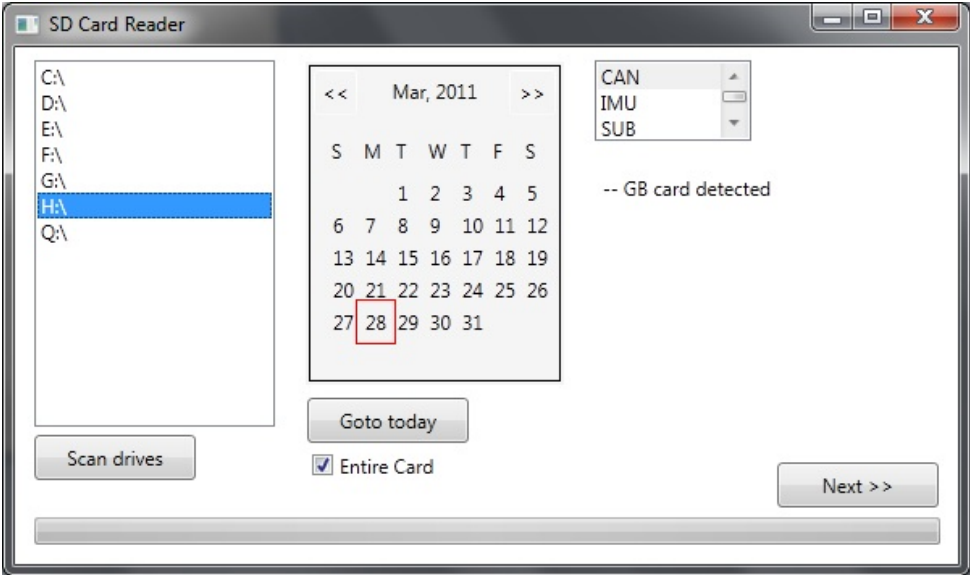
6.3.6 Software

The software used in the work being described in this case study falls into four sections: the real time data viewer, a system for reading the raw SD card data back to a computer, a system for extracting the raw data into a format suitable for processing, and processing / analysis tools.

The first of these four pieces of software, the real time viewer, was an extension of the system used in the energy monitoring case study. The software used the same underlying systems to read the serial data stream, convert it to an intermediate format and then present a subset of that intermediate data for real time display. The system was expanded to include the additional signals present in the more complete version of the instrumentation. This was reasonably straight forward given the lower data rate of these signals compared to the current and voltage data that was already being transferred.

The remaining pieces of software were concerned with the main data repositories, i.e. the three SD cards located at either end and on one bogie of the vehicle. The data transfer and data extraction systems combine together in a multiple stage process to extract the data from the SD cards and present it in a format suitable for analysis.

The first of these pieces of software is used to transfer the data from the SD card onto a computer, and also to provide some indexing of the data in an effort to aid extraction of specific information. The software, shown in Figure 6.10, was written in C# and uses Microsoft's Win32 SafeHandles to access the SD card device at a lower level than any file system stored on it. The

software requires that the user select the drive associated with the SD card and also specify if the card contains data from the CAN bus or the inertial measurement unit. The user should also specify a particular date to retrieve, or tell the system to extract all of the data on the card.
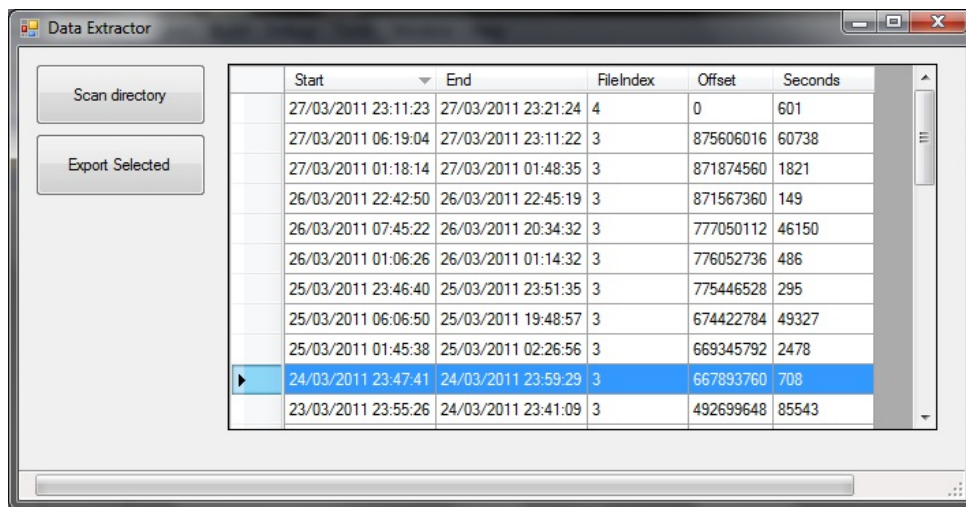
Figure 6.10: Software used to extract raw data from SD cards

Once the user has provided sufficient information, the software determines the card capacity and proceeds with transferring the data to the computer. This process involves transferring sections of data from the card to a local buffer and then advancing through the buffer, one second at a time, checking that the data contains the appropriate time stamp information. If a particular day has been requested, the system checks that the data comes from that day rather than being any previous data left on a card from previous experiments. Once the system has verified that the data is valid it is passed to a file management routine which stores the binary data alongside an index file which contains the timestamp information and the number of bytes into the binary file at which the data associated with that timestamp can be found. The binary files are limited in size to 1 GB. The file management routine automatically splits the data stream into files of this size.

The second phase of the data extraction process uses the binary and index files, generated by the

219

first phase, as inputs and presents the outputs as a series of comma separated files containing timestamped data. A file is produced for each CAN bus message ID or sampling frequency in the case of inertial data. As with the binary files, the size of these files is limited; in this case, the files are divided at a threshold of 100 MB. This apparently low number is to facilitate passing the data to further analysis software.

The software initially asks the user for a directory containing binary and index file pairs. There can be more than one pair in the directory. The index files found in the directory are then scanned to check that they contain timestamp / offset pairs. Each timestamp found is checked against the previous one to identify continuous sections of data. A table is then populated with the start time, end time and duration of each continuous section, along with the corresponding binary file index and the offset into the file at which the section starts. Figure 6.11 shows an example of this table for a small, continuous data set.

Figure 6.11: Software used to extract sections of data from raw data repositories

The software then requires the user to select the row, or rows, of data that they are interested in before requesting that it is extracted. The data extraction process operates on each selected row in the table. The appropriate binary file is opened and the timestamp indicator is used to

differentiate between data recorded from the CAN bus and the inertial measurement unit. The system then activates the appropriate decoder module.

Both the CAN network and inertial measurement decoders start by generating their required output file structures. The decoders then advance through the section of data in one second intervals. If the output files become larger than 100 MB they are split into multiple files. The decoder for the inertial measurement unit data uses deterministic sequence of the data in the buffer to apply the appropriate extraction and calibration routines as it parses each 4096 byte section. The CAN network decoder reads the first entry which contains a standard identifier and a number of data bytes and uses this information to parse the associated number of data bytes before reading the next identifier until the whole second has been processed.

The use of these two pieces of software allows data to be extracted from the SD cards either for a specific date or for the entire range of data stored on the card. This data is presented as a series of files in comma separated variable format. Each file corresponds to a CAN bus message identifier, or a data acquisition frequency in the case of the inertial measurement data. This format is well suited to being a transitory step towards data analysis but is a poor long term storage solution as there is no meta-data associated with the files and separate index files are required to document the data format.

It is anticipated that the data being obtained will not just be used to answer specific questions, but will be used in a variety of different projects. To date, planned uses of the data include the driver style and energy use analysis, DC network energy loss analysis, DC network capacity measurement, track condition monitoring, transition analysis, switch and crossing condition monitoring, and rolling stock profiling. The University of Birmingham is developing a standardised data repository to hold key datasets, such as this one, so that they can be made available to a number of different projects. Ultimately, data processing algorithms will be included in these

databases such that they will become processing tools rather than mere repositories. Once this occurs, data will not need to be exported to be processed. This system, however, requires the data processing algorithms first to have been prototyped and developed and so is not applicable to this project.

The final section of software development is the data analysis software. Analysis is largely performed using MATLAB from Mathworks. MATLAB is particularly well suited to importing the comma separated outputs of the previous two pieces of software, and is ideal for data manipulation and algorithm development. MATLAB works as an interpreted language which means that commands are passed to an interpreter which performs the processing. These commands can be supplied in the form of scripts which are the MATALB equivalent of source code.

6.3.7 Mechanical construction

Unlike the previous three case studies, the instrumentation being described here is designed to be fitted to an in-service vehicle for an extended period of time. Both aspects of this are different from work carried out previously. In all of the previous instrumentation work, dedicated testing vehicles without passengers have been operated for specific periods of time, and generally in such a way as to maximise the spacing to other traffic. This approach is generally favoured with academic projects as the train operating company and the rolling stock company are not usually heavily involved in the development of the instrumentation, and yet still carry the risk for the vehicle. Undertaking testing in this manner minimises the time that the vehicle operates for, and also minimises the contact that the test vehicle has with the public.

In order for the instrumentation being described in this case study to capture the energy use associated with different driving styles, it was necessary for it to be installed onto a system that would operate in realistic conditions. This means that it was not possible to install it on an

isolated test vehicle because the driver would be aware that they were under observation. Also, the lack of passengers would affect the mass of the vehicle and thus its performance.

For the system to be installed on a passenger carrying vehicle and to operate in amongst other traffic, a number of procedures had to be completed. Various assurances of electrical and electromagnetic compatibility had to be obtained as described in a previous section. In addition to these electrical interfaces, the mechanical interfaces between the vehicle and the system also had to be approved. This approval was originally to be in the form of Vehicle Acceptance Body (VAB) certification, but ultimately came in the form of a "minor modification approval" document confirming that an external consultancy (Atkins) had checked the system for compliance with the relevant railway group standards. A copy of the approvals document can be found in Appendix A. The procedure required for such an endorsement required significant input from the Train Operating Company (Merseyrail) and the ROlling Stock Operating COmpany (Angel Trains). This level of required input meant that the development of the instrumentation, particularly the mechanical aspects and maintenance procedures, was rather more industrially lead than previous projects.

The approvals process focussed on four main elements:

- Electrical and electromagnetic compatibility as described previously.

- Mechanical construction and attachment of the instrumentation clusters.

- Fire resilience.

- Integration with existing maintenance procedures.

The mechanical construction of each of the instrumentation clusters was reasonably similar. Each cluster consisted of a single box containing all of the instrumentation. These boxes were all of aluminium construction, either custom built where size or shape was difficult, or custom-

isations of commercially available rack mounting equipment cases as appropriate. Metal tape was used to seal each of the boxes. This was done for both environmental and electromagnetic reasons. Where possible the equipment inside each box was mounted using anti-vibration mountings, although this was not suitable for the inertial measurement transducers mounted in the fourth instrumentation cluster.

Two of the instrumentation clusters, one and three, were mounted in the equipment racks hung from the underside of the vehicle body. In these cases the base plates were removed and replaced with plates supporting the instrumentation. Calculations were performed based on the specified load requirements for equipment mounted on the vehicle body found in GM/RT 2100 [104]. The most stringent of these requirements is for the equipment to remain attached when subjected to a longitudinal acceleration of 5 times the force of gravity.

The second instrumentation cluster was located in the top of an existing 19" rack case, also hung from the underside of the vehicle body. It therefore had to meet the same 5 g longitudinal acceleration criteria as the first and third instrumentation clusters. For this cluster, a standard EMC shielded 1 U rack case was used to house the 3 printed circuit boards of instrumentation. The rack case was modified, as shown in Figure 6.12, such that it did not obstruct the air flow around the right hand side of the equipment case where the main power supply for the WSP unit was located.

A maximum load requirement of 5 g is not tremendously hard to satisfy, particularly if the equipment is essentially located inside the vehicle and so is unlikely to become detached even if the fastenings fail. The fourth cluster of instrumentation was, however, located on the surface of the bogie and was only attached by four M12 bolts. The requirements specified in GM/RT 2100 for attaching equipment to the bogie are significantly more stringent than those for attaching equipment to the vehicle body. The requirements are shown in Table 6.7. Unlike the vehicle body,

the most stringent requirement is that the equipment should remain attached when subjected to vertical accelerations of 20 g.
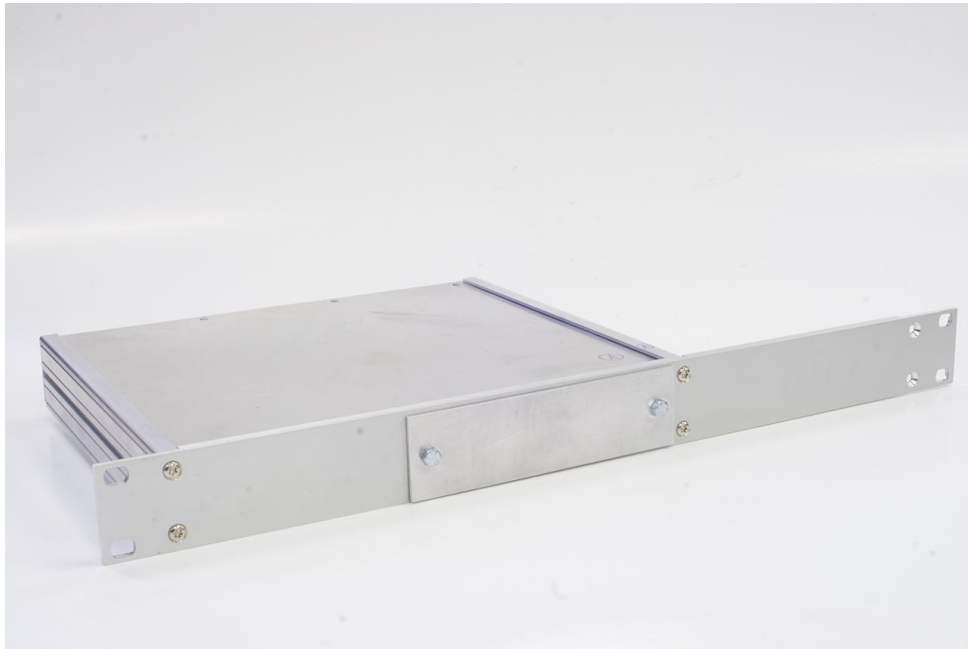
Figure 6.12: Mechanical housing for the second instrumentation cluster

Table 6.7: Mechanical loading requirements for bogie mounted equipment

Axis	Load
Longitudinal	$\pm Mass(kg) \times 5 \times 9.807 ms^{-2}$
Transverse	$\pm Mass(kg) \times 3 \times 9.807 ms^{-2}$
Vertical	$\pm Mass(kg) \times 20 \times 9.807 ms^{-2}$

The calculations to show that the equipment would remain attached to the bogie are not actually much more complex than those for the vehicle body. GM/RT 2100, however, also specifies fatigue life limits of not less than 10^7 cycles with a probability of failure of not more than 2.5% when subjected to the forces shown in Table 6.8. Testing procedures for the GM/RT 2100 requirements are specified in the 1999 edition of BS EN 61373. Lloyds Register Rail were contracted to design the housing for the fourth equipment cluster so that it met all of the appropriate requirements.

Table 6.8: Fatigue loading requirements for bogie mounted equipment

Axis	Load
Longitudinal	$\pm Mass(kg) \times 0.5 \times 9.807ms^{-2}$
Transverse	$\pm Mass(kg) \times 1.5 \times 9.807ms^{-2}$
Vertical	$\pm Mass(kg) \times 10 \times 9.807ms^{-2}$

All equipment being mounted onto a railway vehicle must adhere to flame resistance standards. In this case, the nature of the vehicle's route which takes it through single bore tunnel sections meant that particularly stringent fire standards needed to be applied. BS EN 6853 [95] is the main fire resistance standard that must be adhered to for equipment to be allowed to be attached to a railway vehicle. The standard has various degrees of severity, the most strict of which is known as "category 1a". The instrumentation developed for this project was required to meet this most strict variant of the standard.

The boxes housing the instrumentation were all of metallic construction and so did not pose a fire risk. The standard includes a clause whereby single items enclosed in metal containers and of mass less than 100 g do not need to be considered. Fortunately this clause could be used in relation to the instrumentation itself. A major aspect of all fire standards relates to the spread of fire. For this reason, any connections between different parts of the vehicle are considered particularly closely. In this system, only three types of connection existed: cables, fibre optic cables and conduit. The cables were all selected from Huber and Suhner's Radox and Tenius ranges which meet BS EN 6853 category 1a. The fibre optic cables used did not explicitly meet the standards but did comply with the American UL VW-1 standards which are comparable, if not quite as strict, and the total mass of any given fibre optic cable was less than 100 g. The conduit used needed to be both environmentally and mechanically rigid and so metallic conduit was not deemed suitable. Instead a plastic conduit called Nylofix IRT was used. This conduit

meets the London Underground standard LUL 2-01001-002 [117] which is even more stringent than BS 6853 category 1a.

