THE UNIVERSITY OF
NEW SOUTH WALES

SCHOOL OF ELECTRICAL ENGINEERING
AND TELECOMMUNICATION

Sunswift IV Strategy for the 2011 World Solar Challenge

by

Alexandra Boulgakov

Thesis submitted as a requirement for the degree
Bachelor of Engineering (Electrical Engineering)

Submitted: January 10, 2012 Student ID: z3207632
Supervisor: Gernot Heiser, Ray Eaton Topic ID:GH91

Abstract

Sunswift IV is the UNSW Solar Racing Team's most recent solar vehicle, and their entry into the 2011 World Solar Challenge, a 3000km race from Darwin to Adelaide. In order to be competitive in this race, the team needed to improve the accuracy of the models and reliability of the instrumentation used in their strategy system, which is responsible for calculating the optimal speed at which solar cars run the race. During the course of this thesis, roll-down tests and analysis is carried out to model the mechanical systems of the car, and an accurate first-order model for the solar car battery is developed. The current integrator device, which measures the solar car battery state of charge, is redesigned, and generally performs well in bench and field tests. The accuracy of the mechanical model requires some improvement, and suggestions are made as to how to use the roll-down method in the future to do this.

Acknowledgments

Firstly, I would like to thank my supervisor Gernot Heiser, not only for his support throughout this project, but also for providing much-appreciated encouragement and resources to the *Sunswift* electrical team for many years. Thank you also to my co-supervisor Ray Eaton, for supporting the team and making the administrative process of writing a thesis during a World Solar Challenge race semester so easy.

I am extremely grateful to David Snowdon, for being an indispensable mentor, for sharing so much of his knowledge with me, and for building and improving the *Scanalysis* user interface, which was invaluable during the 2011 World Solar Challenge. Thanks also for being there and providing support in the lead up to, and during the race.

I would also like to thank all members of the *Sunswift* team, for the time and effort that you all invested into *IVy* and the race. In particular, thank you to Tommy Heyser, Daniel Friedman and Irving Tjiptowarsono, without whose hard work and motivation I would not have been able to take the car testing and collect data as much as I did. Thank you also to Campbell McLaren for being my co-strategist during the race, and for discussing many ideas during the course of the year: I could not have done it without you. In addition, a special thanks to Etienne Le Sueur, for organising this year's *Sunswift* electrical team, encouraging us and providing much-needed technical help. I'd also like to extend a big thank you to Joe Wolfe, for being the team's academic supervisor during another race adventure, and for many useful discussions.

I am also very grateful to everyone who took interest in my thesis work and for the many insightful conversations I was lucky to have during the course of the year. In particular, thanks to Graham Doig for your encouragement and guidance during the roll-down tests we did, and for answering so many of my aerodynamics questions. Thank you also to Nina Naroditsakaya, for taking the time to help me understand optimisation methods.

A big thanks to Scott McDaid, for always being there when I needed motivation, a bouncing board for ideas, or desperate technical help at odd times of day and night. Thanks also to David Favaloro and Robyn Hutchinson, for providing much needed support and comic relief during that last month of thesis writing.

Last but not least thank you to my family, for supporting me, sharing my excitement and encouraging my involvement with the *Sunswift* project over the last four years.

Contents

Chapter 1

Introduction

1.1 The World Solar Challenge

Hans Tholstrup, Larry and Garry Perkins made history in the summer of 1983, successfully crossing the Australian continent from Perth to Sydney in the world's first solar car, *The BP Quiet Achiever*. This event marked the birth of solar racing as a sport. Their transcontinental journey drew significant media attention, encouraging Hans to set up the *World Solar Challenge* (WSC). The first WSC, a 3000 km race from Darwin to Adelaide for solar powered cars, was held in 1987, attracting 23 entries from 7 countries [Roche et al., 1997]. The WSC is now held every two years, and is considered to be the premier endurance competition for solar-electric vehicles.

The event aims to promote sustainable transport and to provide a competitive environment in which to develop efficient vehicle technology. Although the solar cars built for the event may never become a practical transport solution themselves, the race has lead to many technological developments that trickle down into the electric vehicle industry. For example, some of the most efficient electric motor controllers, such as the Tritium company's *Wavesculptor* series, were initially designed for solar cars and later upgraded to suit larger electric vehicles [Tritium Power Engineering Group, 2008].

1.2 The sport of solar racing

In order to perform well in the WSC, teams have to focus their efforts on two main factors. The first of these is the optimisation of the design of their car. Solar cars are built for efficiency and speed, and the top vehicles can travel at over 80 km/h on 1 kW of power. To do this, they must be light, aerodynamic, and have very a efficient electric drive train: almost all the solar energy collected by the car's solar panels reaches its motor. On top of this, the cars must be able to run continuously and reliably for hundreds of kilometres at a time.

Figure 1.1: World Solar Challenge Route. *Image courtesy of [Sharp-World, 2011]*

The second deciding factor is how well the the car's energy use is managed during the race, i.e. the solar car racing strategy. Solar cars in the WSC carry a battery pack, which makes strategic planning possible. The battery acts as an energy buffer, where collected solar energy can be stored in times of excess, and retrieved when needed. While there is charge in the pack, the solar car can travel at any chosen speed, despite variations in weather and terrain [Mocking, 2006]. Without it, the only optimal choice would be to travel at the speed that uses all the power collected by the solar car's solar panels at that instant.

Since a solar car's energy intake and expenditure is greatly affected by changing and uncertain weather conditions, a single optimal race strategy cannot be planned ahead of time. Instead, teams develop a *strategy system*: a collection of car and environment models, sensors, optimisation methods and simulation software whose purpose it is to calculate an optimal race speed profile.

1.3 The UNSW Solar Racing Team and *Sunswift IV*

The UNSW Solar Racing Team (from here on referred to as *the Team* or *Sunswift*) was founded by 4th year engineering student Byron Kennedy, and has competed in every WSC since 1996. During this time, the Team has raced four different car designs, with the most recent, *Sunswift IV*, achieving first place in the Silicon Challenge class of the 2009 WSC and fourth place overall.

Although undoubtedly a good result, after the 2009 race the Team identified several areas in which the car's performance could be improved, with the strategy system being one of the main priorities. The two main strategy-related issues were a lack of accurate model data, and reliability problems with one of the most important sensors in the car, the *current integrator*, which measures the car's battery level.

Carrying out these improvements to the strategy system in preparation for the 2011 WSC became the motivation for this thesis. As the Team was not building a new car from scratch for this race, but merely improving our *Sunswift IV*, we believed we had a rare opportunity to collect model data and to test the car months in advance. The 2011 WSC race, which was held in October of this year, also presented the perfect chance for me to take the role of strategist for the Team, and test my knowledge and the improved system first hand.

1.4 Outline

This document firstly introduces the background knowledge necessary for understanding solar racing strategy within the WSC context. It then outlines the work that was done during the course of this thesis, firstly on modelling the mechanical and electrical systems of the solar car, and secondly on implementing a redesigned version of the current integrator device. Lastly, this document presents a log book of the 2011 WSC from my perspective as the race strategist, and outlines valuable conclusions and suggestions for further work.

Chapter 2

Background

The purpose of this chapter is to introduce the reader the World Solar Challenge race context, the *Sunswift IV* solar car and the structure of a typical solar racing strategy system. The principles of solar racing strategy, and the related practical considerations are then presented. Finally, using case studies, the constraints and common shortcomings of past solar racing strategy systems are discussed.

2.1 The WSC race format

All solar cars which compete in the WSC must abide by the event rules, which put hard boundary conditions on race strategy [World Solar Challenge, 2010a].

At the beginning of the race, the cars leave the starting line two minutes apart, in an order which is determined by the car's performance in a timed qualifying lap. The cars start the race with a full battery pack, but after this point it may only be recharged from the car's solar cells.

The cars are allowed to race between 8am and 5pm each day, after which time the teams have a 10 minute time buffer to pull off the highway and stop for the night. The cars must also stop for 30 minutes at 9 designated control points along the route (shown in Figure 1.1). The teams are allowed to continue charging their battery packs from the array when the car is stationary.

When on the road, the solar car must be escorted by a lead vehicle (which must be within 500 metres of the solar car), and a chase vehicle directly behind (which must be within 3 seconds of the solar car). All fleet vehicles must abide by the public speed limits and road rules. The winners of the race are determined by the order in which the cars cross the finish line.

2.2 *Sunswift IV* description

Sunswift IV, shown in Figure 2.1, was designed and built by the UNSW Solar Racing Team during 2008-09 to compete in the 2009 WSC. The solar car was designed to be competitive within the rules of the race, with efficiency and reliability foremost in mind, as mentioned earlier. During the last year, the Team improved *Sunswift IV* in preparation for racing once again in the 2011 WSC.

Figure 2.1: *Sunswift IV* on the Stuart Highway

2.2.1 Mechanical system

Sunswift IV is a three-wheeled vehicle, with two steerable wheels at the front and a single drive wheel at the back. The car's brushless DC motor is embedded into the rear drive wheel, which saves weight and energy by eliminating the need for a mechanical drive train.

The car's body is a carbon fibre monocoque[1] chassis, which forms both the supporting frame for the car's systems and and its outer aerodynamic shape. This overall wing shape is common to all modern competitive solar cars, which have converged on this efficient design over more than two decades of racing. The top half of the body (from here on referred to as the *top shell*) is not mechanically structural, but gives support and contour to *Sunswift IV*'s $6\,\mathrm{m}^2$ of silicon solar panels (from here on referred to as the *array*). The top shell can be removed from the bottom chassis and tilted at an optimal angle towards the sun when the car is stationary (Figure 2.2).