The final part of the approvals process was concerned with how the modifications to the vehicle might affect the maintenance procedures. According to the modification change approval letter issued by Atkins, the changes required have been deemed "textual" and so a full re-issuing of the vehicle approval body certification is not required. The changes required centre on the fact that the plates forming the base of the equipment cases have now been modified so that they can be removed in sections, thereby allowing access without needing to remove all of the instrumentation. Atkins were particularly concerned that the electrical warning labels associated with these base plates should be duplicated on each of the new plate sections, and that the maintenance documentation should be modified to include maintenance of these new labels. They also indicated a requirement that the maintenance documentation should be amended to include a description of the additional equipment, and a checklist of things to ensure during maintenance, such as the appropriate bolt tightening torques.

6.3.8 Fittings

The instrumentation being described in this case study is by far the most complex system that the University of Birmingham has fitted to a railway vehicle. The distributed nature of the system meant that a number of independent mechanical interfaces had to be developed. As the system is an extension of the work described in the energy monitoring case study, one of the four equipment clusters did not require much mechanical development. The other three, however, needed substantial work. Fortunately the second equipment cluster was housed in a standard 19" rack unit. This dramatically simplified fitting, however some modifications still needed to be made to the equipment case. The third and fourth equipment clusters needed their

mechanical housings designed from scratch but also needed to use the existing fittings of the vehicle. Drawings of these fittings were not available and so a series of measurements and trial fittings were used to develop the interfaces. Even with a number of trial fittings, and one part of the instrumentation having been used before, it was still necessary to make modifications during the final fitting process due to the build tolerances of the Class 508 rolling stock.

6.4   Trials

Unlike the instrumentation described in the previous three case studies, there was no specified trial plan for the duration that the instrumentation was fitted to the vehicle. Instead of a single day of trials, or even multiple days of trials, the instrumentation remained fitted to the vehicle as it operated normally for a period of one year. During this time, the instrumentation recorded to the three SD cards located at either end and on the bogie. Interaction with the equipment was as little as changing these SD cards at the vehicle's routine maintenance intervals of 19-21 days.

The routes operated by Merseyrail are shown in Figure 6.13. The Wirral (green) and Northern (blue) routes are DC electrified and so can be serviced by Class 507 and Class 508 vehicles. Maintenance for these fleets takes place at the Birkenhead North train maintenance depot which is situated on the Wirral line. The vehicle wash facility is located at Kirkdale, on the Northern line. It was therefore guaranteed that the vehicle would operate on both branches of the Merseyrail DC electrified network.

Figure 6.13: Map of the Merseyrail network

In addition to monitoring the standard operation of the vehicle for the long term goal of driving style analysis, diagramming was also used to ensure that the instrumented vehicle passed over particular sections of the network. This was more focussed on the industrially specified outcomes of the work, where Merseyrail could use the instrumentation fitted to the vehicle to monitor the condition of the power supply system at various locations within the network. One particular example of this is Garston. The Garston substation provides power to the network at the southern end of the Norhthern line, between the Cressington and Liverpool South Parkway stations. Merseyrail currently have a particular problem where the traction system on their vehicles is being shut down to protect it from particularly low voltages in the vicinity of the

Garston substation. By sending the instrumented vehicle to this region during peak hours, the worst case line voltages can be recorded. As an extension of this, it has been possible to perform a series of trials where the instrumented vehicle is taken out of passenger service and the real time viewer is used to monitor the instrumentation during a series of artificial trial events. In these cases, particular traffic patterns or vehicle duty cycles can be simulated.

Another effect that was not originally specified but that occurs regularly is the combination of two three car trains to make a six car train. This tends to occur during peak periods and, although not originally part of the trial plan, has already happened to the instrumented vehicle during the trial period.

6.5   Results

6.5.1   A and B ends of the vehicle

During the work described in the energy monitoring case study the instrumentation was fitted for a single day's operation to one end of a vehicle. The assumption was made that, despite independent control systems, the two ends of the vehicle would use approximately the same amount of energy as they were physically connected together and were receiving the same demands from the driver. The work being described in this case study involved the instrumentation of both ends of a similar vehicle. It has, therefore, been possible to use the data obtained from this case study to verify the assumption made in the previous one.

The data recorded by the system was tagged using a time index generated by the on board real time clock calendar chips. The two ends of the system operate independently and have separate clock chips. These are highly accurate but, like any clock, still have to be set in the first instance. While adequate for most purposes, some of the data in the system is being sampled 8192 times

per second, and recorded at 256 Hz. At this level, a more precise alignment is required. The demand signals generated by the driver are simultaneously transmitted to both ends of the vehicle. These are used to align the two datasets.

Once the data from the two ends is aligned, it is possible to extract sections relating to specific journeys for comparison. Figure 6.14 shows all of the data recorded for traction power consumption during the 24 hour period covering the 28^{th} February 2011. As the traction power curves are similar, the data recorded from the B end of the vehicle has been plotted mirrored in the x-axis. The figure suggests that the energy used by the two ends of the vehicle is comparable, although the peak values recorded at the A end are usually slightly greater.
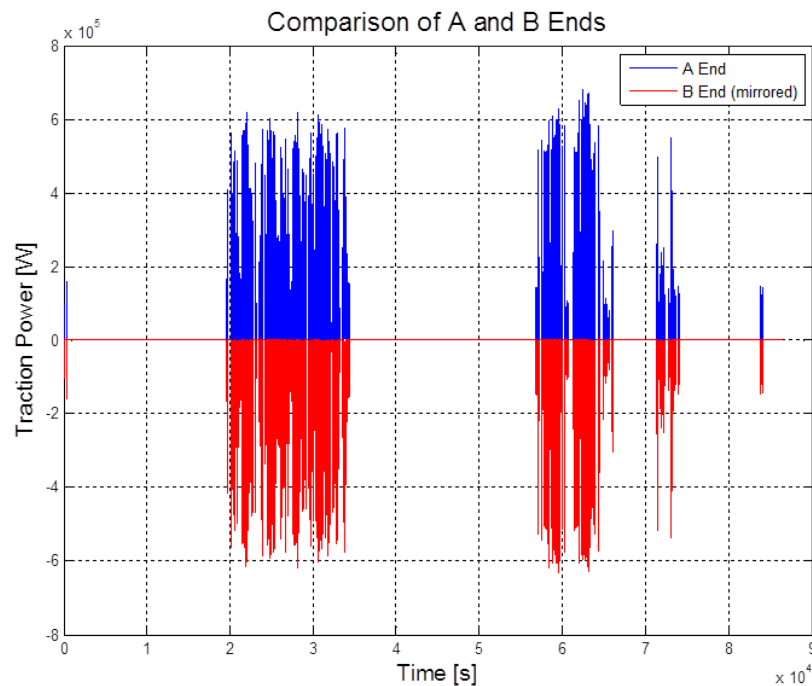
Figure 6.14: Comparison of traction power data taken from A and B ends of the vehicle (B end shown mirrored)

The figure shows the instantaneous traction power over the course of a day. It is possible to calculate the total energy consumed by each end of the vehicle during this day. The values for the total energy consumed by the two ends of the vehicle are shown in Table 6.9.

231

Table 6.9: Energy consumption by each end of the vehicle for the sample day

A End	B End	Difference
1894.5 MJ	1915.5 MJ	38.7 MJ

The difference in energy consumption between the two ends of the vehicle is approximately 1% of the total energy consumed. The transducers themselves have a linear performance guaranteed in the order of 0.7% but are usually better than this. The difference is unlikely to be due to transducer offset as the instrumentation has been calibrated to remove such effects.

As the two ends of the system operate independently, it is only possible to align the data to within the precision that it is sampled. The driver's demand signals, which are used for this alignment, are sampled at 16 Hz. Figure 6.15 shows the power use for both ends of the train for one station-station journey recorded on the 28^{th} February 2011. The point-point difference between the two signals is displayed in the lower part of the figure. This indicates the significance of the alignment of the data. The "difference" recorded is far greater than the actual difference between the signals when perfectly aligned.
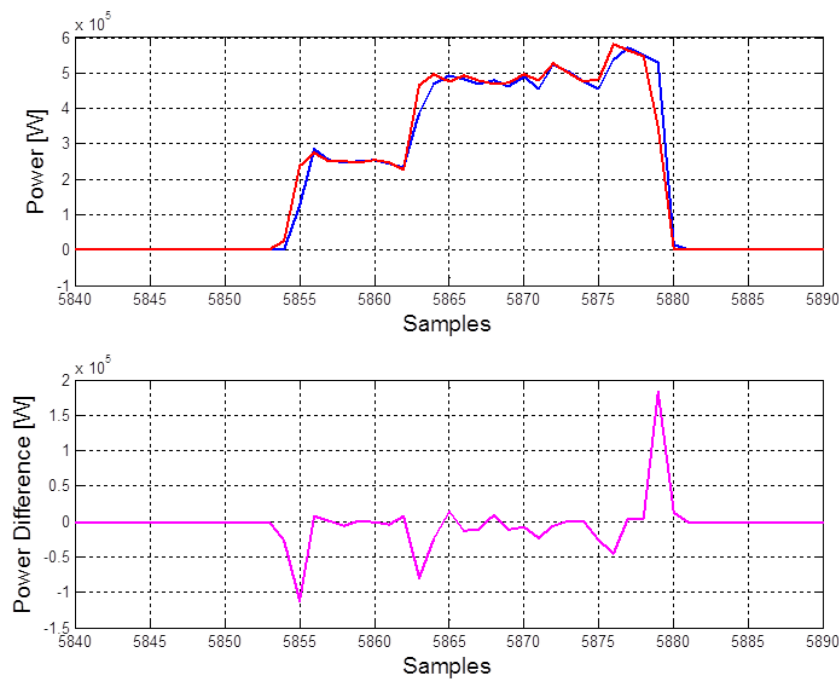
Figure 6.15: The significance of precise alignment on point to point energy metering result comparisons

Another example of the quality of the alignment between the A and B ends of the vehicle is shown in figure Figure 6.16. The figure shows two station-station journeys, first from Liverpool Central to James Street, and then from James Street to Hamilton Square. Again, the B end power consumption has been mirrored in the x-axis for clarity. A scaled and shifted version of the speed signal has also been overlaid for context. The figure shows that the low frequency components of the power use align reasonably closely between the two ends of the vehicle, but that the higher frequency components vary slightly. This is particularly evident during the peak traction period between 58200 and 58250 seconds.

One reason for this slight difference could be which of the two ends of the vehicle is leading. The vehicle is not rigidly coupled and so there would be some flexibility in the connection between the two ends. Given that the vehicle reverses, and each unit is outward facing, it is difficult to define what constitutes "forwards" and what constitutes "reverse". It is possible, however, to

gain some understanding of vehicle direction from the driver's direction requests. Looking at the data from the 28^{th} February 2011, to correspond with the energy data, it is possible to say that, when direction requests were being made, 47% of them were in one direction, and 53% in the other.
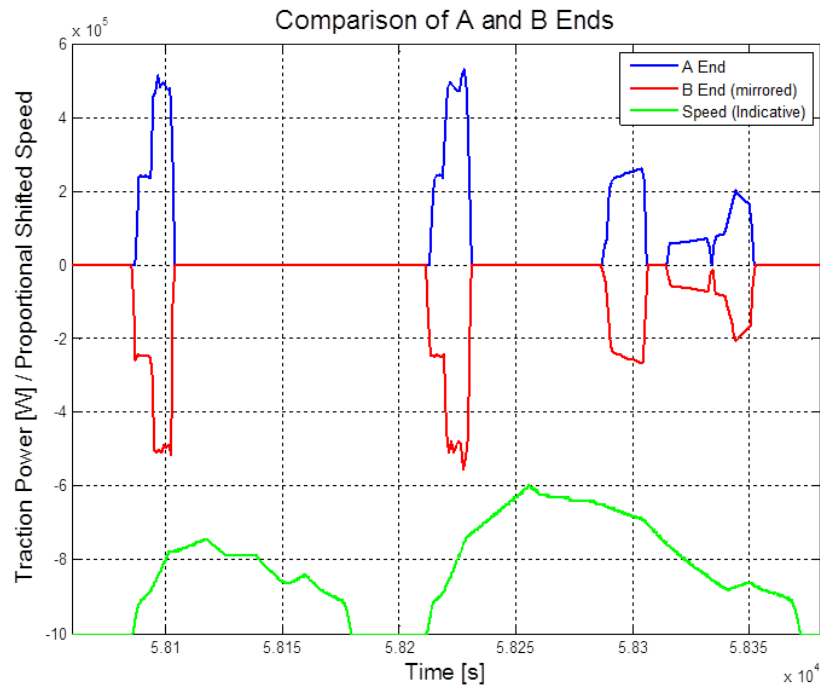
Figure 6.16: Comparison of power consumption for A and B ends of a vehicle in light of journey profile

On the 28^{th} February 2011, the instrumented vehicle was routed on the Wirral line, south of the city. Most of the journeys were between Liverpool central and one of either Cheter, Hooton, or Ellesmere Port. The vehicle, however, started the day at the train maintenance depot in Birkenhead North, and ended it in Kirkdale. These journeys are of differing lengths and so it is likely that this can account for the difference in direction, and thus possibly energy consumption.

6.5.2 3 car and 6 car configurations

One thing that the instrumentation being described in this case study does not measure is the configuration of the train. At peak times, the vehicle may be coupled to another compatible unit to form a 6 car train. In this configuration, the control lines are coupled between the two units so that all four traction systems (eight motored axles) operate as a single unit. The instrumentation does not include a signal to directly identify the configuration of the train.

In practice, because the vehicles can only be joined and split at certain locations for timetabling reasons, the six car operation is a bigger percentage of all journeys than anticipated. Table 6.10 shows the use of different configurations over the Hunts Cross to Southport route for the month of March 2011.

Table 6.10: 508131 running configurations for the Hunts Cross to Southport route in March 2011

Date	Journeys	3 car journeys	6 car journeys
1/3/2011	15	12	3
14/3/2011	11	5	6
15/3/2011	14	8	6
16/3/2011	10	4	6
17/3/2011	13	10	3
19/3/2011	14	14	0
Total	77	53	24

The table shows that approximately 31% of all journeys made by the instrumented vehicle over the sampled route in March 2011 were made while coupled to another, comparable, vehicle. In the most extreme case, the 16^{th} March, 60% of the journeys were made in this configuration.

As mentioned, the data used to populate table Table 6.10 cannot be obtained from the instru-

mentation directly. Instead, using the links that this project has to industry, it was acquired

from outputs from the web interface to the Network Rail / Gemini Applications Ltd GEMINI

database [118]. This is a database that interfaces with the GENIUS train allocation system,

and the Rail And VEhicle Records System (RAVERS) and PAssenger Billing System (PABS)

databases. The GEMINI system has a mainframe component and a web interface. A screen

shot from the web interface is shown in figure Figure 6.17. The final column of data indicates

the presence of a second, coupled, three car unit where applicable. It should be noted that the

times in the GEMINI system are timetable ideals and may not occur in practice.

Figure 6.17: Example output from GEMINI database

With the increased regularity of six car trains present in the data being recorded, it becomes

important to consider ways of automatically identifying three and six car configurations. One

solution for this would be an automatic hook into the GEMINI database, however, a more self

contained solution is desirable from an academic perspective.

Figure 6.18 shows the speed profiles taken from the tacho units on the A end of three vehicles

travelling from Moorfields to Liverpool Lime Street and then on to Liverpool Central. The blue

line on the figure represents a six car train, while the red and green lines are from three car trains

at off-peak and peak times respectively. Particularly noticeable in the first part of the journey, from Moorfields to Liverpool Lime Street, the six car unit has a far smoother acceleration profile and takes longer to get up to the top speed. Braking into the stations is also a smoother process for the six car unit. All of these effects are likely to be due to the increased mass, and thus inertia, of the six car trains.
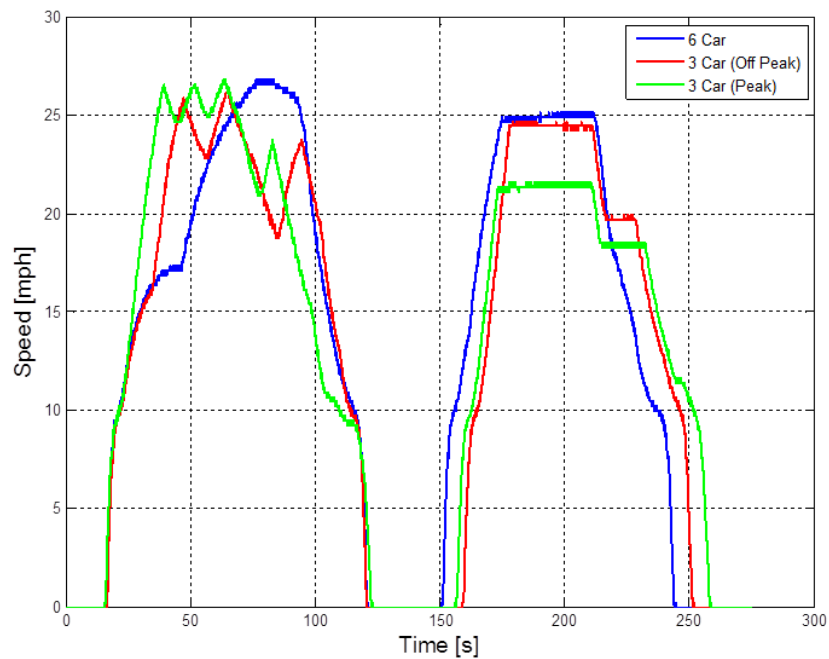
Figure 6.18: Velocity profiles for 6 car, 3 car and 3 car peak time vehicle configurations / journeys

The signal from the AETW system, which is used to give a measure of the pressure in the secondary air suspension and thus the weight of the train, shows that the peak time three car train is heavier than its off-peak counterpart. Following the line of reasoning taken with the six car train, the speed signal should therefore be smoother; however, this is not the case. For a heavier train to have a more variable speed profile, there must be a significant difference in the driver's demands.

Figure 6.19 shows the same three journeys from Moorfields to Liverpool Lime Street and on to Liverpool Central but also includes the driver's handle position and the traction energy consumed by the vehicle. As predicted, the two three car journeys have different driving styles. In the first leg of the journey, only the three car peak time journey uses the full extent of the four possible traction positions and only uses notches three and four. The three car off peak journey steps the demand as the vehicle accelerates and doesn't use the fourth notch. The peak time unit has three points when it comes fully off the power mid-journey compared to two for the off peak unit. This causes the variation in the speed profile but is generally compensated for by the more aggressive notching back onto the power. The overall journey times for the first leg for each of these two trains are comparable. In the second phase of the journey (Liverpool Lime St. to Liverpool Central) both three car units use a single burst of notch four. The off peak unit maintains this for longer to achieve a higher speed and thus requires some additional braking in the deceleration phase. The off peak unit is faster in this leg, but is held at Liverpool Lime Street for longer, hence it is only slightly faster over the whole two-leg journey. The six car journey rarely exceeds the second notch and doesn't engage notch four at any point. This, more conservative demand structure leads to a smoother traction power demand with lower peak values, along with the smoother speed profile described previously.

Despite the more aggressive driving style used in the two three car journeys, the actual energy consumption is greatest in the six car journey, where the power drawn is more continuous. Table 6.11 shows the energy consumption for the three different vehicle / journey types between Moorfields and Liverpool Central.
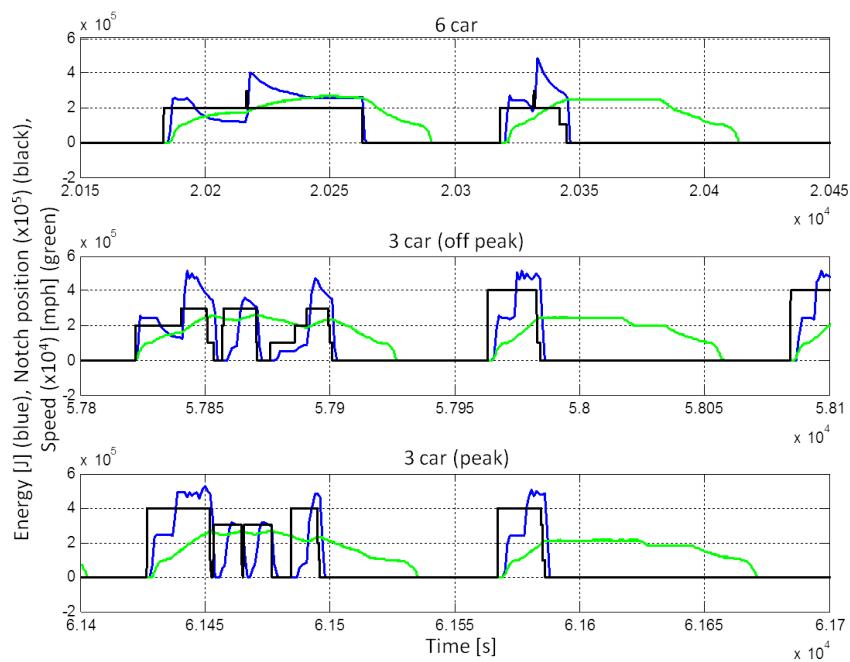
Figure 6.19: Energy consumption and velocity profiles for 6 car, 3 car and 3 car peak time vehicle configurations / journeys including drivers demands

Table 6.11: Energy consumption for 6 car, 3 car and 3 car peak time vehicle configurations / journeys

Configuration	Energy used [MJ]
6 Car	25.7
3 Car Off Peak	23.4
3 Car Peak	22.4

The 77 repeated journeys recorded along the route between Hunts Cross and Southport in March 2011 can be used to further consider any correlation between vehicle configuration, journey time and energy use. Slack in the timetable, however, tends to mask the journey time and so single station-station journeys, or "hops", should be considered. Table 6.12 shows the energy use and journey times for three and six car journeys between Hunts Cross and its neighbour Liverpool South Parkway, and between Southport and its neighbour Birkdale.

Table 6.12: Energy consumption for different vehicle configurations over a single station journey

	Hunts Cross / Liverpool South Parkway		Southport / Birkdale	
Configuration	3	6	3	6
Time [s]	201	213	216	227
Energy [MJ]	7.5	5.9	18.4	17.3

The table suggests that in both cases, the six car journeys took slightly longer than the three car journeys, and in both cases consumed slightly less energy. The differences in energy are proportionately greater than those in journey time. However, neither the difference in journey time or in energy consumption is significant enough to base a classification of the configuration of the train on.

The only reliable markers for the difference between the three and six car configurations appear to be the "smoothness" of the speed profile, and the driving style. Figure 6.20 shows an example of this for a vehicle travelling between Hunts Cross and Liverpool South Parkway. The 6 car configuration, shown in blue, has a continuous demand for power at notch 2 which translates to a smooth acceleration up to approximately 30 mph, followed by a smooth and consistent braking phase. The three car train (green) uses bursts of an equivalent demand level to accelerate to the 30 mph point. This acceleration takes longer and so the driver further increases speed to compensate. This is done with a short request at notch level four. This type of burst of increased speed is only seen in the 3 car unit, which has lower inertia.
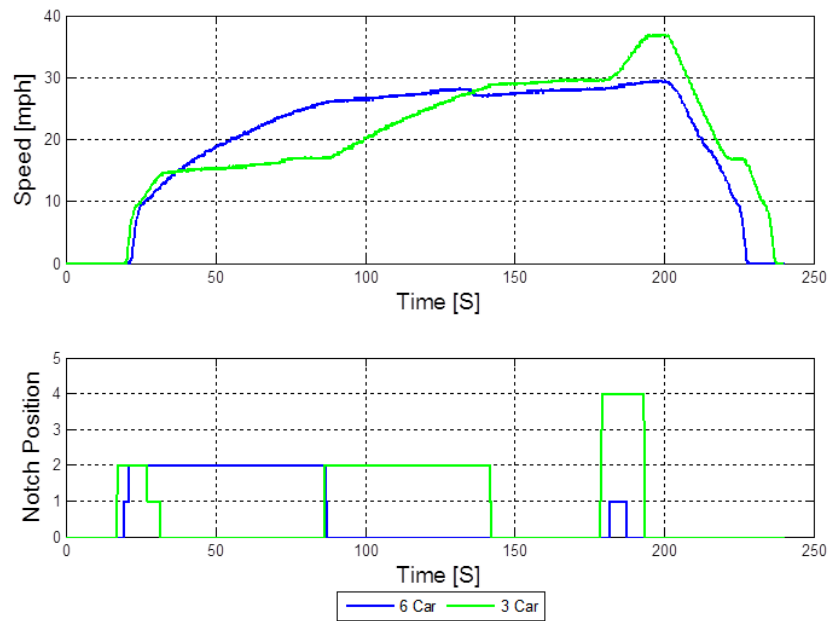
Figure 6.20: Velocity and driver profiles for 3 car and 6 car configurations over a single station journey

It is possible to establish some metric for driving style based on the percentage of a given journey that the driver makes each type of demand for. Table 6.13 shows, for each vehicle configuration, the compiled percentages for all of the data from March 2011 for the 2.1 km single "hop" journey from Hunts Cross to Liverpool South Parkway. The most significant difference is that the six car train spends more time coasting than the three car train, as indicated by the use of notch position zero. The three car train spends an almost equivalent amount of time in notch two.

Table 6.13: Numerical comparison of driving style in 6 and 3 car configurations

| Notch | % Time spent in notch | | % Difference: 3 car to 6 car |
	3 car	6 car	
0	60.0	69.7	+9.7
1	6.7	7.5	+0.8
2	31.3	20.3	-11.0
3	1.4	0.4	-1.0
4	0.7	2.0	+1.3

6.5.3 Driving style

Although the instrumentation being described in this case study directly monitors many of the driver's interactions with the vehicle, a straightforward analysis of these interactions is not the most revealing way to analyse driving style. The true metrics for driving style relate to the speed at which a vehicle completes a journey, and the amount of energy that is consumed in the process of it doing so. The driver's inputs to the system only become relevant once an interesting journey has been identified.

While journey time and energy consumption are widely accepted as the primary metrics for journey quality in a metro system [119], other factors such as jerk rate are also significant and, to some extent, within the driver's control. Even if only the two main metrics are being considered, the significance attached to each is also an important factor in "optimising" the driving style. The investigation into driving style presented here does not go as far as making recommendations based on a full optimisation model. Instead, discussion will focus on the difference that driving style can make to the two main metrics and broad recommendations for

improving both.

The instrumented vehicle made 53 journeys between Hunts Cross and Southport during March 2011 while operating in a 3 car configuration. Although other routes were undertaken, the layout of the Northern Line of the Merseyrail network means that approximately half of these journeys were made in each direction. The Northern and Wirral lines come together at Liverpool Central; this is therefore an obvious location at which to divide the journeys. The journeys have been considered in four sections as shown in Table 6.14, as well as in their full lengths.

Table 6.14: Hunts Cross / Southport journeys considered in four sections

Starting Station	Ending Station
Hunts Cross	Liverpool Central
Liverpool Central	Southport
Southport	Liverpool Central
Liverpool Central	Hunts Cross

It is reasonably straight forward to consider a journey in terms of its duration and the amount of energy that making it required. The raw data being obtained from the instrumentation's transducers is in the time domain and is recorded in a format that inherently preserves timestamps. When the journeys are divided into sections, such as those in Table 6.14, the timestamps are again preserved. These can be used to identify the duration of the journey. Energy use information is simply calculated from the power requirements at each sample time between the start and end timestamps associated with the journey. The driving profile information may sometimes be confused by the driver modifying his or her behaviour part way through a journey based on timetabling or signals. While sometimes difficult to identify, this can provide extra information regarding the consistency of driving style throughout a journey.

Figure 6.21 shows a subset of the data, normalised to the slowest journey and the one with the

highest energy requirement. As one might expect, there is a strong correlation between shorter journeys and larger energy requirements. This is illustrated with journeys 12 and 13. Journey 12 took the longest time recorded for the six stops between Liverpool Central and Hunts cross. While it didn't use the least energy over the route, it was certainly one of the more efficient journeys. Journey 13, in contrast, was one of the faster journeys recorded but required the most energy. One of the more interesting journeys displayed on the figure is journey 3. This journey was both one of the faster journeys and also one of the more efficient in terms of energy use.

The same journeys shown in Figure 6.21 are also represented in Figure 6.22. In this figure, instead of the journey time and energy consumption, the data represented is the percentage of the journey time that the driver spends in each of the four traction notches, and the zero notch position. Journey 12, the slow journey, appears not to use very much of the fourth notch position compared to journeys 13 and 3. Journey 13, the journey with the high energy consumption, has a comparable level of notch 4 use as journey 3 but coasts much less. Journey 3 reduces the time spent in the lower notches to increase the time spent coasting, thereby suggesting that use of notch 4 may be an efficient substitute to using the lower notches.
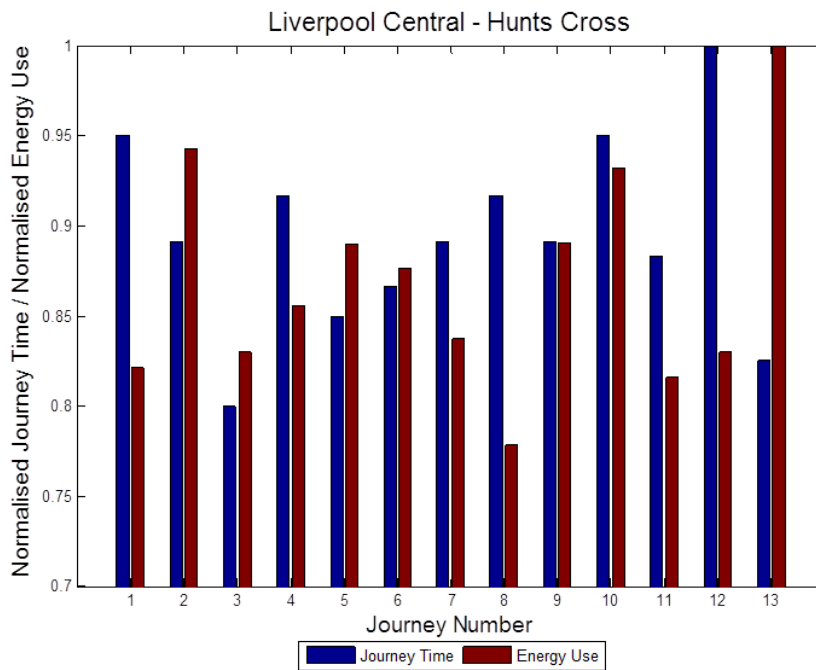
Figure 6.21: Normalised journey time and energy consumption for the Liverpool Central to Hunts Cross route
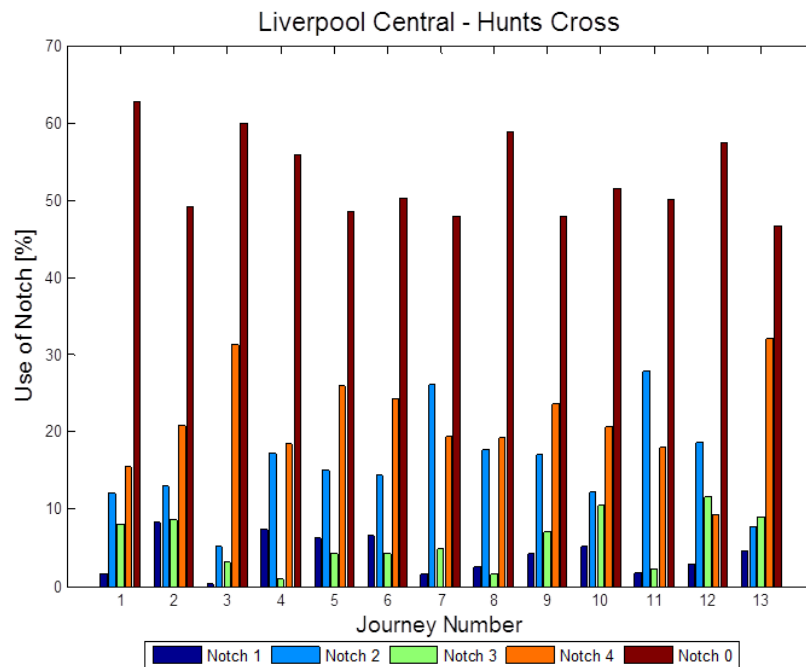
Figure 6.22: Driver's demands for the Liverpool Central to Hunts Cross route

While Figure 6.21 is good for summarising large numbers of journeys for easy comparison, the information contained therein is an average driving style over the six station stops that the route contains. An equivalent figure for the route between Liverpool Central and Southport would be even more obfuscated by the fact that the route contains sixteen station stops. Figure 6.23 shows the actual speed and demand profiles for the three journeys being discussed above. The figure shows that the driver for the slower journey, journey 12, was far more dynamic with their demands. The handle position was modified throughout the vehicle's acceleration curve such that the vehicle was always close in operation to the demand. This should not necessarily be confused with optimising for vehicle efficiency, however, as the vehicle operates more slowly in comparable sections of the journey which leads to a large portion of the efficiency savings.
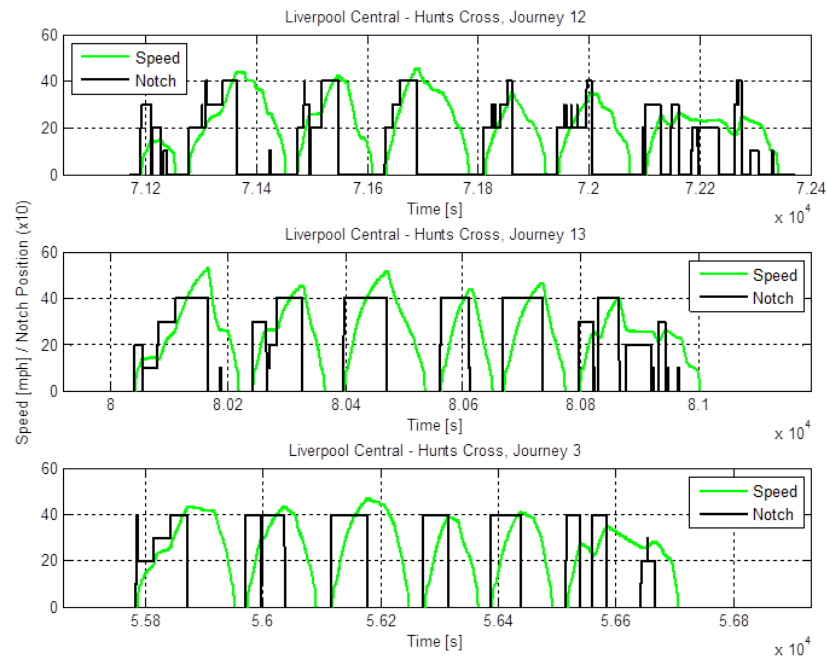
Figure 6.23: Individual velocity and demand profiles for Liverpool Central to Hunts Cross journeys

The figure shows that journeys 13 and 3 have far fewer demand changes. These journeys have comparable durations but vastly different energy requirements. In part the savings come from

the increased willingness to coast shown in journey 3. In almost every case, the driver ceases to demand traction earlier in the acceleration phase. The speed profiles are, however, comparable suggesting less braking is used. This is supported by several rapid speed reductions in journey 13, followed by bursts of traction.