[1]A monocoque chassis is a vehicle construction method which uses the outer shape of the vehicle as its structural support. No internal structural frame is employed.

The driver compartment is positioned directly in front of the rear wheel and holds one person. The car is controlled with two hand-operated throttles mounted to the steering wheel: one for acceleration and one for regenerative braking. A foot-operated mechanical braking system is also in place for emergency use.

Figure 2.2: The *Sunswift IV* topshell tilted towards the sunrise to catch the day's first sun rays.

2.2.2 Electrical system

The car carries a 2 kg lithium-ion battery pack, which can store approximately 5 kWh of energy: enough to power *Sunswift IV* at 80 km/h for about 5 hours. Three maximum-power-point-trackers (MPPTs) ensure that solar energy collected by the array is stored efficiently in the battery. The solar array is rated to deliver 1.3 kW of power in 1-sun conditions[2]. A *Wavesculptor 20* motor controller manages energy flow between the battery and the motor. This controller, designed especially for solar cars by power engineering company Tritium, has inbuilt cruise control and regenerative braking functionality. The interaction of these power components within the solar car is described by Figure 2.3

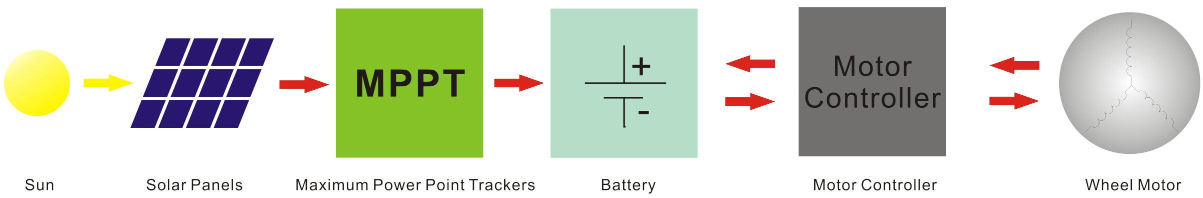

Figure 2.3: Power system of a solar car. Arrows indicate the permitted flow of energy between components. *Image courtesy of Irving Tjiprowarsono.*

[2] 1-sun conditions are equivalent to a solar insolation of 1 kW/m^2 hitting the ground.

In addition to the power system described above, *Sunswift IV* has a network of sensors (from here on referred to as the *telemetry system*) which collect data on the car's operation, such as battery voltage, motor power and array power. The telemetry system sensors communicate using a *controller area network* (CAN) bus. The collected data is transmitted wirelessly by to a laptop in the chase vehicle behind the solar car, where it can be analysed. This is achieved using an onboard embedded computer, the *Sunswift IVy Obsever Node* (SION) [Tjiptowarsono, 2011].The purpose of the telemetry system is to provide accurate and reliable feedback on the operation of the car, which is vital for running a comprehensive race strategy.

2.3 A typical strategy system

The aim of solar car race strategy is to minimise the total time taken to travel the 3000 km of the WSC. This ideal strategy can be described by a speed profile along the race route, which the solar car must follow to optimise its energy collection and expenditure, arriving in the shortest possible time. The role of a solar car strategy system is to calculate this optimal speed profile.

Figure 2.4 outlines the main components of a typical solar car strategy system and their interactions.

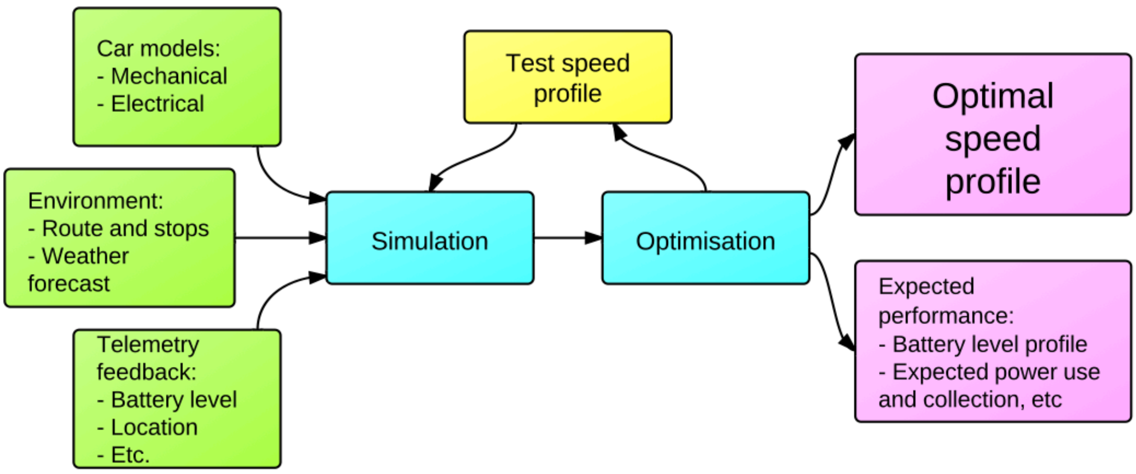

Figure 2.4: A typical solar car strategy system.

Firstly, before it can attempt to calculate an optimum strategy, the system must be provided with sufficient information about the solar car and the race environment, which includes:

- a model of the car's mechanical and electrical systems, which allows us to estimate how much power the car uses and collects, given a particular speed and conditions;

- the route that the solar car must follow, including the road conditions (such as variation in gradient) and any constraints placed by the race format and traffic regulations (discussed in Section 2.1 above);
- the likely weather conditions along the route, which will affect the solar car's power consumption and collection;
- the current starting conditions, which include the car's location along the route and its remaining battery charge.

This information can then be fed into the simulation and optimisation software, which lies at the heart of the system. This part of the strategy system can be implemented in many different ways, and can vary greatly in complexity and capability.

The role of the simulator is to calculate the performance of the solar car along a particular route (or part there of), given the models described above and a *test speed profile* as input. The optimisation software must then assess whether the result of the simulation was satisfactory, and adjust the test speed profile if there is room for improvement. The cycle is repeated until an optimal speed profile is found for the current conditions.

An effective strategy system should also output the expected performance of the solar car for the given optimal speed profile, which at minimum includes the expected battery level, power consumption and power collection along the route. As discussed later in Section 2.5.4, this performance profile, along with telemetry feedback, allows us to ensure that the solar car is successfully tracking a calculated strategy.

2.4 Optimal local strategies

The solar car race optimisation problem has been approached mathematically in numerous publications, most notably by Peter Pudney in his 2000 PhD thesis.

Pudney used optimal control methods to analyse the energy consumption of a solar car system in a variety of constrained scenarios. He arrived a set of techniques which optimise solar car performance over particular local road and weather phenomena.

This section outlines some of these techniques, as well as relevant contributions from other sources.

2.4.1 Power consumption model

Pudney's analysis, as well many other solar racing texts such as *The Speed of Light* [Roche et al., 1997] and *The Leading Edge* [Tamai, 1999] all use slight variations of the same equation as a starting point for modelling a solar car's power consumption (Equation 2.1 and Figure 2.5). This approach splits the power required to drive the car at a given speed into three independent components: gradient losses, rolling resistance and aerodynamic drag.

$$P = P_{gradient} + P_{rolling} + P_{aero}$$
$$= \frac{1}{\eta}[mgv(\sin\theta + C_{rr}) + \frac{1}{2}C_d A\rho(v - v_w)^3] \tag{2.1}$$

where:
P is the instantaneous power required (W)
η is the drive train efficiency (%)
m is the mass of the car (kg)
g is gravitational acceleration (m/s^2)
v is the instantaneous speed of the car (m/s)
v_w is the windspeed in the direction of of the car's travel (m/s)
θ is the road gradient (degrees)
C_{rr} is the coefficient of rolling resistance (dimensionless)
C_dA is the aerodynamic drag area (m^2)
ρ is the density of air (kg/m^3)

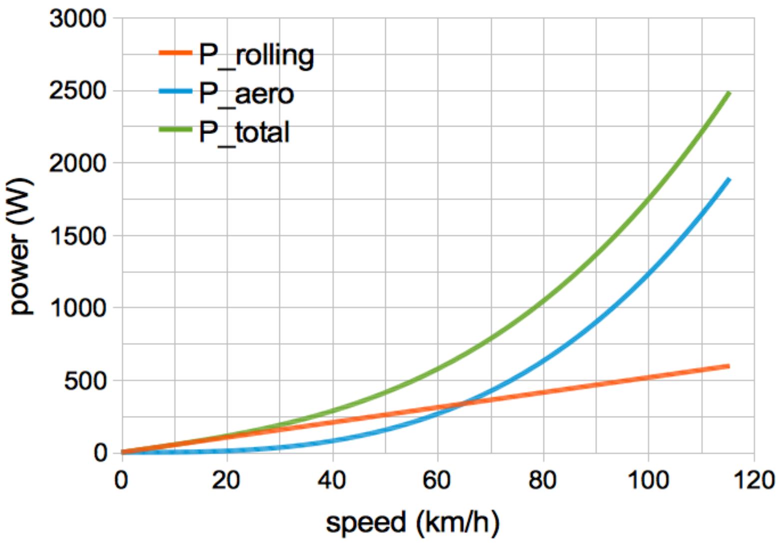

Figure 2.5: Estimated power consumption *vs* speed for *Sunswift IV* on a flat road with no wind, using Equation 2.1. The mass of the car and driver is taken as 237 kg, the drive train is assumed perfectly efficient, and we use $C_{rr} = 0.008$ and $C_dA = 0.09\,\text{m}^2$, which are estimated values derived from various *Sunswift IV* performance data.

2.4.2 Flat terrain

Using the power consumption model described by Equation 2.1, we can calculate the optimum strategy for a solar car in the simplest scenario: travelling on a flat road with no wind.