The single journey from St. Michaels to Aigburth, the third sub journey in Figure 6.23, is shown in detail in Figure 6.24. In addition to the driver demand and speed profiles shown previously, the figure also shows the cumulative energy consumption throughout the journey. Journey 12, the manually notched driving style, does have the lowest energy consumption, but takes longer to complete the journey as it has a lower top speed. Journey 3 is shown to be slightly more energy efficient than Journey 13 due to using a lower top speed and more coasting compared to a slightly higher top speed and more braking. This approach can be identified by a driving style that consists mainly of notches 4 and 0, and a smoother speed profile.
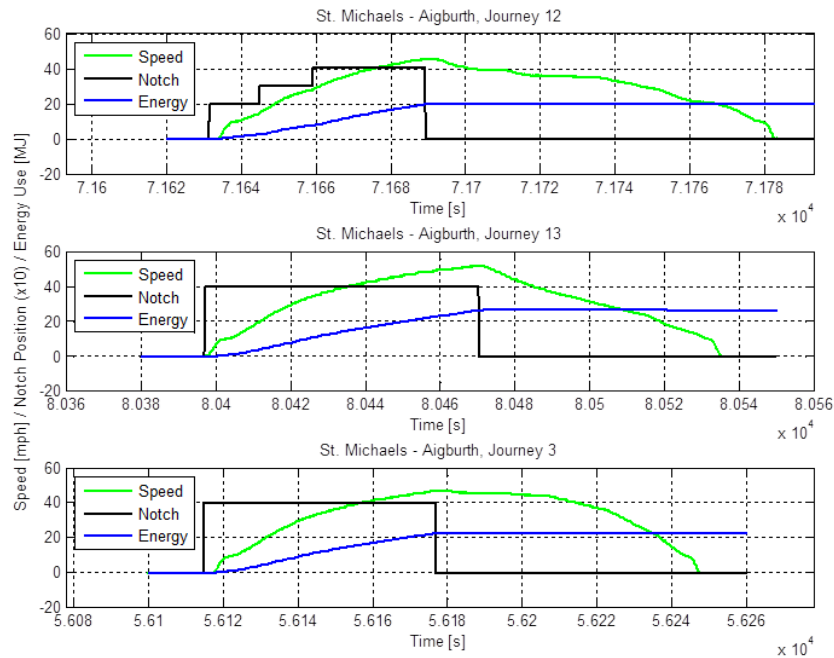
Figure 6.24: Individual velocity, demand and energy proficles for St. Michaels to Aigburth journeys

This trend towards efficiency is found repeated in some of the longer routes inspected. Figure 6.25 shows the journey time vs. energy use balance for journeys from Southport to Liverpool Central, a route consisting of sixteen station stops. The figure shows that none of the journeys take a particularly long or short time compared to the others. This is likely due to the route being long enough, approximately 45 minutes, that it is possible to meet the timetable regardless of a small level of disturbance. While none of the journeys are particularly fast or slow, one of them, journey 8, is far more energy efficient than the others.
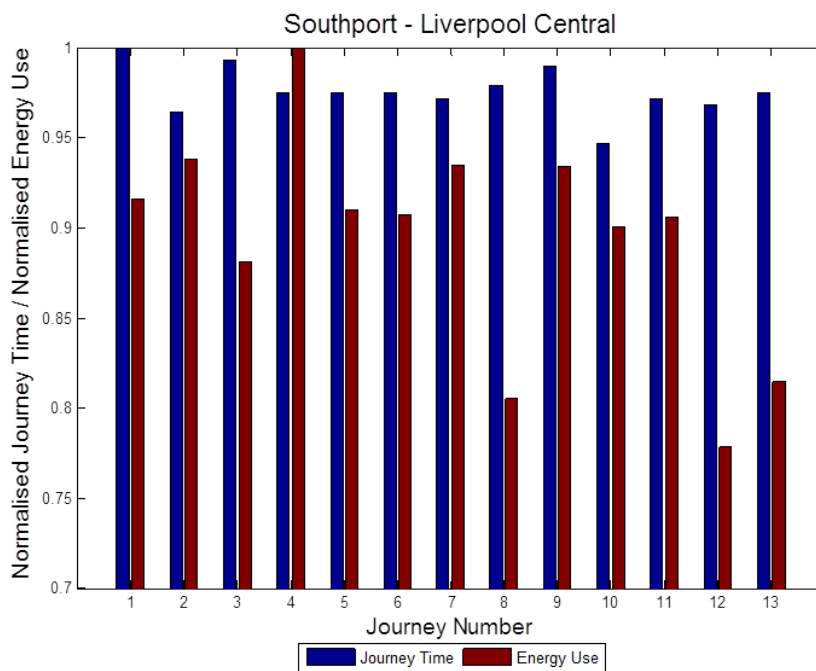
Figure 6.25: Normalised journey time and energy consumption for the Southport to Liverpool Central route

The driver demands and the full speed profile for journey 8 are shown in Figure 6.26. Although the figure shows a large portion of time it is possible to observe that the majority of the time when a traction demand is being made, the driver is requesting notch position 4. The notable exceptions to this are at the terminus stations such as Southport. At these locations, speed restrictions and signalling would limit the driver's options to maintaining lower speeds (e.g.

20 mph around Southport). In these cases, the driver would not want to engage the parallel or field weakening operating modes of the vehicle, and so notch 2 would be adequate. It is noticeable that, in these cases, the driver does not make heavy use of notch 1, moving instead directly to notch 2.
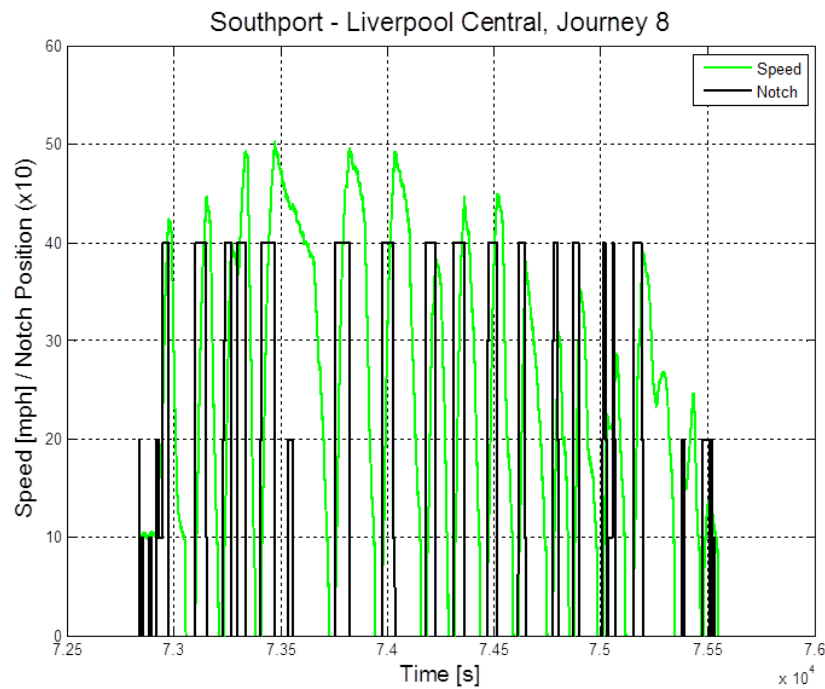
Figure 6.26: Extensive use of notch position four in one example of the journeys on the Southport to Liverpool Central route

There are other occasions when it is not appropriate to use more than the second notch of the traction system. Some stations are sufficiently close that the higher speed modes of operation are not appropriate. One such example can be found in the journeys between Hunts Cross and Liverpool South Parkway. Three example journeys for this route are shown in Figure 6.27. The driver profiles for these three journeys are all similar. Journey 1 used a single burst of notch 2 followed by an extended period of coasting. Journey 10 used a similar, although slightly shorter, burst of notch 2 stepped down to notch one before coasting, and a second shorter (~20 seconds) burst of notch 1. The driving styles used in journeys 1 and 10 yield similar speed profiles and

use similar amounts of energy. Journey 5 began in a similar manner to journey 10 but with even less power demanded early on. This had the effect of a lower initial speed which required the driver to demand more power towards the end of the journey. This extra demand made journey 5 significantly less energy efficient than the other two. From this, it would appear that the most efficient driving strategy for short journeys is to use fewer, larger bursts of power to maximise initial speed and therefore the potential for coasting.
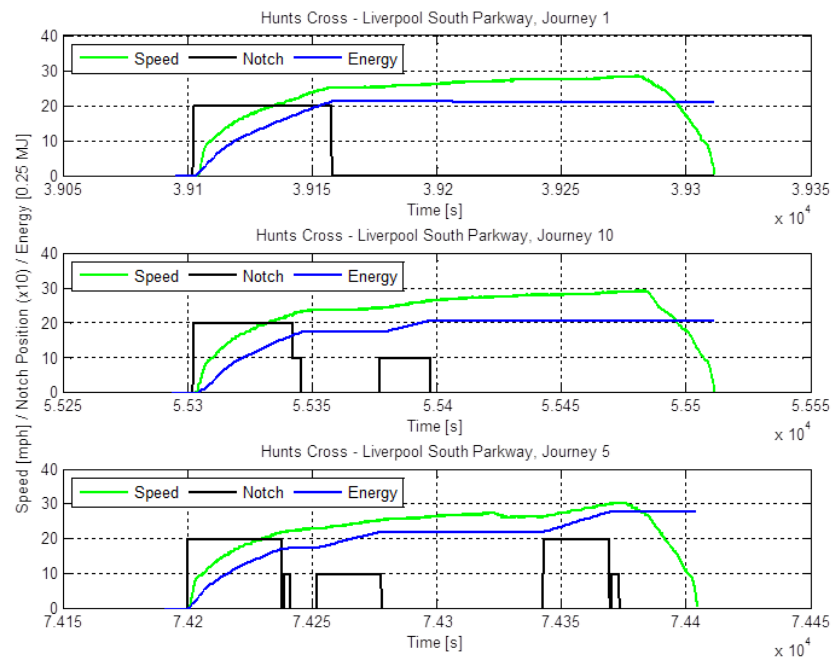
Figure 6.27: Efficient use of lower notch position settings for shorter journeys

6.5.4 Gradient effects

One of the most significant ways that a driver can save energy is to make intelligent use of the environment in which they are driving the vehicle. Previous sections have discussed the use of coasting rather than repeated application of power and braking, but an even more significant technique is to take advantage of the gradient of the track on which the vehicle is moving.

250

There will always be a minimum energy requirement to move a vehicle from one position to another as accelerating (or maintaining the speed of) a vehicle uses energy that is not recovered during the braking phase. The source of that energy is where the savings can be made. When accelerating on a section of track with a positive gradient, the traction system must supply all of the energy. If the gradient is negative, the traction system will be supported by the effects of gravity acting to accelerate the vehicle.

The energy monitoring case study presented in chapter 3 showed results in which different acceleration profiles were considered. One of these results (Figure 3.15) demonstrated a gravity assisted acceleration from a stationary position. The result was classified as being gravity assisted based on the location and the energy consumption compared to the speed profile. Inertial sensors were not available to confirm the gradient.

Work presented in the inertial measurement case study (chapter 5) explains why a vertical sensing accelerometer is not suitable for anything other than very high wavelength vertical alignment monitoring. Instead, a combination of a pitch rate gyroscope and a separate speed signal are used to consider alignment over the standard 35 m and 70 m sections. While this technique is more resilient than the accelerometer based approach over longer wavelengths, there is still a drift component introduced by the gyroscope itself. This drift means that it is only possible to consider vertical irregularity over short time periods.

To consider vertical position over longer time periods, such as those representative of the acceleration phase of a vehicle, it is necessary to produce a more accurate estimate for pitch. A complementary filter can be used to combine multiple estimates which are accurate in particular frequency ranges to give a single, more accurate, estimate for a whole signal [120]. In this case, the complementary filter selected will base one estimate on recordings from a longidudinally sensing accelerometer and the other on gyroscope measurements. The bogie will have some

pitching resonance around 10 Hz, this will be present in the accelerometer signal but should be minimal below 1 Hz. Hence the filter makes use of the accelerometer derived signals at frequencies beteween DC and 1 Hz. The gyrocope solution has little drift and so is good down to very low frequencies but the gyroscopes are not accurate at DC, hence the filter will use the gyroscopes for signals above 0.01 Hz. It is appreciated that this brings the 10 Hz resonance back into the calculation, but the gyrocopes are better suited to this due to their processing chain as described in chapter 5. The author wishes to acknowledge Dr Paul Weston for his support with the development of the this processing chain.

The inertial measurement unit fitted in this case study has two mechanisms for estimating pitch. The first relies on the fact that when an accelerometer is not perfectly aligned to its operational axis it will show some component of gravity and that the low frequency component of longitudinal acceleration can be obtained by differentiating the speed signal:

$$\ddot{x} = g\sin(\phi) + \frac{dv}{dt} + \ddot{x} \text{ vibration at} > 1 \text{ Hz} \tag{6.1}$$

Hence, for low frequencies (<1 Hz):

$$\hat{\phi}_1 = \frac{\ddot{x} - \frac{dv}{dt}}{g} \tag{6.2}$$

The second estimate is based on the direct measurement of pitch. The accuracy of this measurement, however, relies on the gyroscope pitching perfectly around the axis of lateral movement. Some railway track, notably in curves, is canted to allow vehicles to travel faster throughout the curve without succumbing to centripetal force. This cant effect means that the vehicle is rotated around the axis of longitudinal movement, i.e. it has rolled. Taken to the extreme, if the vehicle rolled onto its side the gradient would appear in the yaw rate measurement rather

than the pitch.