It is trivial to demonstrate that the most efficient way for a solar car to cover a stretch of flat road is to hold one constant speed [Pudney, 2000]. Referring to Figure 2.5 and Table 2.1 we can see that to travel at a steady 80 km/h for one hour, *Sunswift IV* would use just over 1043 Wh of energy. If the car were, instead, to travel at 60 km/h for the first 30 minutes and then 100 km/h for the remaining 30, it would average the same speed but use $288 + 873 = 1161$ Wh, or 11% more energy.

Speed (km/h)	Instantaneous power (W)	Energy used in 30 min (Wh)
60	575	288
80	1043	522
100	1746	873

Table 2.1: Estimated power consumption of *Sunswift IV* at different speeds on a flat road with no wind, with reference to Figure 2.5.

Although this is quite an aggressive example, it highlights the importance of choosing a constant speed to cover relatively flat terrain. Any large deviations from the average will lead to unnecessary energy loss. Over the WSC route, which has relatively shallow gradients, the solar car speed should only depart from a set cruise speed in the case of significant changes in terrain or weather.

2.4.3 Undulating terrain

When a solar car travelling at a given cruise speed is faced with an approaching steep incline or decline, Pudney suggests that the optimum strategy is the *hill anticipation* method. This method takes into account the inefficiency of the solar car's drivetrain.

When approaching an incline, the solar car should increase its speed, reaching a maximum at the foot of the hill (Figure 2.6(a)). It should then disengage cruise control, and allow its speed to drop over the course of the climb. When the car is once again on level ground, it should accelerate back to its previous cruising speed.

A similar method should be used with a steep decline, except that the car should decrease its speed in anticipation of the gradient. This approach can be extended to calculate the optimal speed profile over a hill or a valley (Figure 2.6(b)).

This scenario was also analysed in a paper on solar car strategy in the Sunrayce event[3] by [Daniels and Kumar, 1997], reaching the same conclusions. The method for calculating the precise *hill anticipation* speed profile required is presented in a book on efficient train control by Howlett & Pudney (1995).

[3]The Sunrayce event, now known as the North American Solar Challeng, is an endurace competition for solar vehicles similar to the WSC, held biennially in North America [Daniels and Kumar, 1997].

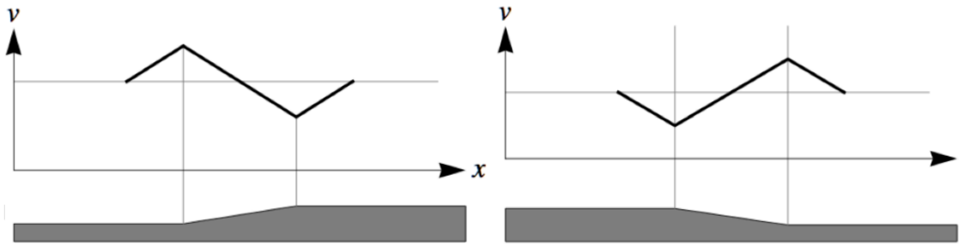

(a)

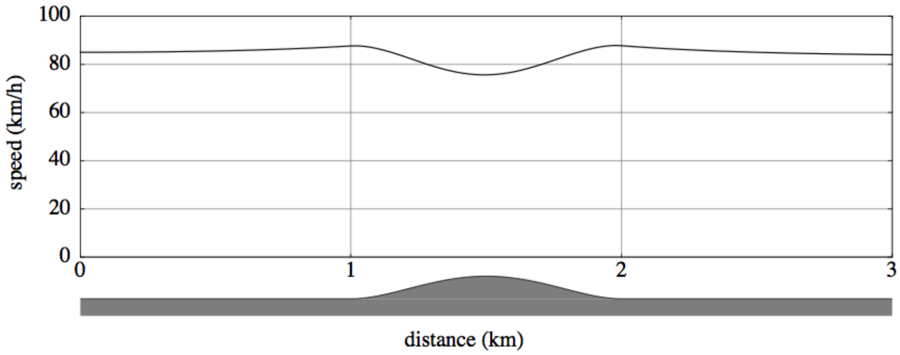

(b)

Figure 2.6: Optimal solar car strategy over undulating terrain. The speed *vs* distance plots in (a) describe the *hill anticipation* approach for steep declines and inclines. (b) describes the optimal speed profile over a hill. *Images courtesy of Peter Pudney [Pudney, 2000].*

2.4.4 Variable insolation

Another important problem which solar strategy needs address is variable insolation along the race route.

Pudney's work suggests that the optimal strategy for dealing with patches of cloud in an otherwise sunny environment is the *sun chasing* technique. This method involves increasing the solar car speed when under the cloud, and slowing down during good weather. This way, the solar car spends less time in unfavourable conditions, and collects more energy during sunny periods [Pudney, 2000].

Shimizu *et al.* from the Honda solar racing team reached a similar conclusion in their 1996 paper on solar car cruising strategy [Shimizu et al., 1998]. Figure 2.7 shows the example they used to analyse the problem. In their scenario, the solar car has to travel from 8 am to 5 pm, covering a flat route which has alternating patches of sunny and cloudy weather every 20 km. They evaluate three different strategies under these conditions:

A. Constant battery power, which involves changing the solar car speed to account for both temporal and spatial variation in insolation throughout the day.

B. Constant motor power, i.e. constant speed on a flat road with no wind.

C. Two set motor powers (speeds), one used for sunny patches and one for cloudy patches.

The two values were determined using optimisation methods not outlined in detail in their paper.

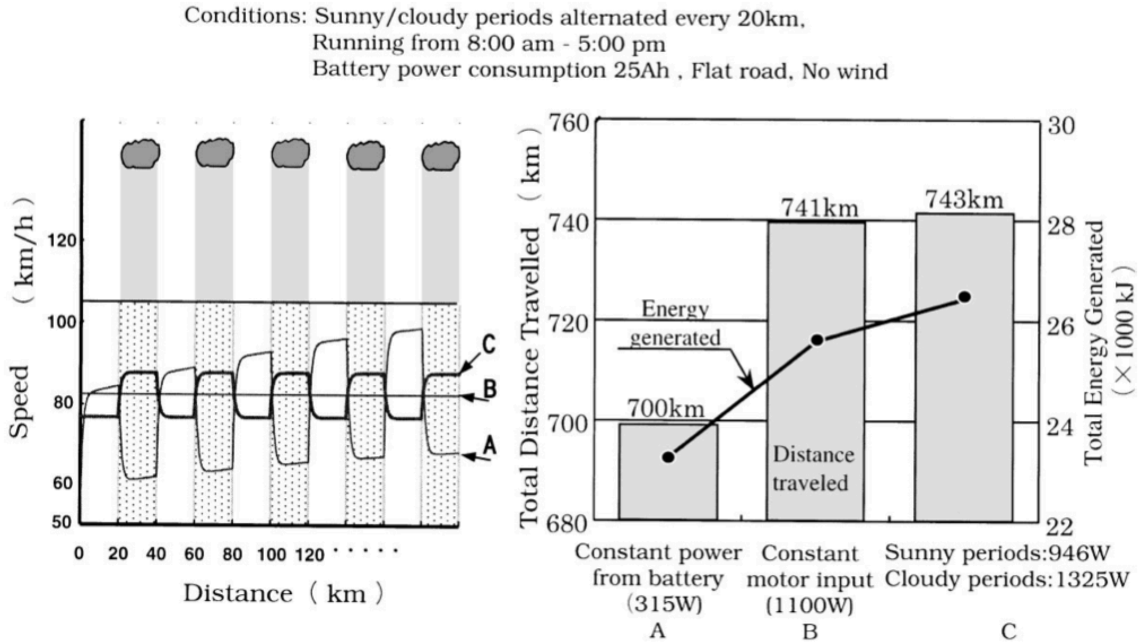

Figure 2.7: Performance of different strategies in spatially variable insolation conditions. *Image courtesy of Shimizu et al. [Shimizu et al., 1998]*

They find that the optimal technique is given by strategy C, which involves travelling slower in the sunny patches and faster under clouds. Using this technique, the solar car travels further during their defined race day, and collects more total solar energy. The difference in performance between this strategy and a constant speed strategy is, however, very small.

This difference is so small, perhaps, that slightly different model parameters and calculations could lead to a different result. Contradicting both Pudney and Shimizu *et al.*, the analysis of the same problem by [Daniels and Kumar, 1997] concludes that a *slower*, not faster speed in cloudy conditions is optimal.

It is likely that Daniels & Kumar reached this different conclusion as a result of using a more pessimistic battery efficiency model than both Pudney and Shimizu. Their calculations, therefore, show that it is more favourable to use solar energy immediately as it is collected, instead of storing in an inefficient battery.

The only conclusion that can really be drawn from these differing analyses is that the optimal technique is highly dependant on the models used and on the specific formulation of the problem. While changing the solar car speed in response to variable insolation can lead to a slightly better result than a constant speed strategy, the optimal solution must be determined individually for each situation.

2.5 Practical strategy considerations

While the techniques discussed in the previous section give good insight into the optimal behaviour of a solar car, the real challenge of the race is to use this knowledge in practice. A given optimal strategy only holds in the ideal world inside the strategy software. Outside the simulation, the real behaviour of the car will always deviate from a calculated strategy due to uncertain and changing weather, model inaccuracies and unexpected incidents. To run an effective strategy during the WSC, we need to understand and compensate for these deviations.

2.5.1 Weather uncertainty

Weather is the biggest uncertainty faced by solar racing teams during the WSC, and introduces an inevitable element of luck into the competition. There are several reasons why weather can cause the car's performance to deviate from a planned strategy:

- The weather forecast cannot be completely accurate, especially when predicting weather more than 1-2 days into the future. Even the top teams take at least 4 days to complete the WSC route, meaning that at the starting line, there is always some uncertainty about the weather along the rest of the route.
- The solar car's performance may be affected by local weather phenomena that is not covered by weather forecasts, such as low-lying smoke clouds resulting from bush fires, which affected most entrants in the 2011 WSC.
- The effect of some weather phenomena on the solar car's performance is difficult to model correctly. For example, the effects of crosswinds on the aerodynamic drag of a solar car are difficult to quantify, even though wind forecast information may be readily available. Similarly, the effect of thin cloud layers and diffuse insolation on the performance of the solar array is difficult to model, but has significant effects on the energy that can be collected.

The top solar racing teams usually have access to very detailed forecasts along the route, and some even employ a team meteorologist to help mediate weather uncertainties. It is also common for teams to send a weather scouting car 100-200 km ahead of the main fleet during the race, which carries some basic weather instruments and can report the local weather conditions.

Nevertheless, this is a variable that cannot be completely controlled. An effective strategy system needs to be able to detect and adapt to unexpected changes in weather.

2.5.2 Limited model accuracy

Weather uncertainties aside, the optimal calculated strategy can only be valid for the car that was described by the mechanical and electrical models used as inputs into the simulator. If these models don't correspond to the real power consumption and collection of the solar car on the road, following the calculated speed profile will lead to unwanted results. For example, if the solar car actually uses more power than is estimated by the model, and keeps running at the calculated speed, the battery will be flattened prematurely. The team will then to have to slow down significantly, or even stop, until some charge can be collected. As was discussed in Section 2.4.2, such variations in speed waste valuable energy.

However, while it is undeniably important to have a good model of the car, there are practical limitations to just how accurate this model can be. As is the case with any complex system that interacts with the outside world, we must make approximations in our models in order to keep calculations manageable. Furthermore, due to both time and resource constraints, it is often difficult to obtain and verify accurate model parameters in time for the race.

2.5.3 Unexpected interruptions

In addition, during the course of the WSC, a solar car that is following a set strategy is likely to face some unplanned interruptions. This could be an unexpected technical problem, such a flat tyre, or a malfunctioning component in the car. It could also be something beyond the team's control, such as roadworks along the route.

Any such event will cause the solar car to deviate from its planned strategy, and it is important for the strategy system to be able to adjust quickly to this change.

2.5.4 Importance of telemetry feedback

In order to detect any of the deviations discussed above, and to monitor how well a solar car is tracking its expected performance, we need to have sufficient feedback about the car's operation. This is provided by the telemetry system, as mentioned previously in Section 2.2.2.

At minimum, the telemetry system should give us information about the car's location on the route, the level of charge remaining in its battery and its instantaneous power consumption and collection. The strategist can compare these to the expected values calculated by the strategy system, and make adjustments to compensate for any discrepancies. For example, if the car is consistently using more power than is expected, the mechanical model provided to the simulator is likely to be too optimistic. Upon noticing this deviation, the strategist can adjust the model parameters, and recalculate a more realistic strategy. Similar adjustments can be made if the car's solar power collection is different to what is expected, or if an emergency stop is required and the team falls behind its calculated position along the route.

Telemetry feedback also extremely useful for detecting faults in the solar car's operation. For example, if the power used by the car increases suddenly, it is a tell-tale sign that a component in the system has started to malfunction. Telemetry feedback allows many on-road faults that are otherwise hard to notice to be detected and fixed quickly.

Battery state of charge tracking

One of the most useful pieces of information that we can get about the car's operation is its current battery level or *state of charge* (BSOC). While other feedback helps to verify the instantaneous performance of the car, the BSOC gives us an assessment of how well we have been following the planned strategy as a whole. A solar car's BSOC is equivalent to the fuel gauge of a conventional car: it is the cumulative sum of the energy that has been put in and taken out of the battery, and allows us to estimate how long we can expect to be able to keep driving.

Peter Pudney suggests that a practical way of following a calculated strategy is to drive the car to the predicted BSOC profile, instead of the calculated speed profile. If the car is performing better than strategy expects, we will drive faster than the optimal speed profile in order to stay on the predicted BSOC curve. Conversely, if the car performs worse, we will slow down. This way, we can ensure that we don't empty the battery pack before the end of the race, despite inaccuracies in the strategy models and weather forecasts [Pudney, 2000].

This method was used successfully by Pudney in the 1999 WSC with the *Aurora 101* solar car. He tracked the BSOC curve of the car when possible, and recalculated strategy whenever it deviated too far from the expected battery level [Pudney, 2000].

Local weather measurements

In addition to feedback on the performance of the car, it is also useful to have a means of checking the local weather conditions against the predicted forecast used in strategy simulations. At minimum, a windspeed sensor and a reference solar cell (to measure insolation) mounted to the roof of the chase vehicle can be extremely helpful. As mentioned previously, scouting the local weather ahead of the fleet can also be very valuable, especially in variable conditions.

2.6 Previous solar car strategy systems

Different solar racing teams have approached the strategy problem in a variety of ways over the years, with different resources and implementations of the simulation and optimisation software. However, many of the shortcomings and practical race experiences seem to be common to all teams. Researching and evaluating these different systems helped me to identify where I should focus my efforts in improving the *Sunswift IV* strategy system for the 2011 race.

2.6.1 UNSW *Sunswift IV*, 2009 WSC

The *Sunswift* strategy system has been in development since the team's first races in the late 1990s. By the 2009 WSC it had all the elements of a sophisticated strategy system, although it still lacked in some common areas, such as model accuracy.

Simulation and Optimisation

The simulation software used by the strategy system was written by members of the team, and implemented in the C programming language to maximise computation speed. The software could run a fast simulation of the race, and easily allowed the addition or adjustments of models. This same system was used again in 2011 without major modifications, and is discussed in detail in Appendix A.

During the 2009 WSC, the *Sunswift* strategy system only used constant-speed optimisation. The system found the fastest constant speed at which the car could run the course of the race without flattening its battery pack prematurely, but performed no improvements beyond that. Any major variations in weather which made a constant speed strategy suboptimal had to be accounted for by the strategist by "gut-feeling".

Some attempts were made at implementing the optimisation techniques discussed in Section 2.4, but these did not work successfully before the start of the race.

Car model

In the lead up to the 2009 WSC, the Team ran behind in the *Sunswift IV* build schedule and did not have time to obtain good models of the car's systems.

The power consumption of the car was modelled using Equation 2.1. At the starting line, these mechanical model parameters were mostly estimates, derived from theoretical values calculated during the aerodynamic design process [Beeves and Doig, 2009]. They were adjusted during testing just before and over the course of the race, but were never highly accurate.

The battery model was based mostly on data-sheet values, as well as some data collected during last-minute tests before and during the race. No comprehensive battery discharge or load tests had been completed, so the strategist did not have accurate knowledge of the battery capacity and impedance.

The solar array was modelled simply as a flat rectangular surface covered in solar cells. A lot of work had been put into a more comprehensive model that accounts for the curvy shape of the array, but this was not completed in time. The array model did not account for the fact that during stationary charges at the end of each day and at control stops, the array surface would be tilted directly towards the sun. This lead to some possible underestimation of the collected solar power [Snowdon, October 2011].

The MPPTs, motor and motor controller were modelled as having constant efficiencies, which were mostly estimated from data-sheet values.

Environment model

Since 1996, the Team has had access to a very detailed road survey of the WSC route. The survey was carried out by students from the UNSW School of Geomatic Engineering, and contains information on the gradient, altitude and road roughness for over 90,000 GPS locations along the 3000km route (a resolution of about 30m) [Wong and Rizos, 1996]. This is one of the Team's most prized strategy resources. During the 2009 WSC, the gradient information from the survey was used in the strategy's power consumption model. The altitude and road roughness data was not used.

Thanks to a sponsorship agreement with meteorology company Weatherzone, the Team had access to detailed and frequent weather forecasts during the race. These forecasts included wind, insolation, precipitation and cloud cover information. Some of Weatherzone's meteorologists also accompanied the team on the WSC trip, helping to interpret the data and forecasts. Due to time constraints before the race, only the insolation data was explicitly included in the strategy simulations.

Telemetry feedback

As discussed in Section 2.2.2, *Sunswift IV* was equipped with a comprehensive telemetry system, which provided feedback about the car's operation. This information included the car's GPS speed and location, the motor and array power and the battery voltage, current and state of charge. Other information about the operation of individual components in the car was also transmitted.

Unfortunately, the telemetry device responsible for measuring the battery state of charge (BSOC), the *current integrator*, showed very unreliable behaviour for the entire duration of the race. As a result, the Team strategist could not trust its readings and did not have accurate knowledge of the battery level. Problems with this device are discussed in detail in Chapter 4.

Another significant problem was the unreliability of the GPS device in the solar car. Aside from recording the car's location, the GPS is responsible for synchronising the timestamps of all the other telemetry nodes in the network. During the 2009 race, the GPS device did not correctly update the timestamp on its data messages, therefore making it difficult to sync the car's location to other telemetry data. To overcome this issue, the strategist used a back up GPS device installed in the chase vehicle.

Additionally, the Team did not have good means of measuring the local weather. A windspeed sensor was mounted to the chase vehicle, but did not work as expected. The reference solar cell used for measuring local insolation was not well calibrated.

The telemetry data was displayed using the *Scanalysis* user interface, developed by David Snowdon, longtime *Sunswift* member and lead strategist during the 2009 race. This software was used again in the 2011 WSC, and is discussed in detail in Appendix A.

General comments and performance

Although the Team performed very well in the 2009 WSC (1st place in the Silicon Challenge class and 4th overall), the strategy system had been problematic in many ways.