Hence, for higher frequencies (>0.01 Hz):

$$\hat{\phi}_2 = \int \left(\dot{\phi} \cos(\hat{\theta}) - \dot{\psi} \sin(\hat{\theta}) \right) dt \tag{6.3}$$

These two estimates are combined using a complementary filter to give an improved estimate for pitch which then allows an estimate of height to be generated. This is shown in Figure 6.28.
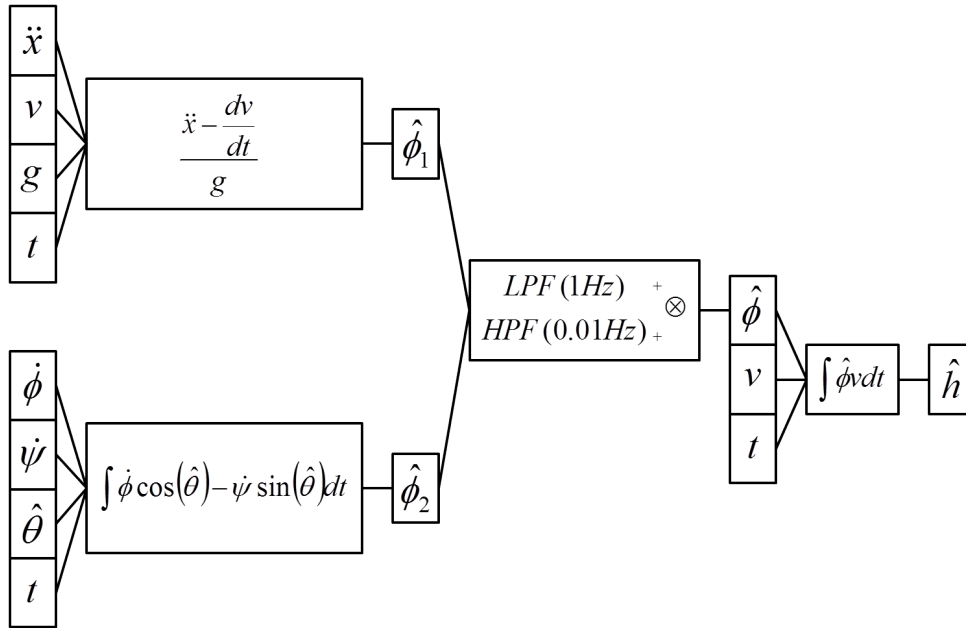
Figure 6.28: Processing chain for estimation of height

To maximise the accuracy of the second estimate for pitch, it is necessary to have the best possible estimate for roll. It is possible to use the complimentary signal technique to produce a multi-part estimate for roll based on the lateral accelerometer and, centripetal acceleration and the roll rate gyroscope. This is shown in Figure 6.29.
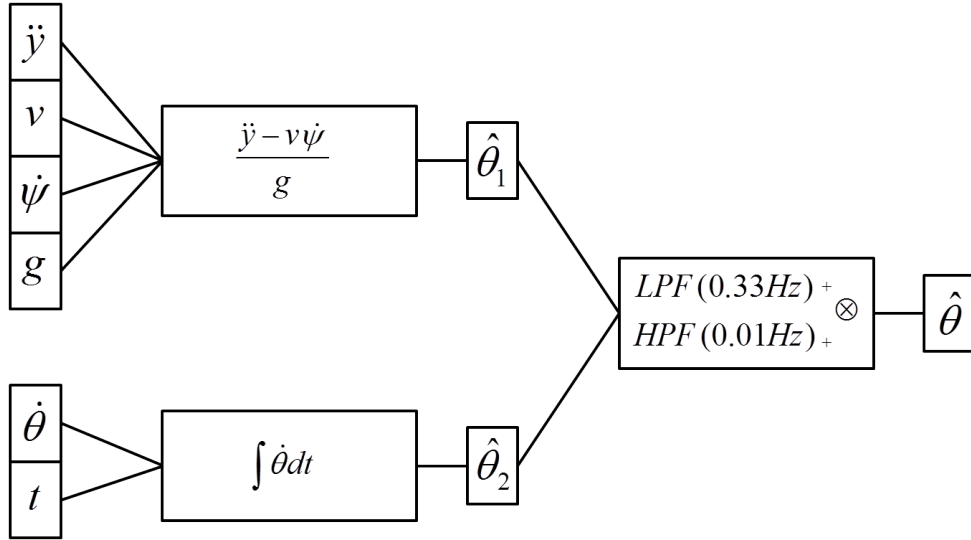
Figure 6.29: Processing chain for estimation of roll

Using these, highly accurate, estimates for roll and subsequently pitch, it is possible to use the data from the inertial measurement unit fitted to the bogie of the instrumented vehicle to contextualise the energy and speed data also recorded. This can then be used to support the development of driving style recommendations based on environment as well as vehicle perfromance. Although attached to the same vehicle, the data storage requirements and system robustness issues mean that the data is recorded in two separate systems. In order to combine the data, the two data repositories must first be combined. The two sets of data are aligned using the GPS and tacho signals which are relayed using a fibre optic interface and are recorded in both subsystems.

Applying the processing described above to the inertial measurement data and using other data from the energy monitoring systems installed on the same vehicle, it is possible to produce graphs such as the one shown in Figure 6.30. The figure was produced using 30 minutes of data recorded on the 28^{th} February 2011, as the vehicle passed through the Mersey tunnel and the Liverpool loop to re-enter service for the evening rush hour trip to Hooton.
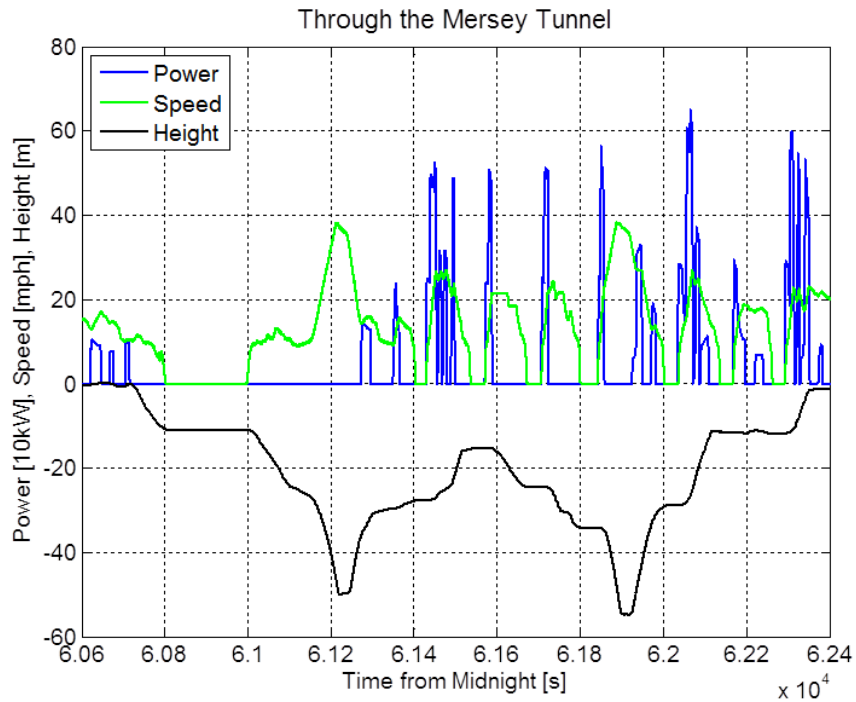
Figure 6.30: Power consumption and speed in relation to vertical height as the vehicle passes under the river Mersey

The figure shows the vehicle power consumption and speed along with the estimation of height derived from the inertial measurement unit. Almost all of the data shown in the figure was recorded in a tunnel but the section where that tunnel passes under the river Mersey is clearly visible; once on the way into Liverpool and once on the way out again. The slight discrepancy between the maximum tunnel depths is due to cumulative error from the transducers (drift) over the duration of the figure. This is also visible with Birkenhead central station being shown as being underground.

During the course of the figure, the vehicle passes through several stations. These are indicated both in the tacho signal and as periods spent at constant height. The vehicle does not stop at all of the stations as it is re-entering service for the evening rush hour. The approximate indexes of the stations are given in Table 6.15.

Table 6.15: Time indexes of station passage

Time from Midnight [s]	Station or point of interest
60800	Signal stop to enter the loop
61100	Hamilton Square (no stop)
61220	Bottom of tunnel
61320	James St. (no stop)
61420	Moorfields
61530	Lime St.
61630	Liverpool Central
61800	James St.
61900	Bottom of tunnel
62000	Hamilton Square
62140	Birkenhead Central
61260	Green Park

The vehicle is forced to stop for 200 seconds as it enters the loop section. This is thought to be signalling to bring it into rush hour operation. The signal is on a sloped section of track with a downwards gradient of approximately 1:42. When the vehicle is released from the signal, the brakes are disengaged and it initially accelerates to approximately 10 mph under coasting. No power is applied. The brakes are used to regulate the speed until the vehicle has passed through Hamilton Square station (without stopping) at which point it is allowed to coast freely. The vehicle continues to accelerate towards the bottom of the tunnel and then decelerates as it climbs out the other side. At this point power is applied to maintain speed through James Street station, where it also does not stop, and again to bring it into Moorfields station.

The counterpart journey, where the vehicle exits the loop section, is also shown. In this case the

vehicle passes through a number of station stops on its way to the main tunnel section. Although Liverpool Central and James Street stations are situated on sloped sections, the comparatively minor gradient and the loaded mass of the vehicle require the driver to apply some power to overcome the vehicle's inertia and to start it moving. Once in motion the power is generally disengaged and the vehicle allowed to accelerate by gravity assisted coasting.

The processing used to generate Figure 6.30 uses complementary filtering and attempts to compensate errors in one sensor using measurements from other ones. This is technically using one data stream to augment the processing of another rather than contextualise it, but is still a valuable integration of (albeit less disparate) signals. This technique is still, however, not completely accurate over such long timescales as the 30 minutes shown in the figure. This is highlighted in the apparent discrepancy between the low point of the tunnel shown on each of the two recorded passes. The processing chain described above is primarily used to generate an accurate estimate for pitch. This is then combined with the vehicle speed and integrated to give the estimate for height. As with the other signals in the system, this integration will amplify any low frequency error. An alternative way to look at the figure, without this final integration, is to consider the power use of the vehicle in relation to the gradient of the track on which it is travelling.

Figure 6.31 shows the pitch of the bogie over the same 30 minute journey. While the bogie remains in contact with the track, this is representative of the gradient. The signal has been offset by -50 mrad to allow it to be plotted on the same figure; the dashed red line represents the horizontal.
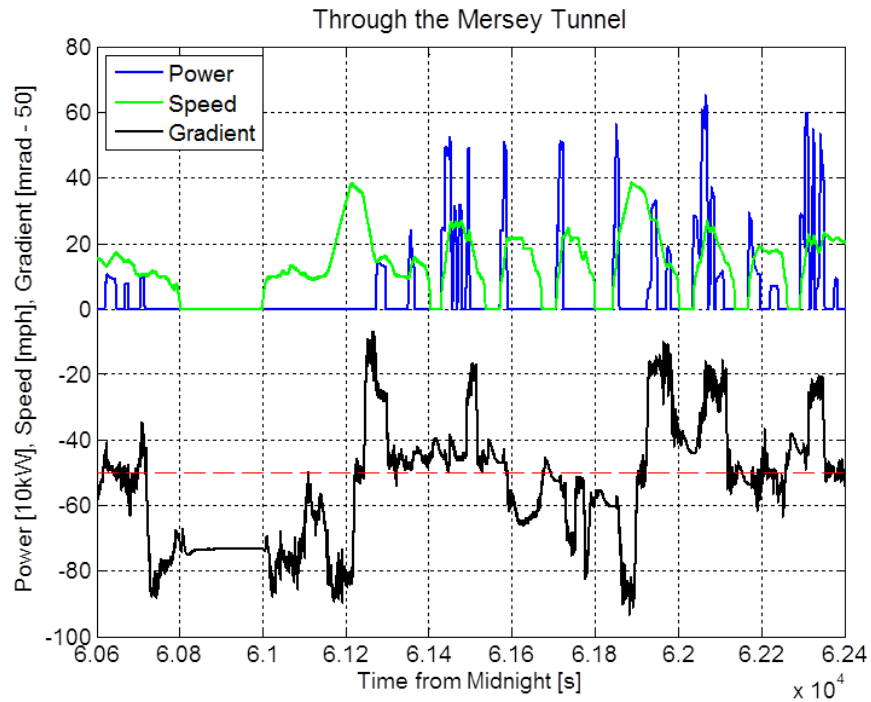
Figure 6.31: Power consumption and speed in relation to gradient as the vehicle passes under the river Mersey

Unsurprisingly, the most severe gradients are found in the deepest parts of the tunnel under the River Mersey. The values estimated for pitch are consistent both before and after the loop section suggesting that the inconsistencies in the estimate of height are introduced by the integration stage. The two most significant gradients recorded are approximately 37 mrad and 33 mrad. In each case the 37 mrad occurs on the side of the tunnel associated with the loop section under the centre of Liverpool, and the 33 mrad on the Wirral side of the river. These values correspond to the 1:27 and 1:30 gradients recorded for these sections of track in Network Rail's 5 mile diagrams for that part of the network. This is an extreme example, designed to convey the concept, but during the steeper decent sections the vehicle can be seen to accelerate without consuming any significant amounts of energy. It is also seen to ascend with a lesser or greater energy requirement depending on the speed carried into that manoeuvre. This is the kind of interaction between complementary signals that enhances understanding of the key data

streams recorded by the vehicle.

The use of the complementary filter will improve the estimate of pitch, and thus height. Ultimately, however, the data that is being used in this processing still comes from transducers which suffer from drift. This means that over time small errors will accumulate such that the pitch and thus height results will appear to vary. If this analysis was being performed to build a gradient map of the network, this would be significant. However, in this case, the analysis is being performed to give context to other measurements obtained at the same time as part of an integrated dataset. In particular, the analysis is being used to contextualise energy use. Providing the cumulative errors in the pitch / height are therefore not significant over the duration of the acceleration / braking phase being considered the process can be deemed acceptably stable.

One way to estimate this error is to repeatedly traverse the same section of track and compare the results in relation to the time between the measurements. In the last few pages, energy use and speed were plotted against estimated height and gradient for a 30 minute return journey under the river Mersey. The height figure (Figure 6.30) shows a total discrepancy for the 30 minutes of less than 1 m, however at the two passes to to lowest point the discrepancy appears greater (as much as 7 m). As mentioned above, it is worth noting that the height estimate is the final stage in the processing chain and includes an integration of pitch. Considering the pitch figure (Figure 6.31), which is arguably more important as it gives the gradient for comparison with energy, the key reference points are the gradients of the main tunnel sections. These differ from each other and from the design specification by less than 1 mRadian over the 12 minutes between the two steepest sections. The length of the longest continuous gradient section in these figures is 1 minute. It is therefore considered acceptable to use this approach to contextualise individual acceleration / braking phases.

6.5.5 Traction system

The control system of a Class 508 is essentially a large digital logic circuit operating on 110 V. The driver's demands are input to the logic system and the requests are interpreted by the cam shaft unit which controls the distribution of power to the rest of the traction system.

External to the work being undertaken here, the University of Birmingham has a number of vehicle simulators which can be used for, among other things, power system dimensioning [121]. These simulators take as their inputs a vehicle model as well as the specifications of the infrastructure on which the vehicle simulation is to be undertaken. Development of these vehicle models requires a thorough understanding of the traction system of the vehicle.

During the energy monitoring case study, a detailed inspection of the traction current waveform was undertaken. At this time, it was speculated that the notches visible in the traction current would align with the operation of the camshaft and the activation of different components of the traction system. The instrumentation being described in this case study includes components to monitor the operation of the camshaft. The magenta components in Figure 6.32 represent the camshaft shift signal which is used to trigger a cam shaft state advancement. The shift signals appear to align with the notching in the traction current waveform.

This information can be used to assist with the development of vehicle models, and also suggests that future instrumentation systems could be simplified by inferring the camshaft signals from the traction current waveforms.
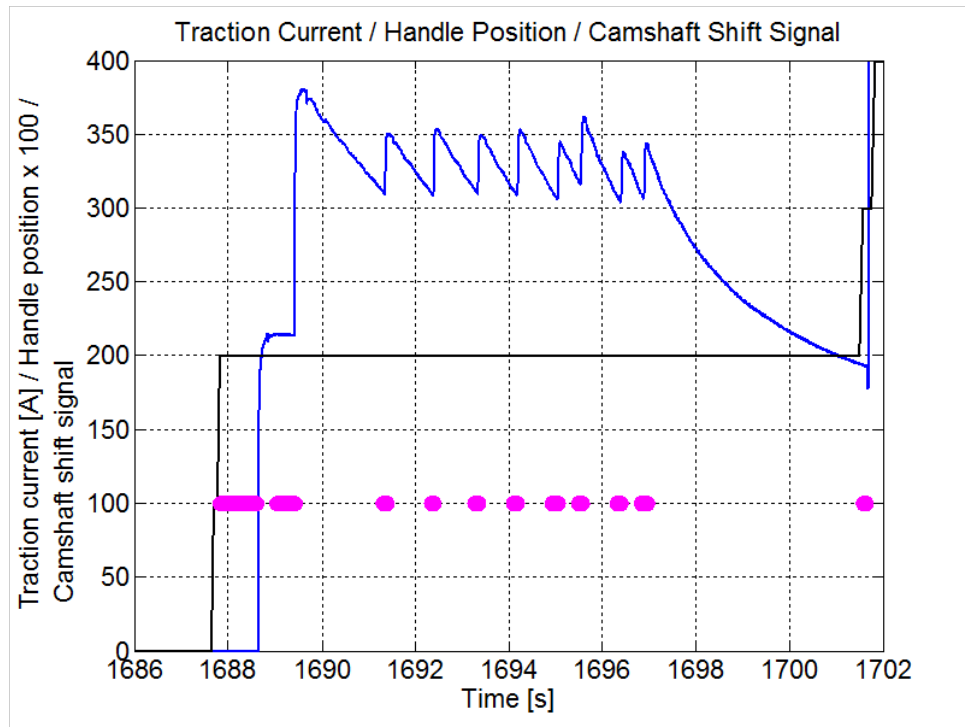
Figure 6.32: Traction current notches aligned with camshaft operation

6.5.6 Auxiliary systems

In addition to the instrumentation installed on the vehicle to monitor the traction system, transducers are used to monitor the auxiliary current drawn. While this is usually small in comparison to the traction current, it is drawn almost continuously while the vehicle is switched on. The auxiliary system is used to power several subcomponents of the vehicle, and so the level of current drawn varies. To some extent this is under the driver's control and may be highly seasonal (e.g. the heating systems).