Due to inaccuracies in the mechanical model, the car's power use deviated from the the calculated strategy. The strategist compensated for this by tracking an approximate BSOC curve instead of the calculated speed profile (as discussed in Section 2.5.4).

However, since the Team did not have an accurate battery model and the current integrator device was unreliable, the strategist only had an approximate idea of the car's battery level. The Team was forced to run a slightly conservative strategy to make sure that the battery pack was not emptied before the end of the race.

It is estimated that at the finish line, the battery had about 10% of its total charge remaining. Optimally, the battery should be completely empty at the end of the race, signifying that the solar car had run at the fastest speed possible to get there. That last 10% of charge, if used over the course of the race, could have let the car travel about 0.5 km/hr faster on average, finishing the race 15 minutes earlier[4]. Although not a huge loss, it could in some cases make the difference between overtaking a competing solar car or staying behind them. In the 2011 WSC, for example, there were many close finishes between teams, some as close as 7 minutes [World Solar Challenge, 2011].

2.6.2 University of Twente *SolUTra*, 2005 WSC

Because of the competitive nature of the WSC event, it is not easy to obtain detailed information about the strategy systems of other teams. The University of Twente's strategy for their *SolUTra* solar car in the 2005 race is one of the exceptions to this rule. The race strategy software system was the subject of a Master's Thesis by a student on the team, and contains valuable information about its implementation and performance [Mocking, 2006].

Simulation and optimisation

The team used custom strategy software called PALLAS, which was developed by Mocking for his thesis. PALLAS uses the commercial dynamic modelling software 20-sim for

[4]This calculation assumes that 10% of the battery pack is about 500 Wh. If used evenly over the course of the race, which takes about 40 hours to complete, this is equivalent to 12.5 W of added instantaneous power. *Sunswift IV*'s average speed in the 2009 WSC was 76.0 km/hr, and with 12.5 W extra power it could travel at an average of about 76.5 km/hr (using the power consumption curves from Figure 2.5).

simulation and optimisation calculations, and MATLAB for its graphical user interface, database connections and telemetry monitoring. During the 2005 WSC, the PALLAS program was capable of determining constant speed strategies only [Mocking, 2006].

Car model

The team did not collect enough data to accurately model the car's power consumption. They used the Arnhem Highway just outside of Darwin for testing before the race, but couldn't get enough data to build good models, and did not have time to test them.

The team used a low-current charging method to obtain the equilibrium voltage curve for their battery (see Section 3.2.1). They did not have time, however, to collect data over the entire operating range of the battery, and estimated the rest of the curve by extrapolation. [Mocking, 2006]

Environment model

The team obtained gradient information along the route using GPS measurements. However, they did not include this data in their strategy system because its inclusion significantly slowed down the performance of their simulation calculations.

The team received weather forecast information, although not a lot of detail is available regarding these. Only insolation was considered in their calculations, wind information was not included. [Mocking, 2006]

Telemetry feedback

SolUTra appears to have had quite a comprehensive telemetry system. The BSOC was estimated with current and voltage measurements, but Mocking states that they did not have a lot of confidence in these measurements. The current measurement (and integration) device was not calibrated properly and the battery level reading it gave was subject to significant drift.

The Twente team also had at least two significant telemetry drop-outs, where the chase vehicle lost all telemetry communication with the solar car for several hours. [Mocking, 2006]

General comments and performance

From Twente's race log, it seems that it was very difficult for the strategist to check whether or not the solar car was correctly tracking its calculated strategy. This was due to the fact that their BSOC readings could not be accurately verified until each morning, when the the battery level could be inferred from a reading of the 'at-rest' battery voltage. The fact that their strategy simulations did not include the effect of gradient also made it significantly harder to check whether the instantaneous power consumption of the car was as expected.

The team arrived in 9th place, after suffering significant tyre-related problems during the race. The PALLAS program proved to be useful in helping to monitor the car's performance, and in making decisions regarding strategy. [Mocking, 2006]

2.6.3 Summary

In addition to the two systems described in detail above, there is a lot of other valuable information available regarding the strategy systems of other solar racing teams. We can make some very useful observations about solar strategy systems in general, and their typical shortcomings:

- The lack of accurate models of the car's systems seems to be an overwhelmingly common problem. Without good models of the car's power consumption and collection, limited benefit can be extracted from a sophisticated strategy system.
- Telemetry failure is also a common problem. Both University of Twente in the 2005 WSC and the Aurora team (previously known as the Ford Vehicle Association team) in the 1999 WSC had problems with their wireless data link [Mocking, 2006] [Pudney, 2000]. Without reliable telemetry feedback, it is difficult to verify whether a calculated strategy is being followed correctly.
- Battery state of charge tracking (Section 2.5.4) provides a very effective realistic way of following a set strategy [Pudney, 2000]. However, in order to do this the car's battery level must be measured accurately and reliably. Both *Sunswift* in 2009 and the University of Twente in 2005 failed to do this well.
- Despite the fact that small advantages can be gained by employing the strategy techniques discussed in Section 2.4, most teams that I could find information about only employed constant-speed optimisation during the race. This strategy was typically calculated each morning, and a set speed established for the day.
- Little can be gained from a highly optimised strategy if the solar car is unable to follow it correctly. Generally, model inaccuracy, uncertainty in weather and telemetry problems introduce a larger error than the benefit that can be gained from local strategy optimisations. It is therefore important to ensure that the system is working well in these areas first, before concentrating efforts on sophisticated optimisation.

Chapter 3

Modelling the car

As discussed in Section 2.6.3, one of the most common downfalls of solar car strategy systems is their lack of accurate models of the car. It is sometimes tempting for solar car teams to put a lot of effort into solving the mathematical race optimisation problem, while forgoing rigorous verification of the car models.

In the most common scenario, solar car teams simply run out of time to build and test these models before the start of the race, and have to resort to estimated or theoretical model parameters. The typical time budget for each WSC is 18 months: from the release of race regulations to the starting line. There is a lot of work to complete in a relatively short period, especially if designing a new car, as the the UNSW Solar Racing Team did in 2009 with *Sunswift IV*.

In 2009, the car was not a complete working system until about 3 weeks before the race, and the first high speed testing was done on the drive up to Darwin. Due to this, not much data was collected on the mechanical operation of the car, and none of it in well-controlled conditions. The build of the electrical system ran similarly close to the deadline: by the time the array and battery pack were built, there was no time left to properly characterise them. As a result, during the race the strategist used largely estimated car model parameters, and had to continually manually compensate for the car's deviation from calculated strategy, as discussed in Section 2.6.1.

One of the main goals of the thesis was to rectify this problem before the 2011 WSC. This chapter outlines the methods which were used to collect data on the mechanical and electrical systems of the car and to build it into models which the strategy software, described in Appendix A, can use. The performance of the resulting models, and difficulties in achieving high accuracy are discussed. For completeness, this chapter also outlines the theory behind modelling some system components which were outside the scope of this thesis, but are recommended for future work.

3.1 Mechanical model

Almost all the energy used by the solar car during operation is spent in overcoming mechanical drag forces. A good model of the car's mechanical properties is required in order for the simulation software, discussed in Appendix A to accurately predict the car's energy expenditure along the race.

As discussed in Section 2.4.1, the power consumed by a solar car to travel at a given speed can be approximated by the following equation:

$$P = \frac{1}{\eta}[mgv(\sin\theta + C_{rr}) + \frac{1}{2}C_d A\rho(v - v_w)^3] \tag{3.1}$$

If we simplify this scenario further, assuming a perfectly efficient drive train, and that the car is travelling on a flat road with negligible wind, our equation reduces to the following:

$$P = mgC_{rr}v + \frac{1}{2}C_d A\rho v^3 = P_{rolling} + P_{aero} \tag{3.2}$$

The power required now consists of two components: the power used to overcome rolling resistance ($P_{rolling} = mgC_{rr}v$) and the power required to overcome aerodynamic drag ($P_{aero} = \frac{1}{2}C_d A\rho v^3$). If we know the car mass m, gravity g and air density ρ, we now only need to determine the coefficients C_{rr} and $C_d A$ to obtain our mechanical model.

3.1.1 Estimated model parameters

A good starting point for determining the model parameters is to look at their theoretical or estimated values. In the worst case scenario, where no empirical data can be collected, these values can be used as an approximate final model and adjusted manually using feedback from *Scanalysis*, the *Sunswift* strategy graphical user interface.

Aerodynamic drag area

A theoretical $C_d A$ of 0.095 m^2 was determined for *Sunswift IV* using *computational fluid dynamics* (CFD) methods during the aerodynamic design of the car [Beeves and Doig, 2009]. The value was calculated assuming the following ideal conditions:

- the car's angle of attack is adjusted to the model used in the CFD calculations
- the surface of the car is perfectly smooth
- the yaw (direction of airflow with respect to the centreline of the car) is zero

While the angle of attack of the car of can be adjusted to a good approximation by changing the ride height of the front suspension [McLaren, 2009], the latter two conditions are impossible to replicate perfectly. The real surface of the car has small unavoidable imperfections, most notably in the connection between the top shell and bottom chassis, and at the interfaces between solar panels in the top shell. The direction of airflow during the race is also never consistently parallel to the body of the car, due to variable wind direction. Fortunately, since the car typically travels at speeds much higher than wind velocity (upwards of 80km/h), the yaw angle is generally very small. In light of these non-idealities, we expect the real drag coefficient to be somewhat higher than the theoretical value.

Rolling resistance

It is difficult to make a good estimate of the real rolling resistance of a solar car without collecting data on their performance on the car. The C_{rr} coefficient is determined by the interaction of the solar car tyres with the road surface. This interaction is sensitive to many factors specific to the car's set up, such as the tyre pressure, the condition of the tread on the tyre (i.e. whether the tyre is new or not) and how the tyre is fitted to the wheel [Roche et al., 1997].

Although tyre manufacturers typically publish a nominal rolling resistance in their product data sheets, these values usually cannot be used for much more than qualitative comparisons. In the 1996 WSC, for example, many teams used the Michelin *Radial* tyres, whose quoted[1] C_{rr} value was given as 0.0025. The Honda team, however, measured a C_{rr} = 0.0036 during on-track testing of the solar car, while the Biel team measured a value of 0.00562 for their car, both using the same Michelin tyres [Roche et al., 1997]. It is evident that the rolling resistance coefficient is largely dependent on set-up, and must be determined experimentally.

3.1.2 Experimentally determining model parameters

To obtain accurate empirical values for the mechanical model parameters, data on the solar car's dynamic behaviour needs to be collected in a controlled environment. Equation 3.2 can then be fitted to the data using various regression methods to extract the required coefficients.

Wind tunnel testing is used in the automotive industry, and by some of the bigger budget solar car teams, to verify the aerodynamic properties of their solar car. The *Sunswift* team, however, did not have access to a big enough wind tunnel facility to use this method. We resorted to other practical alternatives.

Two effective ways of determining the mechanical parameters are through roll-down testing and constant speed testing. These methods, which are detailed in this section, can essentially be performed on any long, flat stretch of road during calm weather conditions.

[1] This rolling resistance was measured by Michelin using a 3m diameter drum with a texture surface, which was spinning at 5km/h.

Roll-down testing

Roll-down testing involves accelerating the solar car to a given speed, then disengaging the accelerator and letting the car coast down to a stop. During the coast down process, the velocity of the car is sampled [Rutman, 2007]. The data collected is a function of the car's velocity with respect to time, $v = f(t)$, which can be related to Equation 3.2 as follows:

$$P = mgC_{rr}v + \frac{1}{2}C_d A\rho v^3 \tag{3.3}$$

then diving through by v, we get:

$$P/v = F = ma = m\frac{dv}{dt} = mgC_{rr} + \frac{1}{2}C_d A\rho v^2 \tag{3.4}$$

now dividing through by m, we get:

$$\frac{dv}{dt} = C_{rr}g + \frac{1}{2}\rho C_d A v^2/m \tag{3.5}$$

where F is the retarding force in Newtons, a is the decelertion in m/s^2, and the remaining variables are defined as per Section 2.4.1.

In light of the above analysis, once we have the $v = f(t)$ relationship from the roll down test, we can extract the C_{rr} and C_dA values from the first derivative of this data, dv/dt.

This method assumes that the road gradient is negligible during the course of the test, and that there is no wind. It is therefore important to perform the test on an extremely flat stretch of road, in very calm conditions.

Constant speed testing

Constant speed testing is very similar in principle to the roll-down method. However, instead of letting the car coast down to a stop, the car is accelerated to a nominal speed and cruise control is engaged. The average power used by the car to hold this speed over some distance is recorded. The test is then repeated for a range of different speeds, to build a dataset of power used as a function of speed [Optimum G Vehicle Dynamics Solutions]. This data can be related directly to Equation 3.2:

$$P = f(v) = mgC_{rr}v + \frac{1}{2}C_d A\rho v^3$$

As with the roll-down method, close to ideal testing conditions are necessary. However, although more time consuming, this method should in theory be less error prone. Each point in the data set is the result of average power and velocity measured over a distance, which should filter out a lot of the inevitable road and wind disturbances.

3.1.3 Pre-race roll-down tests

During August and September 2011, with help from other members of the Team, I carried out multiple roll down tests on *Sunswift IV*. These tests were done on a 900 m stretch of flat road at the Sydney International Regatta Centre grounds in Penrith (Figure 3.1). The speed *vs* time data was taken from readings made by the solar car's motor controller. The intention was to supplement the roll down data with constant power tests, but we could not complete these due to time constraints.

Figure 3.1: Roll-down testing *Sunswift IV* at the Sydney International Regatta Centre

Testing considerations

This stretch of road was not long enough to perform continuous coasting from high speed (over 75 km/h) to a standstill, but after searching widely for appropriate tracks, it was the best Sydney-based venue which we could get access to. This issue was overcome by performing the roll down test in three parts, at high, medium and low speed, and then combining the data to form a continuous curve.

Two students from the UNSW School of Surveying completed a comprehensive survey of the track in September, and verified that it was almost perfectly flat (no gradient variations above 2% [Li and Qiu, 23rd September 2011]. Wind velocity during the tests was not measured, but I partly compensated for its effect by performing each roll down test in both directions along the track and taking the average of the results. We were fortunate to have a few calm testing days with very small differences between roll down tests in the two directions. In retrospect, wind data would have been very useful for determining how much to trust a particular dataset - without it, this was based on qualitative observations.

We had access to the Regatta Centre track almost every week over two months in the lead up to the 2011 WSC. During a good day of testing, 3 or 4 full sets of roll-down data in both directions could be collected. Unfortunately, during this period the mechanical team was still performing work on the car in preparation for the race, and the car's set up changed slightly from week to week. In particular, a new mechanical brake system was installed, and was adjusted over a period of weeks to find a position in which the brake pads did not rub against the wheel of the car. As a result, during many of the early days of testing, the roll down test results were skewed by extra resistance from rubbing brake pads.

Furthermore, the Team had not yet finalised many details regarding the car's race configuration. Race tyres, which are discussed in detail in Section 3.1.3 had not yet been chosen. The suspension had not been adjusted to ensure the desired angle of attack. As a result, even the best data collected during these roll-down tests was not a true reflection of how the car would operate during the race.

In retrospect, I should have made sure to control many of these variables, such as ensuring that the angle of attack of the car was set up correctly from the start. Many others, however, were a consequence of the Team as a whole not establishing a hard deadline to finalise major technical work in advance and leave time for testing. **The importance of having the car ready to be tested well before the race cannot be emphasised enough.**

Extracting model parameters

From the many roll-down datasets collected, we selected and analysed the ones which were taken in the least windy conditions, with minimal brake rubbing. The dark blue line in Figure 3.2 is a plot of a full *Sunswift IV* roll-down curve from a test performed on the 30th of August 2011, which I considered to be a good dataset.

Due to the fact that the curve of the data was quite shallow, it was difficult to directly fit a polynomial curve to it to extract the parameters. It was necessary to perform quite a lot of data manipulation in order to get a result that made sense.

Before beginning any analysis, the data was smoothed using a moving average filter to remove noise caused by road roughness and any other disturbances during the roll-down test.

The first step in the analysis was then to extract the coefficient of rolling resistance from the data. Referring back to Figure 2.5, we can see that below a speed of about 18km/h (about 5 m/s), the effect of aerodynamic drag on the power consumption is negligible. Therefore, if we take the gradient dv/dt of the data below this speed, it should depend only on the C_{rr} coefficient:

$$\frac{dv}{dt} = C_{rr} g \tag{3.6}$$

where gravity $g = -9.8\,\mathrm{m/s^2}$.

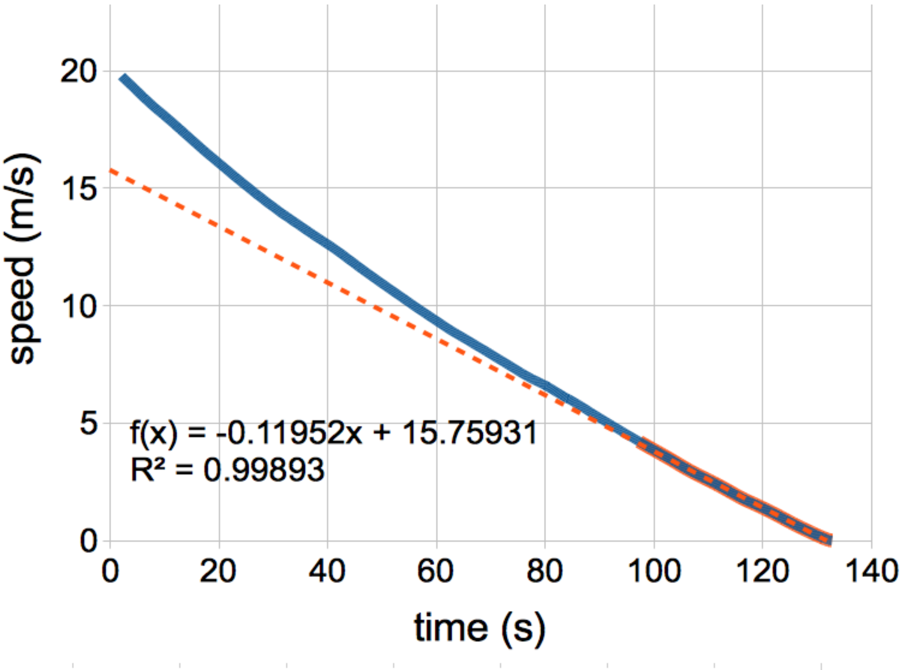

Figure 3.2: High-speed roll down curve for *Sunswift IV*.

As shown in Figure 3.2 (red line), such treatment gives a gradient of $-0.120\,\mathrm{kg/s^2}$, and a C_{rr} of $(-0.120)/(-9.8) = 0.012\pm20\%$. The error is estimated from the spread of C_{rr} results from different datasets.