During the analysis of the journeys between Southport and Hunts Cross, an occasion was observed where the vehicle did not move for a period of almost five hours. During this time, the vehicle was not switched off. Figure 6.33 shows the auxiliary current and the energy consumed by the auxiliary systems of the vehicle during a portion of this time. The current waveform

shows short periods of minor current consumption, along with several more significant events. From calibration testing, the smaller periods could correspond to the driver's cab heater, and the larger events to the main body heaters. The most significant event, around 45000 seconds, could possibly be a combination of the body and trailer heaters. This assumes thermostatic control of the heating systems, and that the heating systems are active. Alternatively, the smaller peaks could correspond to the compressor on the vehicle being activated.
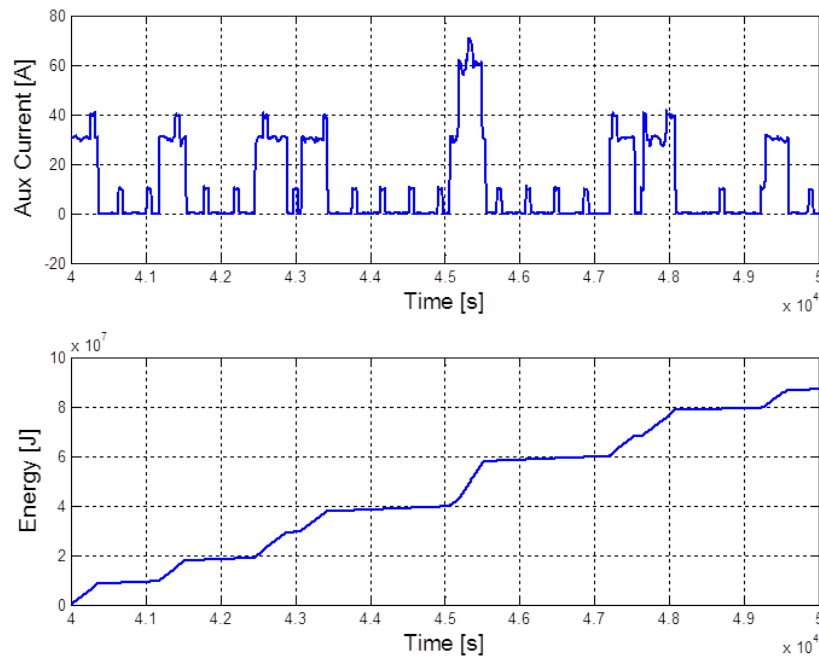
Figure 6.33: Auxilliary current and total energy consumption for a stationary vehicle

The total energy consumption for the period is approximately 90 MJ. This energy could have been saved had the vehicle been switched off. Anecdotal evidence from Merseyrail personnel suggests that vehicles being left switched on is not uncommon, even overnight.

6.6 Conclusions

In this chapter, the author has described an instrumentation framework based on multiple, low cost, simple processors being used to service multiple types of transducers. These "nodes" are used together in a network whereby signals obtained in different locations, at different resolutions, and at different rates can be combined together using a standard interface. Another series of nodes are then used to extract the data presented by all of the various acquisition components so that it can be visualised or stored for further processing. This further processing may involve the interpretation of one signal in light of another.

In many cases, the signals recorded will not in themselves be novel; instead, the effectiveness of this technique arises from the combinations of signals that can be considered based on the richer dataset that a more extensive instrumentation system delivers. The specification and development of an instrumentation system based on this extended framework has been described.

This, final, case study has presented an investigation into the merits of increasing the complexity of instrumentation systems in order to better understand the key signals being recorded as part of a dataset. This is most obvious in the range of sensors used, but also materialises in the sampling and data rates, and the processing and integration techniques applied to the recorded signals. In the case study, the author has described some of the complexities of developing such complex instrumentation systems and the incompatibility between these complex systems and standard acquisition / data recording pair setups. The case study describes the development of a modular instrumentation framework based on a distributed array of embedded processors either acting as acquisition or recording nodes.

The nodal architecture of the instrumentation framework means that it can be used in a distributed fashion over an object the size of a railway vehicle with reduced installation effort and

increased reliability. This has been demonstrated in this case study where an in-service monitoring system has been fitted for an extended period of time. The system developed extends the energy monitoring work presented in the first case study by (among other things) combining it with the inertial systems described in the third case study. The result is an extremely highly instrumented railway vehicle capable of acting as an instrument of measurement for both the electrical and geometrical aspects of the infrastructure as well as a gauge for vehicle energy use, vehicle performance and driving style. The instrumentation systems are capable of operating unattended for extended periods and the nodal framework structure has demonstrated excellent system reliability and fault resilience. The system was designed to be minimally invasive and to operate on in-service passenger vehicles. The significant complexity of the approvals process for putting equipment into passenger service has also been discussed.

The instrumentation framework described in this case study consists of a distributed system of microprocessors, each servicing a range of transducers or other inputs, and each presenting their outputs via a single data bus. This was necessary due to the distributed nature of the inputs being recorded from around the vehicle. The amount of data generated by the system is far in excess of anything described in any of the previous case studies and so consideration has also been given to the data compression format used and the associated storage media.

The concept of contextualising data streams using other data recorded within the same instrumentation framework has been demonstrated, as has the integration of pseudo independent datasets linked by common signals. Applying this context to key data has meant that analysis of it has been able to verify suggestions made in the earlier case studies, particularly the energy monitoring work from which this project evolved. The system has shown that the assumptions made regarding the comparable energy consumption of the two ends of the vehicle are valid. The key factor has been shown to be the precision of the alignment which is achieved using a signal

sampled at only 16 Hz. The system has also validated the suggestion that the notching shown in the traction current waveforms is related to the cam shaft switching its mode of operation. This, along with other data about the operation of the cam shaft can be used to develop vehicle models and so expand the simulators being developed within the University.

The operational parameters of the vehicle have also been considered in the context of being configured as a three or a six car unit. The significance of the six car configuration was underestimated in the original design of the instrumentation and so additional work has been undertaken in the analysis phase to try to identify the configuration from the recorded signals. In general terms, the performance of the traction system is reasonably consistent with only slightly more power demanded in the six car configuration. The demands made of the vehicle by the driver seem to differ between the two configurations with more consistent, lower level, demands being made of the six car configuration. This is reflected in the energy consumption and the speed profile. The maximum frequency of step variation shown in the speed profile is probably the clearest indicator of the configuration of the train.

Courtesy of the long term in-service nature of the trials, the vehicle has been recorded at all stations on both lines of the network under different passenger loading conditions. It is, however, difficult to consider the driving approach when considering different routes and so particular journey sections have been considered in depth. The return route between Hunts Cross and Southport has been divided into four sections for consideration. Analysis has shown that the most energy efficient driving style can be to request only slightly more from the unit than it is delivering at the time. This solution, however, can go drastically wrong and if poorly implemented can cost vastly more in energy than any other style. The preferred solution appears to be to use the highest notch appropriate to the journey and allow the vehicle to manage the traction system. On shorter journeys, with lower top speeds, this may be notch two; on longer

routes, the use of notch four is preferable. For best efficiency, the driver should also be bold with their use of demand during early traction phases, thereby allowing coasting to occur later. Increased initial traction followed by coasting appears to be preferable to reduced initial traction supplemented with traction later in the journey. The use of coasting rather than a combination of braking and subsequent traction is also of significant importance.

The use of the instrumentation framework to integrate multiple types of instrumentation systems into a single dataset has allowed inferences to be made during the processing of one signal based on information obtained from another. A good example of this is the interpretation of the energy data in light of information about vehicle speed and track gradient obtained from the tacho and inertial measurement systems respectively. This has allowed recommendations to be made in the field of driving style. For example intelligent traction, or using a combination of traction and coasting as dictated by the gradient profile, has been shown to yield huge benefits in terms of minimising energy requirements. In particular, the tunnel under the river Mersey has been considered as an area where entire journey sections can occur without the use of the traction system. This was selected as a clear example as the inertial processing demonstrates significant gradients in this region, however, it would be equally applicable in other areas of the network.

A more detailed study into driving patterns particular to locations is required for the train operating company to maximise their efficiency. An understanding of the company's position on the trade off between energy efficiency and journey time would be needed before this could be undertaken. At present, there would be no infrastructure constraints on such an optimisation because the train operating company is not responsible for the infrastructure. As a closed system, however, there is some suggestion that the network may be used as a test case for vertical integration. This, combined with the age, loading and capacity of the power components of the infrastructure may contribute another requirement to any optimisation.

An example of the energy consumption of an inactive vehicle has also been presented. This has been based on data recorded when the vehicle was left unused but not powered off. This data was only recorded because the instrumentation systems devised were so closely linked to vehicle systems and operated whenever the vehicle was powered. How necessary such energy use is currently unknown, however, it would appear that minimising such "losses" would be an obvious step towards minimisation of the energy consumption of the railway.

Due to the unforeseen and extended period required to obtain the relevant authorisations to install the instrumentation on an in-service vehicle, the trials which are scheduled to last one year have been dramatically delayed. The benefits of this kind of integrated instrumentation system are, however, readily visible. Clearly there is still much to be obtained from the instrumentation and statistical approaches, such as the average traction demand over a journey, will become more powerful as the dataset is expanded. Current processing approaches have been acceptable while the dataset has been comparably small, and processing has been being developed, however this will not prove adequate as the volume of data increases. The intention is to develop a database to hold the data and provided tools to extract and analyse various subsets thereof. This kind of ongoing development work is common in academic institutions where the dataset obtained as part of one project may be used in several others. In this case, the inertial measurement data recorded on the bogie is already being used in a UK government funded research project into track stiffness around transitions. In some ways, this extends the inertial measurement work described in the third case study.

In addition to developing a repository for the data, the technologies used in this project will likely be developed further in future instrumentation work. The next logical step would be to upgrade the core processor used repeatedly in the distributed network to allow some of the processing now being conducted at the university to occur on the vehicle. This approach would,

however, potentially require nodes of one type to have access to the data from others. While this is not technically a problem, it would not comply with the framework policy of each type of node being completely independent to allow modularity. A more appropriate solution, which would still operate in the same framework, would be to add processing nodes. These would read the data from whatever nodes were appropriate, perform the processing and then either re-output the data onto the bus to be recorded, or output a flag using a dedicated code in the schema.

The flag based event output solution is generally favoured by industry where vast, unprocessed datasets are seen as inferior to key event markers which can be responded to. While this is supported by the instrumentation framework proposed, there is an existing wealth of railway data already available. From an academic perspective, enriched datasets allow the development of further processing beyond the original scope of the projects that generate them, however recording full datasets in all cases is not ultimately practical. It is likely that in the future some kind of compromise solution will develop where on-board processing nodes are integrated into the framework and multiple output nodes will be used to record both the key event based outputs and (in the shorter term) the full data range available. The University of Birmingham already has a project of this nature under way. The work extends the conductor shoe monitoring work described in the second case study and involves the detection and reporting of defects in the third rail network.

CHAPTER 7

CONCLUSIONS

In this thesis the author has considered the process of extracting information from the physical environment in order to make data driven and informed decisions about the operation of the railway. At present, some decisions are being based on data, however, much of this data is considered out of context. This thesis has shown that by expanding instrumentation systems, and integrating multiple previously disparate data streams, the key data can be contextualised such that better quality decisions, more relevant to the railway system as a whole, can be made.

7.1 Instrumentation Frameworks

This thesis has demonstrated that with modern instrumentation technology it is possible to devise larger, more complex, instrumentation systems that can be distributed over an asset or system (such as a railway vehicle) in order to provide inherent integration of the different data streams required. An example instrumentation framework based on a common backbone and node structure has been presented. This framework delivers the capacity required to interrogate increased numbers of transducers, it allows these transducers to be distributed over a larger asset, it provides the resilience required to operate in the harsh railway environment, and it's modular

nature means that it is adaptable to different scenarios while effectively reducing the number of single points of failure. Each node is developed with a particular role, the failure of which should not disrupt the operation of the remaining components of the instrumentation system. Should the system be re-tasked, acquisition of different signals would simply be a case of changing the acquisition nodes. Should different outputs or real-time processing be required, alternative output nodes could be implemented. Hence the instrumentation framework presented provides the capability to obtain and integrate these different data, while still remaining adaptable for use in different scenarios.

Analysis techniques making use of data from multiple sources distributed over a vehicle have also been shown to provide context to key measurements. In some cases this data has been obtained independently and combined during post processing, but in the more closely integrated systems it has been recorded and integrated within the instrumentation framework. The different case studies presented in this thesis have highlighted the assumptions and caveats that must be made when processing an isolated dataset compared to the processing of an expanded dataset.

Throughout this thesis the concept of instrumentation of railway vehicles and the use of the data obtained to inform decisions within the railway has been considered through a series of four case studies. The first three of these referred to short term data acquisition exercises with highly specific data requirements. All three systems make use of a popular and moderately simple instrumentation architecture consisting of a small number of transducers, possibly with some local management systems, being directly serviced by a single acquisition unit. Different transducers are used in each of the three cases. Their inclusion in the thesis demonstrates some evolution within this simple architecture, but more significantly each case study highlights a limitation of the architecture or instrumentation system used.

The first case study focused on energy monitoring. The range of transducers used was very

limited but the case study shows that by increasing the rate at which data is recorded beyond that mandated it is possible to enrich a dataset such that the level of analysis and interpretation possible is increased. Consideration of the data has indicated things about the operation of the infrastructure and the vehicle which can not be verified using the limited set of instrumentation.

The second case study presented results from a series of trials investigating the performance of conductor shoe equipment. The results show that additional sensors being added to existing instrumentation architectures does not significantly affect the complexity or installation effort required and can lead to significantly more information being obtained from the data. The case study also shows that it is possible to use different transducers to measure common effects, thereby increasing confidence in the results, but highlights that systematic faults in the latter stages of the analysis chain would not be identified.

The third case study referred to an inertial measurement system fitted to a railway vehicle. The instrumentation fitted built on previous work and so was reasonably efficient in terms of sensor selection and used appropriately high sampling rates. The primary limitation highlighted by the case study is the ultimate volume of data collected being too low. This case study made the case for the use of in-service instrumentation rather than dedicated trials as the volume of data available for processing would be massively increased. It also highlighted the fact that the simple architecture used by the three case studies would not be appropriate for such long term data collection exercises, and started to hint at the data management and processing issues that in-service instrumentation would create.

The author has then presented a fourth case study in which the concept of an advanced distributed instrumentation system built upon a standard and extensible framework was developed. The project described made use of multiple instrumentation technologies, several of which could be seen in the first three case studies. The instrumentation framework developed addresses all

271

of the concerns and limitations highlighted in the other case studies, and also expands on the architecture used to present a system more suitable to the modern day railway environment - both physically and in terms of outputs. The distributed instrumentation framework developed is extensible to multiple data sources across the vehicle (or system) and makes use of local nodes to manage transducers and report the data in a standardised format based on a pre-defined but similarly extensible schema. Data processing and reporting is again managed using a nodal structure which can be customised based on the user's requirements to make use of all, or a subset, of the data available within the framework. The data from each individual node is inherently integrated to allow analysis processing to consider multiple data streams in relation to each other, and thus data can be considered in context rather than in isolation. This final stage, as demonstrated in the case study, is the key to truly understanding the data recorded, and thus making informed decisions about either the infrastructure or the rolling stock that makes up the (complex) railway system.

7.2 Case Study Outputs

As part of the core investigation presented in this thesis, a series of case studies have been undertaken and described. Each of these case studies demonstrates some aspect of an instrumentation system or analysis process based on a range of available data. Lessons from each case study have been shown to be combined into the development of a more complex instrumentation framework that allows an integrated approach to data analysis and information extraction that is key to "bigger picture thinking" and decision making within a railway context. The data processing and / or analysis performed as part of each case study has been designed to demonstrate the limitations imposed by the restricted datasets obtained, or the types of capability that the more advanced frameworks can provide. They are not exhaustive investigations into

each project showcased. However, this does not mean that a number of significant and notable railway outputs have not been observed.

The first case study presented was concerned with the energy consumed by a railway vehicle. The project was heavily supported by industry and government and so the outputs were reasonably tightly specified. The system was fitted with the co-operation of the train operating company and used a very limited sensor set. The trials were conducted in a single day, with limited academic input into the routing or train behaviour. The development of the instrumentation, other than where it interfaced with the vehicle, was academically controlled and so it was possible to increase the sampling rates beyond those required by the industrial specifications for the project.