Using this C_{rr} value, the C_dA parameter could then be extracted. A little more algebraic manipulation was required:

$$\frac{dv}{dt} = a = C_{rr}g + \frac{1}{2}\rho C_d A v^2/m \tag{3.7}$$

then, rearranging to solve for C_dA:

$$C_dA = \frac{2m(a - C_{rr}g)}{\rho v^2} \tag{3.8}$$

For each datapoint in the set, I calculated the acceleration a between consecutive data points, using the car mass $m = 230\,\mathrm{kg}$ and air density $\rho = 1.225\,\mathrm{kg/m^3}$. I then calculated the C_dA value for each point, using the above definition [Diasinos and Doig, 22nd August 2011]. As expected, the results had a very wide spread, and the values calculated for low speeds deviated so much that they were basically meaningless (Figure 3.3).

To extract a C_dA value that was close to the theoretical value of $0.095\,\mathrm{m^2}$, I had to reduce the dataset. I limited the range of data points to those at high speeds (above about 30km/h), and removed significant outliers. The resulting C_dA values are plotted in Figure 3.4. Taking the average of these, I calculated a C_dA of $0.1 \pm 0.07\,\mathrm{m^2}$. The large error comes from the standard deviation of the values.

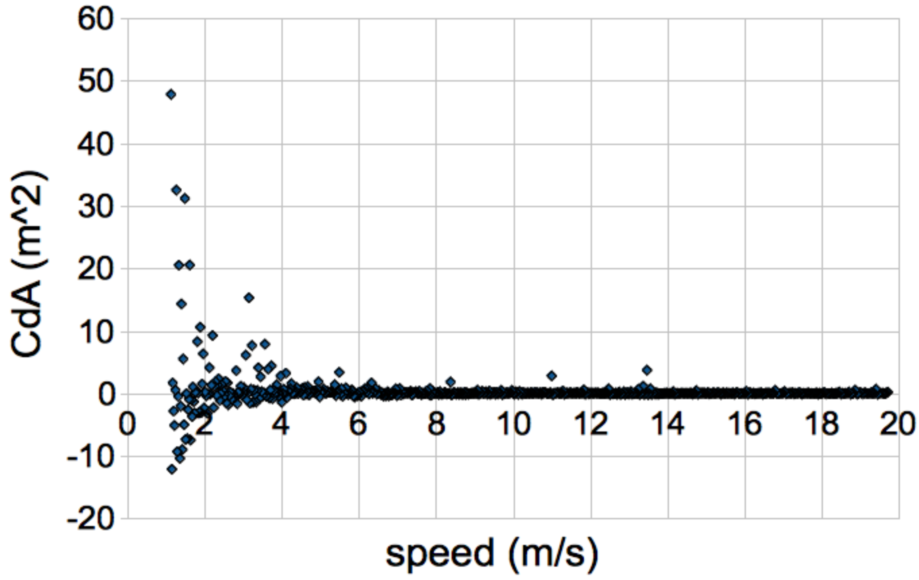

Figure 3.3: Scatter plot of calculated C_dA values before data reduction.

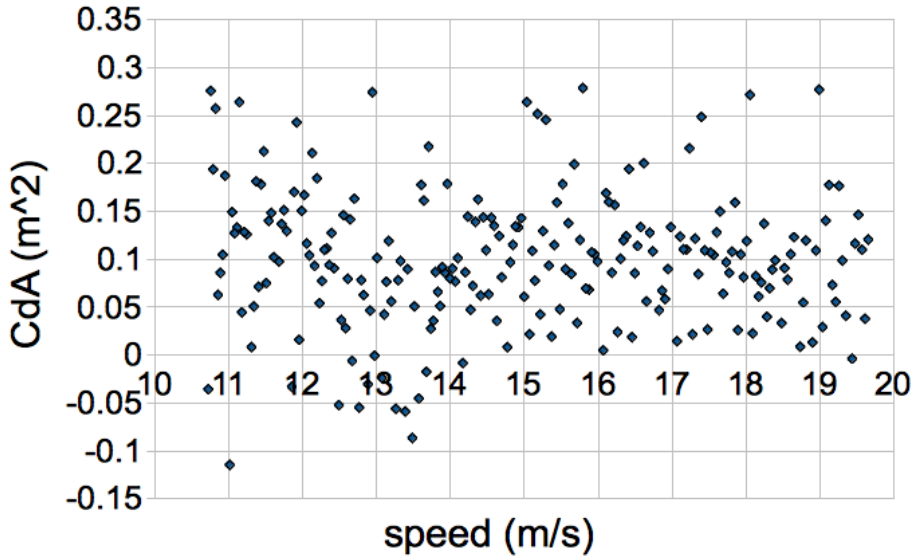

Figure 3.4: Scatter plot of calculated C_dA values after data reduction.

Many repeated roll-down tests in the same conditions and car set-up are required to ascertain whether or not the model parameters can be extracted with repeatability. Unfortunately, by the time the car was in the correct configuration before the race, there was not enough time to do this.

Tyre comparisons

Although the results of these pre-race roll down tests would not be used for the race models without further adjustments, they were useful in helping the Team make informed decisions regarding our choice of race tyres.

Rolling resistance was one of the biggest points of concern in the lead up to the 2011 WSC. New changes to the regulations for this race prohibited all teams from using Dunlop *Solarmax* tyres, a favourite during the 2009 race, as they were deemed inadequate for highway use [World Solar Challenge, 2010b]. *Sunswift IV* used these tyres during the 2009 WSC, and the car's suspension was designed around them. The *Solarmax* tyres were one of the few on the market made specifically for solar cars. They had very shallow tread and a low rolling resistance, estimated from testing performed in early 2011 to be about 0.0085 [Snowdon, January 2011].

This rule change left the Team with few good options. In the lead up to the race the Team had two choices: the Sava *MC2* moped tyres, and the harder-to-obtain Michelin *Radial X* tyres, which were another custom product for solar cars. The *MC2*s had very deep tread and were expected to push the rolling resistance up significantly, above 0.01. The *Radial X*s, however, carried the major disadvantage that they were too large to fit in *Sunswift IV*'s front suspension without a redesign. They were, however, expected to have a rolling resistance similar to the *Solarmax* tyres used in 2009.

Since we could not obtain any Michelin tyres in time for testing, we decided to compare the Dunlop *Solarmax* and the Sava *MC2* tyres to get an approximate idea of their relative performance. Due to some technical problems on the day that we did the testing, we had very little time on the track. We performed low-speed roll-down tests with *MC2* tyres on all three of *Sunswift IV*'s wheels, but did not have time to change all of them to the *Solarmax*s. We changed the rear tyre only to the Dunlop *Solarmax*, and performed some roll-down tests in this configuration, in the hope that a significant difference could be seen.

Figures 3.5(a) and Figure 3.5(b) show the results of two roll-down tests in opposite directions along the track, using *MC2* tyres on all wheels. Figure 3.6(a) and Figure 3.6(b) show the roll-down test results, using a *Solarmax* tyre on the rear wheel, and *MC2* tyres on the front two wheels. Even with only one tyre changed, the difference in performance is significant.

Using the same method as before, we assume that for these low speeds aerodynamic drag is negligible, and so $dv/dt = C_{rr}g$. Due to a visible non-linear component in some of the data at very low speeds (below 2 m/s), I took the gradient for estimating the C_{rr} for this data consistently from the linear region between 2 and 4 m/s. This was okay since this data was only to be used for qualitative comparisons, and we did not expect to extract an accurate coefficient value.

I took the average of the gradients in both directions for each tyre set-up. The rolling resistances were calculated as $C_{rr} = 0.012 \pm 20\%$ for the full *MC2* tyre configuration and a $C_{rr} = 0.009 \pm 15\%$ for the case where the rear tyre was changed to a *Solarmax*.

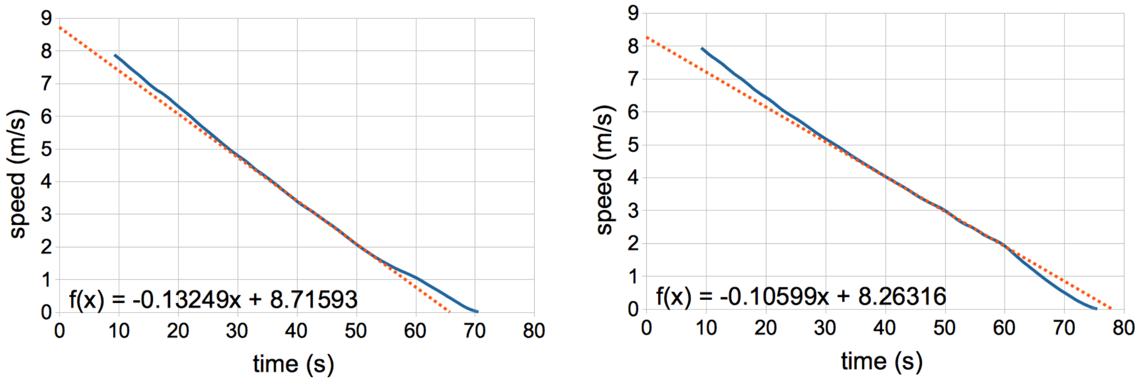

Figure 3.5: *Sunswift IV* low speed roll-down tests in both directions along the Regatta Centre track, using Sava *MC2* tyres.