From the enriched dataset obtained by increasing the sampling rates of the various transducers, it was possible to extract far more information than the original specifications for the project required. The results of the case study show an understanding of the efficiency of the vehicle based on several different modes of operation. From this it has been possible to hypothesise about the efficiencies of different driving styles and the importance of factors external to the vehicle such as the gradient of the route. Preliminary modelling of the traction system has also been possible based on high frequency sampling of its outputs at peak operating capacity.

The results of the first case study also showed that it is possible to use an instrumented vehicle as a measurement tool for the power network on which it operates. This output was made possible by the combination of increasing the sampling rate and controlling the routing and driving style of the vehicle. It has also been possible to quantify the differences in the outputs of research based and mandated systems by extracting from the high frequency dataset only the information available to the lower frequency approach.

The second case study presented was an investigation into the dynamics of the mechanisms

by which third rail electric vehicles collect their power. The instrumentation developed falls into the category of developmental but the project itself evolved from previous troubleshooting work. The industrially specified outputs of the system were to verify the performance of existing shoegear in the context of existing standards and to develop a greater understanding of the forces and displacements involved in mechanism to inform future standard development work.

In addition to these outputs, the inclusion of one extra strain gauge bridge in the instrumentation allowed an investigation into the torque within the shoegear system. This allowed the results from the force component of the system to be verified by the displacement component and also went towards a greater understanding of the mechanisms by which conductor shoes clear ice from the third rail which was the previous, related, piece of work. Data analysis, beyond that required by the industrially specified outputs, has also lead to quantification of hysteretic behaviour and operating modes in the shoegear system. This information has been discussed, at length, with the shoegear manufacturers.

The data has also, beyond the original scope, been used to indicate how components of the instrumentation could be used for condition monitoring of the third rail and spawned a subsequent project in this area. The data from the system described in the case study has been shown to compare favourably with previous third rail measurement and gauging work.

The instrumentation for the third case study was closely linked to the second one, but the objectives were focused on condition monitoring of the running rails and track systems rather than general development work undertaken on subsystems of the railway vehicle. Processing of the dataset revealed some areas of track in poor condition and large portions of the route have subsequently been re-laid. Furthermore, feature extraction more closely linked to the original S&C objectives allowed identification of transitional areas of the railway and has subsequently lead to further work focused on transitions.

The final case study showed the integration of many of the components seen in the earlier work. Systems were developed around an instrumentation framework which allowed far greater instrumentation complexity and thus the richness of the dataset was dramatically increased. Furthermore, the new framework meant that the trial process could be upgraded to include full in-service trials which dramatically increased the volume of data recorded and the potential benefits from analysis. This added complexity, however, required a full VAB process to be undertaken.

The ongoing nature of in-service trials means that the dataset obtained from the instrumentation is still expanding. Already, however, lots of information has been extracted from the data. Results are being used to support recommendations towards efficient driving style, both in terms of driver behaviour and the use of the environment. The traction efficiencies of the vehicle are, after 30 years, finally being understood allowing driver training to be standardised. The routes around the network are being considered to draw a balance between journey time and energy efficiency. And the efficiency of the network is even being questioned with further instrumentation work being undertaken to identify the losses throughout the entire DC system.

7.3   The future of Instrumentation Systems

The railway is an increasingly complex system; this is set to become increasingly clear in the near future with the introduction of things like faster trains and new signalling systems, the increasing prevalence of technologies such as platform screen doors and automatic train control, and of course increasing capacity to support growing popularity. As the complexity of the system increases, so the need to understand it better and thus better advise decisions taken within the railway is enhanced. This thesis has shown that is is possible to use more complex instrumentation systems based around advanced instrumentation frameworks to improve the

level of understanding of what currently occurs in the railway and thus hopefully to improve the quality of decisions made in relation to it. The integration of multiple, previously disparate data streams to allow data processing and analysis in a contextual way is key to the development of the "data driven railway".

The systems showcased in this thesis are based around rolling stock, but the framework proposed could equally be applied to infrastructure monitoring. Instrumentation of both infrastructure and rolling stock is now becoming more prevalent with fixed assets such as points machines now coming with some embedded condition monitoring systems, and new rolling stock having increasingly advanced on train monitoring equipment. The integration of the signals from different data streams within or between these monitoring systems is key to obtaining maximum benefit from this increased level of instrumentation within the railway. Frameworks, such as the one described in this thesis, will be needed to ensure compatible instrumentation systems between which processing benefits can be shared, and modular instrumentation components can be used. Without this, custom instrumentation and processing systems will need to be developed for each and every application. While not a major undertaking for simple architectures, this would be hugely detrimental to overall efficiency if applied to the larger instrumentation systems required to obtain the benefits of multiple signal integration and contextual processing.

In the future, more onboard processing systems will undoubtedly be found as part of instrumentation systems attached to railway vehicles. It is also likely that, without large data transfer and management overheads, more railway vehicles will have instrumentation systems fitted. This will lead to a more highly instrumented railway network, about which more is known, and thus the condition of which can be maintained more highly.

Long term, in-service, trials are clearly of benefit to strategic projects, however, many trouble shooting projects are generally better served by shorter term instrumentation. While the work

presented here suggests that an increased dataset is beneficial to development work, it is worth considering that a dataset can be populated through repeated single day trials as well as in-service measurement. Single day trials allow instrumentation to be prototyped or modified, possibly through the addition of nodes to modular systems, based on intermediate findings, and improve the possibility for controlling trials through routing etc. Significantly, the approvals processes are also reduced. Clearly some projects, such as the driving style analysis from the fourth case study, are better suited to in-service trials, but they are not suitable for every application, and vital industrial support may be more forthcoming with the reduced commitment that single day trial work involves.

The collection of huge datasets will always have a place in the development of technology as, by definition, it is not always possible to know all the questions that one will want to ask in the future at the point of designing the instrumentation. More and more, however, the volume of data being returned is becoming harder to manage and process while at the same time embedded processing power is increasing. As a result, systems that process the raw transducer signals and record only results are becoming more popular. This is particularly the case in troubleshooting and condition monitoring systems where specific questions are being asked. In this thesis, brief consideration has been given to how this kind of in-line / real-time processing could be achieved within the proposed framework. This kind of extensibility, combined with the modular nature, is crucial to the uptake of such systems which may have service lives akin to that of the vehicles of systems to which they are fitted.

From the work described in the four case studies that form the core of this thesis, several further projects have been developed. The first case study, performed some time ago, developed to allow the fourth to occur. The second has lead to a further third rail instrumentation project which is now a key part of the Network Rail third rail maintenance strategy. The third has lead

to instrumentation work looking at transition sections of the infrastructure as well as a general track condition monitoring project. Finally the fourth has lead to infrastructure instrumentation looking at the efficiencies of third rail networks, as well as its ongoing data processing work. This ongoing development of work is essential to the research group and allows continuity in both the development of instrumentation systems and the themes in which they are used.

APPENDIX A

APPROVALS DOCUMENTS FOR DISTRIBUTED

EMBEDDED INSTRUMENTATION SYSTEM

A.1   VAB Approvals Document

[Not available in the digital version of this thesis]

APPENDIX B

RELATED JOURNAL PAPERS

B.1   Using bogie-mounted sensors to understand the dynamics of third rail current collection systems

B.2   Condition monitoring opportunities using vehicle-based sensors

[Not available in the digital version of this thesis]

GLOSSARY

A

AETB   Air-to-Electric Transducer - Braking.

AETW   Air-to-Electric-Transducer - Weight.

ATP   Automatic Train Protection.

AVI   Automated Video Inspection.

C

CAN   Controller Area Network.

CARS   China Academy of Railway Science.

CCD   Charge-Coupled Device.

CDMA   Code Division Multiple Access.

D

DfT   Department for Transport.

DLC   Data Length Code.

DLR   Docklands Light Railway.

DMA   Direct Memory Access.

DMU   Diesel Multiple Unit.

DSP   Digital Signal Processor.

E

EMC   Elecro-Magnetic Compatibility.

ERCIR   Enhanced Rail Contribution through Improved Reliability.

F

FAT   File Allocation Table.

FAT32   File Allocation Table - 32 bit.

FIFO   First In First Out.

FOG   Fibre-Optic Gyroscope.

G

GIS   Geographic Information System.

GPS   Global Positioning System.

H

HABD   Hot Axle Box Detector.

HARL   Height Above Rail Level.

I

I2C   Inter-Integrated Circuit.

ICE   Inter-City Express.

IM   Infrastructure Manager.

IMU   Inertial Measurement Unit.

INS   Inertial Navigation System.

ISTec   Institute for Safety Technology.

L

LED   Light Emitting Diode.

M

MA   Motor / Alternator.

MAS   Measurement & Assessment System.

MELF   Metal Electrode Leadless Face.

MEMS   Micro Electro-Mechanical System.

MOSFET   Metal Oxide Semiconductor Field Effect Transistor.

N

NCDT   Non-Contact Displacement Transducer.

NMEA   National Maritime Electronics Association.

NMT   New Measurement Train.

O

OBB   Austrian Federal Railway.

ORR   Office of Rail Regulation.

P

PABS   PAssenger Billing System.

PWI   Permanent Way Institute.

R

RAVERS   Rail And VEhicle Records System.

RC   Resistor / Capacitor.

ROSCO   Railway Operating Support COmpany.

RSSB   Railway Safety and Standards Board.

RTCC   Real Time Clock Calendar.

RWCOMOS   Railway COndition MOnitoring System.

S

S&C   Switch and Crossing.

SD   Secure Digital.

SDHC   Secure Digital High Capacity.

SDXC   Secure Digital eXtended Capacity.

SID   Standard IDentifier.

SPI   Serial Peripheral Interface.

SSB   Sensor System Box.

T

TGRIS Track Geometry Real-time Instrumentation System.

TMD   Train Maintenance Depot.

TOC   Train Operating Company.

TRL   Technology Readiness Level.

TSI   Technical Standard for Interoperability.

TTL   Transistor to Transistor Logic.

U

UAV   Unmanned Aerial Vehicle.

V

VAB   Vehicle Acceptance Body.

VSG   Vibrating Structure Gyroscope.

W

WCML   West Coast Main Line.

WCRR   World Congress on Railway Research.

WILMS Wheel Impact Load Measurement System.

WSP   Wheel Slide Protection.

BIBLIOGRAPHY

[1] P. Hooper, "Information on technical standards for interoperability (TSIs)," tech. rep., June 2009.

[2] J. Young, W.R. and W. Wilson, "Monitoring electric vehicle performance," in *Southcon/95. Conference Record*, pp. 298 –302, mar 1995.

[3] Director, Research Directorate, US Department Of Defense (DOD), "Technology readiness assessment (TRA) deskbook." Download from:
http://www.dod.mil/ddre/doc/DoD TRA July 2009 Read Version.pdf, 2009.

[4] V. Gupalov, A. Mochalov, and A. Boronachin, "Application of a ring laser for measurement of track geometrical parameters," *Proceedings of the SPIE*, vol. 4316, no. 13, 2000.

[5] Campbell Scientific, "Cr-800 series measurement and control systems." Download from:
http://www.campbellsci.com/documents/product-brochures/b_cr800.pdf.

[6] LEM, "EM4T energy meter for traction usable on each network." Download from:
http://www.lem.com/images/stories/files/Products/1-5-3_traction_trackside
/CH_26100e_em4t.pdf.

[7] B. Cene, A. Radosavlejevic, and B. Garvrilovic, "Driving energy consumption of electric traction vehicles," *Communciations Reveiw*, no. 3, pp. 48–54, 2008.

[8] M. Computing, "LGR-5320 series: Stand-alone, high-speed, multifunction data loggers." Datasheet available from:
http://www.mccdaq.com/products/LGR-5320-Series.aspx.

[9] N. Instruments, "NI USB-6009: 14-bit, 48 ks/s low-cost multifunction daq." Datasheet available from:
http://sine.ni.com/nips/cds/view/p/lang/en/nid/201987.

[10] N. Instruments, "Ni compactrio: Rugged and reconfigurable control and monitoring system." Datasheet available from:
http://www.ni.com/compactrio/.

[11] A. McNaughton, "Track inspection: Moving from man to machine," *Modern Railways*, vol. 61, pp. 66–74, March 2004.

[12] W. M. Kaufman and J. E. Borntraeger, "Automated gauge inspection," *Industry Applications, IEEE Transactions on*, vol. IA-10, pp. 567 –572, sept. 1974.

[13] R. H. Hronik, "Apparatus for high-speed measurement of track geometry," March 1970.

[14] Y. Jin, Y. Goto, Y. Nishimoto, H. Naito, Y. Fukumoto, and A. Iwake, "Advanced dynamic obstacle-detecting system for railway surroundings using a highly accurate laser-sectioning method," in *Industrial Electronics, Control, Instrumentation, and Automation, 1992. Power Electronics and Motion Control., Proceedings of the 1992 International Conference on*, pp. 1599 –1603 vol.3, nov 1992.

[15] M. Benfield, "The modern gauging process," *Journal and Report of Proceedings of the Permanent Way Institution*, vol. 127, no. 3, pp. 128–130, 2009.

[16] E. Yazawa, "Track inspection technologies," *Railway Technology Avalanche*, January 2003.

[17] S. Ono, A. Numakura, and T. Odaka, "High-speed track inspection technologies," *JR East Technical Review*, vol. 2, pp. 9–13, 2003.

[18] H. McAnaw, "The system that measures the system," *Non Destructive Testing and Evaluation International*, vol. 36, pp. 169–179, 2003.

[19] R. Lewis, "Track-recording techniques used on british rail," *Electric Power Applications, IEE Proceedings B*, vol. 131, pp. 73 –81, may 1984.

[20] A. Jablonski, "T353 scoping and development of the strategic vehicle-track interaction model: report on the scoping study," 2004.

[21] G. Presle, "The em 250 high-speed track recording coach and the em-sat 120 track survey car, as networked track geometry diagnosis and therapy systems," *Rail Engineering International*, vol. 41, pp. 31–35, September 2001.

[22] B. Lichtberger, "The certification procedure for track maintenance machines," *European Railway Review*, vol. 2, pp. 1–8, 2008.

[23] P. Strange, "Absolute track geometry on the west coast route modernisation," *Journal and Report of Proceedings of the Permanent Way Institution*, vol. 127, no. 3, pp. 119–127, 2009.

[24] K. Takeshita, "Track irregularity inspection method by revenue-earning train," *Railway Technical Reseearch Institute Report*, vol. 9, January 1995.

[25] Y. Sunaga, I. Sano, and T. Ide, "A method to control the short wave track irregularities utilizing axlebox acceleration," in *Railway Technical Research Institute, Quarterly Reports*, vol. 38, pp. 176–181, 2002.

[26] Q. Tang, "Railway track geometry real-time inspecting system," in *Instrumentation and Measurement Technology Conference, 1992. IMTC '92., 9th IEEE*, pp. 656 –660, may 1992.

[27] D. Hentschel, B. Frankenstein, K. Frohlich, P. Liewers, and M. Trepte, "Approaches to the continuous monitoring of bogies in service," *Rail Technology Review*, vol. 4, pp. 17–20, 2004.

[28] M. Bryan, "Methond and apparatus for monitoring railway defects," February 1999.

[29] R. Sunder, A. Kolbasseff, K. Kieninger, A. Rohm, and J. Walter, "Operatinal experiences with onboard diagnosis system for high speed trains," in *Proceedings of the 5th World Congress on Railway Research*, 2001.

[30] M. Saglitz and R. Sunder.

[31] A. J. Savage, "Modern methods of track inspection," *Transport Research Laboratory Reports*, pp. 1–9, 2002.

[32] Y. Misu, K. Kaneko, H. Hikma, and K. Hirabayashi, "Unusual behavior detection and non-dismantling diagnosis by monitoring bogie conditions," in *Proceedings of the 5th World Congress on Railway Research*, 2001.

[33] M. Plainer, C. Kral, and G. Singer, "Investigation of the longitudinal dynamic effects before and after the derailment of a freight train," in *Rail Conference, 2004. Proceedings of the 2004 ASME/IEEE Joint*, pp. 75 –81, april 2004.

[34] K. Yoshimoto, K. Kataoka, and K. Komaya, "A feasibility study of train automatic stop control using range sensors," in *Intelligent Transportation Systems, 2001. Proceedings. 2001 IEEE*, pp. 802 –807, 2001.

[35] M. Lockyear, "The application of a transmission based moving block automatic train control system on docklands light railway," in *Developments in Mass Transit Systems, 1998. International Conference on (Conf. Publ. No. 453)*, pp. 51 –61, apr 1998.

[36] S. Morar, "Evolution of communication based train control worldwide," in *Railway Signalling and Control Systems (RSCS 2010), IET Professional Development Course on*, pp. 281 –289, june 2010.

[37] British Rail Engineering Limited (BREL), "Class 508 maintenance manual," 1979.

[38] M. G. Wing, A. Eklund, and L. D. Kellogg, "Consumer-grade global positioning system (gps) accuracy and reliability," *Journal of Forestry*, vol. 103, no. 4, pp. 169–173, 2005.

[39] J. Zywiel and G. Oberlechner, "Innovative measuring system unveiled," *International Railway Journal*, no. 3, pp. 14–16, 2000.

[40] P. Weston, C. Ling, C. Roberts, C. Goodman, P. Li, and R. Goodall, "Monitoring vertical track irregularity from in-service railway vehicles," *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, vol. 221, no. 1, pp. 75–88, 2007.

[41] E. Sackl, "The em-sat 120 track survey car, an integrated part of the track geometry data base of the austrian federal railways öbb," *Rail Technology Review*, no. 2, pp. 39–43, 2004.

[42] F. Mahmood, "Rr06-08: Real-time vertical track modulus measurement system from a moving railcar," *Federal Railroad Administration, Research Results*, pp. 1–4, 2006.

[43] B. Nejikovsky and E. Keller, "Wireless communications based system to monitor performance of rail vehicles," in *Railroad Conference, 2000. Proceedings of the 2000 ASME/IEEE Joint*, pp. 111 –124, 2000.

[44] P. Ackroyd, S. Angelo, and J. Stevens, "Remote ride quality monitoring of acela train set performance," in *Railroad Conference, 2002 ASME/IEEE Joint*, pp. 171 –178, april 2002.

[45] J. D. III, W. M. Zavis, D. G. Toth, S. K. Punwani, M. F. Stewart, and M. C. Edwards, "Revenue service demonstration of on-board condition monitoring system," *ASME Conference Proceedings*, vol. 2003, no. 37270, pp. 115–120, 2003.

[46] T. Tsai and S. Sill, "Rr00-02: Developed wheel and axle assembly monitoring system to improve passenger train safety," *Federal Railroad Administration, Research Results*, pp. 1–3, 2000.

[47] R. Deuce and J. Rosinski, "Ultra miniature data loggers for testing and in-service monitoring of railway axles and other rail vehicle components," in *Railway Condition Monitoring, 2006. The Institution of Engineering and Technology International Conference on*, pp. 57 –60, nov. 2006.

[48] P. Keen, R. Phillpotts, and S. Conway, "An instrumented pantograph for high speed current collection measurements," in *Railway Condition Monitoring, 2008 4th IET International Conference on*, pp. 1 –5, june 2008.

[49] S. Mittal, S. Kalra, M. Shamishi, and B. Sharma, "Micro-controller based oscillation monitoring system for the safety of railway vehicles with high storage capacity and real time warning facility," *Journal of Scientific and Industrial Research*, vol. 63, no. 9, pp. 752–757, 2004.

[50] N. Nenov, T. Rouzhekov, G. Mihov, and E. Dimitrov, "Strength sensor for dynamic wheel load measuring of railway carriages," in *Electronics Technology: Integrated Management of Electronic Materials Production, 2003. 26th International Spring Seminar on*, pp. 260 – 265, may 2003.

[51] P. Wolfs, S. Bleakley, S. Senini, and P. Thomas, "An autonomous, low cost, distributed method for observing 'vehicle track interactions," in *Proceedings of JRC2006 Joint Rail Conference*, American Society for Mechanical Engineers, pp. 279–286, April 2006.

[52] G. Shafiullah, A. Gyasi-Agyei, and P. Wolfs, "Survey of wireless communications applications in the railway industry," in *Wireless Broadband and Ultra Wideband Communications, 2007. AusWireless 2007. The 2nd International Conference on*, p. 65, aug. 2007.

[53] Z. Liu, "A protocol analysis method of acceleration sensor in maglev train," in *Networking, Sensing and Control, 2006. ICNSC '06. Proceedings of the 2006 IEEE International Conference on*, pp. 342 –345, 0-0 2006.

[54] M. Edwards, I. Donelson, J., W. Zavis, A. Prabhakaran, D. Brabb, and A. Jackson, "Improving freight rail safety with on-board monitoring and control systems," in *Rail Conference, 2005. Proceedings of the 2005 ASME/IEEE Joint*, pp. 117 – 122, march 2005.

[55] H. M. Hawari, *Minimising Track Degradation through Managing Vehicle / Track Interaction.* PhD thesis, Faculty of Built Environment and Enginnering, School of Urban Development, Queensland University of Technology, 2007.

[56] C. Wang, Q. Xiao, H. Liang, Y. Liu, and X. Cai, "On-line monitoring of railway deformation using acceleration measurement," in *Proceedings of the 6th World Congress on Intelligent Control and Automation*, pp. 5828–5832, June 2006.

[57] C. Alippi, E. Casagrande, F. Scotti, and V. Piuri, "Composite real-time image processing for railways track profile measurement," *Instrumentation and Measurement, IEEE Transactions on*, vol. 49, pp. 559 –564, jun 2000.

[58] C. Alippi, E. Casagrande, M. Fumagalli, F. Scotti, V. Piuri, and L. Valsecchi, "An embedded system methodology for real-time analysis of railways track profile," in *Instrumentation and Measurement Technology Conference, 2002. IMTC/2002. Proceedings of the 19th IEEE*, vol. 1, pp. 747 – 751 vol.1, 2002.

[59] F. Attivissimo, A. Danese, N. Giaquinto, and P. Sforza, "A technique to charactererize the wheel-rail interaction by means of a diagnostic railway vehicle," in *Instrumentation and Measurement Technology Conference, 2004. IMTC 04. Proceedings of the 21st IEEE*, vol. 3, pp. 1680 – 1685 Vol.3, may 2004.

[60] F. Attivissimo, A. Danese, N. Giaquinto, and P. Sforza, "A railway measurement system to evaluate the wheel & rail interaction quality," *Instrumentation and Measurement, IEEE Transactions on*, vol. 56, pp. 1583 –1589, oct. 2007.

[61] J. Kawasaki and K. Youcef-Toumi, "Estimation of rail irregularities," in *American Control Conference, 2002. Proceedings of the 2002*, vol. 5, pp. 3650 – 3660 vol.5, 2002.

[62] A. Mirabadi, N. Mort, and F. Schmid, "Application of sensor fusion to railway systems," in *Multisensor Fusion and Integration for Intelligent Systems, 1996. IEEE/SICE/RSJ International Conference on*, pp. 185 –192, dec 1996.

[63] T. Maly, M. Rumpler, and K. Schweinzer, "Joining sensor systems for railway vehicle inspection," in *Sensors, 2004. Proceedings of IEEE*, pp. 12 – 15 vol.1, oct. 2004.

[64] T. Maly, M. Rumpler, H. Schweinzer, and A. Schoebel, "New development of an overall train inspection system for increased operational safety," in *Intelligent Transportation Systems, 2005. Proceedings. 2005 IEEE*, pp. 188 – 193, sept. 2005.

[65] F. Leahy, M. Judd, and M. Shortis, "Measurement of railway profiles using gps integrated with other sensors," in *Vehicle Navigation and Information Systems Conference, 1993., Proceedings of the IEEE-IEE*, pp. 706 –709, oct 1993.

[66] P. Weston, C. Ling, C. Goodman, C. Roberts, P. Li, and R. Goodall, "Monitoring lateral track irregularity from in-service railway vehicles," *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, vol. 221, no. 1, pp. 89–100, 2007.

[67] S. Saab, "A map matching approach for train positioning. i. development and analysis," *Vehicular Technology, IEEE Transactions on*, vol. 49, pp. 467 –475, mar 2000.

[68] S. Saab, "A map matching approach for train positioning. ii. application and experimentation," *Vehicular Technology, IEEE Transactions on*, vol. 49, pp. 476 –484, mar 2000.

[69] J. Yang, Q. Huang, and W. Li, "A bench study on lateral semi-active control system for high-speed railway vehicle," in *Mechatronics and Automation, Proceedings of the 2006 IEEE International Conference on*, pp. 1335 –1339, june 2006.

[70] H. Zamzuri, A. Zolotas, and R. Goodall, "Intelligent control approaches for tilting railway vehicles," *Vehicle System Dynamics*, vol. 44, pp. 834–842, 2006.

[71] B. Allotta, L. Pugi, and F. Bartolini, "Design and experimental results of an active suspension system for a high-speed pantograph," *IEEE/ASME TRANSACTIONS ON MECHATRONICS*, vol. 13, no. 5, pp. 548–557, 2008.

[72] S. Kia, H. Henao, and G.-A. Capolino, "Mechanical health assessment of a railway traction system," in *Electrotechnical Conference, 2008. MELECON 2008. The 14th IEEE Mediterranean*, pp. 453 –458, may 2008.

[73] Y. Han, K. Kim, S. Seo, C. Park, S. Han, and Y. Kim, "Traction system characteristics of railway vehicle," in *Proceedings of ICCAS2005 - International Conference on Control Automation and Systems*, June 2005.

[74] Y.-J. Han, S.-W. Kim, K.-H. Kim, T.-H. Lee, and H.-M. Koo, "A study on mechanical characteristics for induction motor of rolling stock," in *SICE-ICASE, 2006. International Joint Conference*, pp. 718 –721, oct. 2006.

[75] N. Lehrasab and S. Fararooy, "Intelligent multiple sensor early failure warning system for train rotary door operator," in *Proceedings of the 1996 IEE Colloquium on Multiple-Sensor Data Fusion*, pp. 1411–1419, November 1996.

[76] N. Lehrasab, S. Fararooy, and J. Allan, "Fault detection in intelligent early failure warning sensors system for train rotary door operator,"

[77] S. Hillmansen, "Memorandum to parliament: Low carbon technologies in a green economy." Available from:
http://www.publications.parliament.uk/pa/cm200910/cmselect/cmenergy/193
/193we30.htm, 2009.

[78] Parliamentary Comiitee on Climate Change, "Rail transport comission to the committe on climate change." Available from:
http://www.theccc.org.uk/pdfs/080521%20CCC%20rail%20submission.pdf, May 2008.

[79] G. McFadden, "Sustainability and carbon reduction on electrified railways," in *Railway Electrification Infrastructure and Systems, 2009. REIS 2009. 4th IET professional Development Course on*, pp. 246 –254, june 2009.

[80] Railway Safety and Standards Board (RRSB), "GM/RT2132 on-board energy metering for billing purposes," tech. rep., September 2010.

[81] British Standards, "BS EN 50463 energy measurement on board trains," tech. rep., 2007.

[82] British Standards, "BS EN 16001 energy management systems," tech. rep., 2009.

[83] C. J. (Editor), "Metering validates regeneration saving," *Railway Gazette International*, p. 486, August 2008.

[84] Association of Train Operating Companies (ATOC), "List of members." Available from: http://www.atoc.org/train-companies.

[85] G. Hull, *Simulation of Energy Efficiency Improvements on Commuter Railways.* PhD thesis, School of Electronic, Electrical and Computer Systems Enginnering, College of Engineering and Physical Sciences, The University of Birmingham, 2009.

[86] British Standards, "BS EN 50163 supply voltages of traction systems," tech. rep., 2007.

[87] LEM, "DV1200/SP2 dv series voltage transducer for insulated high voltage measurements." Download from:
http://www.lem.com/images/stories/files/Products/1-5-3_traction_trackside
/LEM_DV_Leaflet_9_9_08.pdf.

[88] British Standards, "BS EN 50155 electronic equipment used on rolling stock," tech. rep., 2007.

[89] British Standards, "BS EN 50121-3-2 electromagnetic compatibility," tech. rep., 2006.

[90] British Standards, "BS EN 50124-1 insulation coordination," tech. rep., 2001.

[91] LEM, "HTR 600-SF current transducer datasheet." Download from:
http://www.lem.com/docs/products/ltc%20600-sf%20e.pdf.

[92] LEM, "LTC 50-SB current transducer datasheet." Download from:
http://www.lem.com/docs/products/htr%2050%20to%20500-sb%20e.pdf.

[93] M. Bundy and T. Ohlemiller, "Bench-scale flammability measures for electronic equipment," tech. rep., July 2003.

[94] L. Technology, "LTC 1859 analogue to digital converter datasheet." Download from:
http://cds.linear.com/docs/Datasheet/185789fa.pdf.

[95] British Standards, "BS 6853 codes of practice for fire precautions in the design and construction of passenger carrying trains," tech. rep., 1999.

[96] A. Anderson, "Robert davidson - father of the electric locomotive," *Proceedings of History of Electrical Engineering Summer Weekend, Proc. IEE*, no. 8, pp. 1–17, 1975.

[97] M. C. Duffy, *Electric Railways, 1880-1990 (IEE history of technology series).* The Institution of Engineering and Technology, 2003.

[98] T. Lascelles, "The city and south london railway," 1955.

[99] D. Hartland, "Electric contact systems - passing power to the trains," *Proceedings of the IEE Second Residential Course on Railway Electrification Infrastructure and Systems (REIS)*, May 2005.

[100] Network Rail Company Standards, "RT/E/C/27010 compatibility between electric trains and electrification systems," tech. rep., 1997.

[101] Railway Safety and Standards Board (RRSB), "T942 pantograph sway acceptance requirements and methodology," tech. rep., October 2011.

[102] J. Chiltham, "Southeastern trains third rail height survey," tech. rep., 2004.

[103] H. Bai, C. Roberts, and C. Goodman, "A generic fault detection and diagnosis approach for railway assets," *Proceedings of the IET International Conference on Railway Engineering*, vol. 1, pp. 164–171, 2008.

[104] Safety and Standards Directorate Railtrack PLC, "GM/RT 2100 design scrutiny for the engineering acceptance of rail vehicles," tech. rep., December 1997.

[105] P. F. Weston, C. Roberts, C. J. Goodman, and C. S. Ling, "Condition monitoring of railway track using in-service trains," in *Railway Condition Monitoring, 2006. The Institution of Engineering and Technology International Conference on*, pp. 26 –31, nov. 2006.

[106] C. Ward, P. Weston, E. Stewart, H. Li, R. Goodall, T. Mei, G. Charles, and R. Dixon, "Conditition monitoring opportunities using vehcile-based sensors," *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, vol. 225, no. 2, pp. 202–218, 2011.

[107] E. Stewart, P. Weston, S. Hillmansen, and C. Roberts, "Using bogie-mounted sensors to understand the dynamics of third rail current collection systems," *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, vol. 225, no. 2, pp. 219–227, 2011.

[108] P. Li, R. Goodall, P. Weston, C. S. Ling, C. Goodman, and C. Roberts, "Estimation of railway vehicle suspension parameters for condition monitoring," *Control Engineering Practice*, vol. 15, no. 1, pp. 43 – 55, 2007.

[109] S. Grassie, "Measurement of railhead longitudinal profiles: a comparison of different techniques," *Wear*, vol. 191, no. 1-2, pp. 245 – 251, 1996. 4th International Conference on Contact Mechanics and Wear of Rail-Wheel Systems.

[110] R. Barrett, "Low frequency machinery monitoring: Measurement considerations," tech. rep., August 1993.

[111] E. Stewart, P. Weston, S. Hillmansen, and C. Roberts, "The merseyrail energy monitoring project," *Proceedings of the 9th World Congress on Railway Research*, 2011.

[112] Network Rail, "EC4T Metering Rules," tech. rep., 2011.

[113] U. Kiencke, S. Dias, and M. Litschel, "Automotive serial controller area network," *SAE International*, pp. 1 – 6, 1986.

[114] British Standards, "BS ISO 11898-5 controller area network (CAN)," tech. rep., 2007.

[115] C. Landi, M. Luiso, and N. Pasquino, "A remotely controlled onboard measurement system for optimization of energy consumption of electrical trains," *Instrumentation and Measurement, IEEE Transactions on*, vol. 57, pp. 2250 –2256, oct. 2008.

[116] SD Group Technical Committee of the SD Card Association, "SD Specifications Part 1 Physical Layer Simplified Specification Version 3.01," tech. rep., May 2010.

[117] K. Payne, "LUL 2-01001-002 fire safety performance of materials," tech. rep., Apriil 2005.

[118] Network Rail, "Catalogue of railway code systems," tech. rep., December 2010.

[119] Y. Bocharnikov, A. Tobias, C. Roberts, S. Hillmansen, and C. Goodman, "Optimal driving strategy for traction energy saving on dc suburban railways," *Electric Power Applications, IET*, vol. 1, pp. 675 –682, sept. 2007.

[120] R. Lewis and A. Richards, "A compensated accelerometer for the measurement of railway track crosslevel," *Vehicular Technology, IEEE Transactions on*, vol. 37, pp. 174 –178, aug 1988.

[121] B. Mellitt, C. Goodman, and R. Arthurton, "Simulator for studying operational and power supply conditions in rapid transit railways," *Proceedings of the Institution of Electrical Engineers*, vol. 125, no. 4, pp. 298–303, 1978.