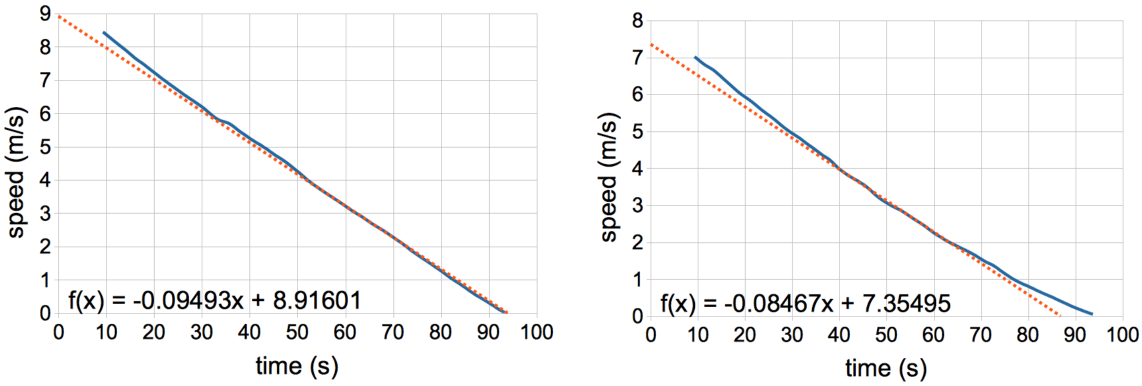

(a) (b)

Figure 3.6: *Sunswift IV* low speed roll-down tests in both directions along the Regatta Centre track, using a Dunlop *Solarmax* tyre on the rear wheel and Sava *MC2* tyres on the front two wheels.

This difference was much higher than we expected. Using the weight distribution of the car on the three wheels (Appendix F), I extrapolated that changing all the tyres to the Dunlop *Solarmax* would push the C_{rr} value down to 0.003. This was unrealistically low, when compared to the 0.0085 value estimated from testing the car using a full set of *Solarmax* tyres during January 2011.

I had to conclude that the data we collected that day was not particularly reliable. It had been a windy day, which made results in opposite directions significantly different, and probably contributed to the non-linear effects at low speeds. Due to time constraints, we also hadn't collected enough datasets to be confident in the results.

The trend, however, was undeniable - the *MC2* tyres were significantly worse than the *Solarmax* (and hence, we assumed, the Michelin *Radial X*). As a result of these roll-down tests the Team seriously considered obtaining *Radial X* tyres for the race, and rebuilding our suspension to fit them. After a lot of discussion, however, we realised that this was simply unrealistic, as the race was barely one month away at the time. We decided that we would use the *MC2* tyres and suffer the cost of the higher rolling resistance.

This was not, however, the end of the Team's tyre-related adventures. In Darwin, just days before the start of the WSC, we had the unexpected opportunity to buy 12 Schwalbe *Energizer S* tyres from the Stanford Solar Racing Team. These tyres were a new product, and not many teams had known that they were available. They fit our suspension and, according to testing performed by the Stanford team, they were expected to have a similar rolling resistance to the *Solarmax* and *Radial X* tyres. The main disadvantage was that we would be using these tyres on the race with no prior highway testing, without knowing their C_{rr} value at the starting line. In the heat of the competition, the Team decided to take this risk, and to determine the rolling resistance of the tyres during the course of the race.

3.1.4 Determining model parameters from telemetry feedback

Despite our attempts to collect mechanical model data before the start of the race, the Team did not have verified mechanical model parameters at the starting line. I had a good idea of what the approximate C_{rr} and C_dA values should be, but especially due to the last-minute change of tyres, I knew that they would have to be adjusted.

Fortunately, some improvements made to the *Scanalysis* strategy interface (see Appendix A) by David Snowdon in the weeks before the race provided a very elegant method of making these adjustments using telemetry feedback. Snowdon added a new functionality to the graphical user interface program, which plots the real-time power consumption of the car on the same graph as what is expected by the strategy simulations. If the car is following the set strategy speed, and a consistent difference is observed (Figure 3.7(a)), this is a sign that the mechanical model parameters are incorrect. Using my knowledge of the models and of the car, I could then adjust these parameters in the strategy model definitions, and calculate a new strategy which should be more correct. I would repeat this process a few times to get the two power curves to match well. Figure 3.7(b) shows a plot of the two power curves after a successful model adjustment.

We found this method to be incredibly effective during the 2011 WSC. Once we adjusted the parameters at the beginning of a given race day, the power curves would not deviate much at all as long as we had no technical problems and held an approximately constant speed (within 5 km/h or so).

The model parameters that we determined this way were only valid over a limited range of speeds. However, since the solar car speed doesn't typically vary much over the course of a WSC race, this method is still valid in practice. The 2011 WSC was an exception to this, with unusually variable weather forcing the Team to change our speed significantly from day to day. Even then, by re-adjusting the model parameters a few times a day, I could make good use of the strategy simulations.

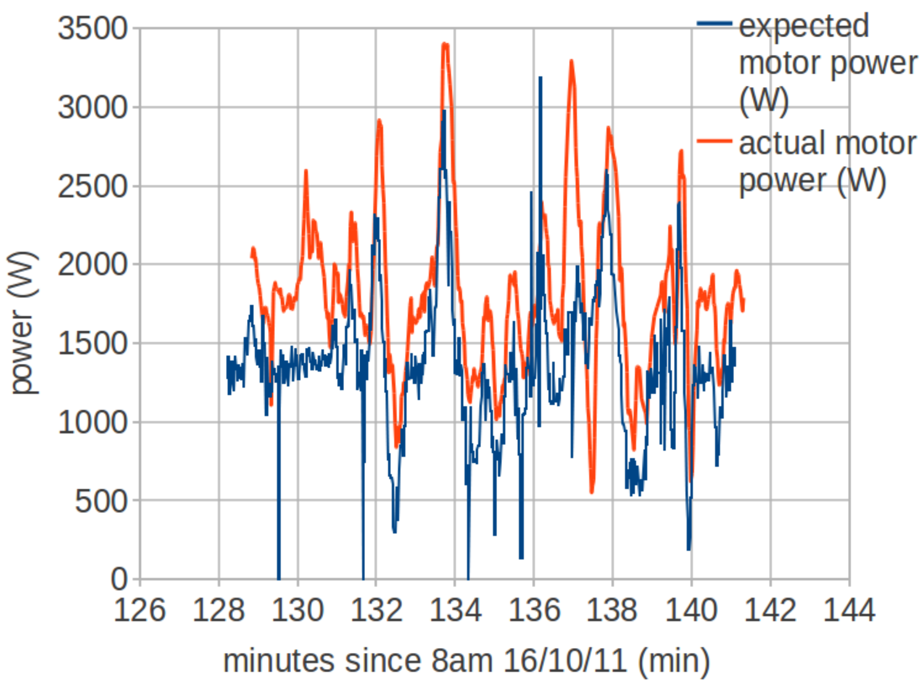

(a)

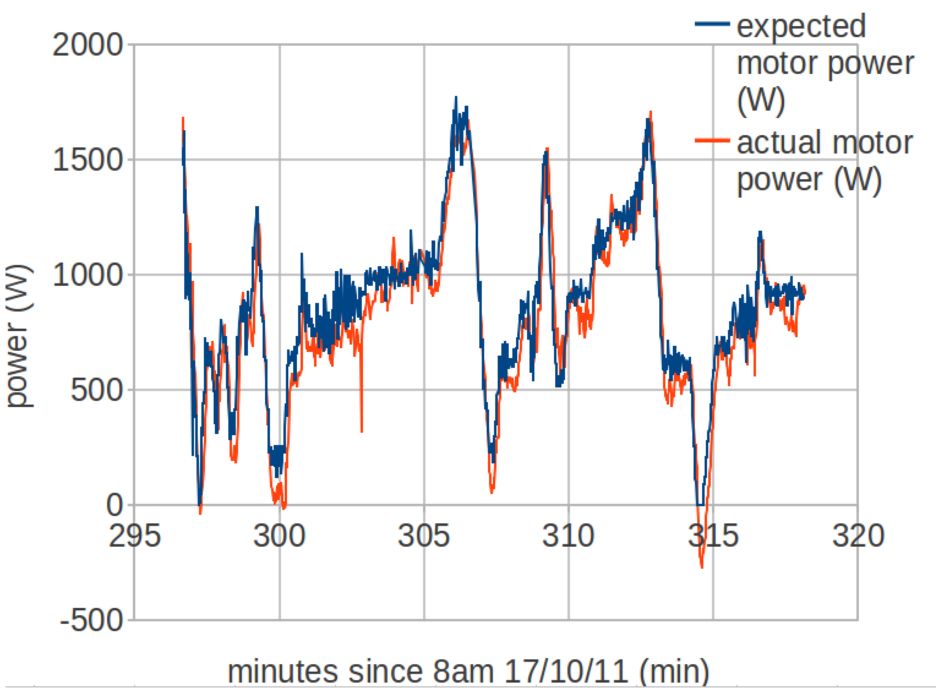

(b)

Figure 3.7: Real power use vs expected power use of *Sunswift IV*, for a strategy using inaccurate model parameters (3.7(a)) and accurate model parameters (3.7(b)).

3.1.5 Conclusions

In light of the work discussed in this section, some important conclusions can be drawn about the realistic value of different methods of determining the mechanical model parameters.

- In order for roll-down and constant speed testing to give accurate results that can be used in the strategy system, the car needs to be in its race-ready configuration when the tests are performed, weeks in advance. Until all systems have been finalised, these techniques can only really be used to make qualitative comparisons.
- The qualitative comparisons that these methods provide can be very useful for evaluating the effect of potential changes to the car, and should be used to this effect. While in our case the roll-down method was used to compare the performance of different tyres, it could (and should) also be used in perfecting several adjustments, such as
  - angle of attack
  - wheel adjustments
  - weight distribution
- Since the car usually travels within a small range of speeds during the race, a limited but effective mechanical model can be determined using feedback from the car's telemetry while on the road. Such a model may not be accurate over the full range of speeds, but may be the perfect solution for the strategy system, especially if last-minute changes to the car need to be accounted for.

3.2 Electrical model

The main purpose of the electrical model is to allow the strategy system to simulate how energy is collected and stored during the course of the race. The models of the battery and the array are therefore the most important. The effect of other components of the electrical system, which include the maximum-power-point-trackers (MPPTs) and the electric drive train can usually be reduced to a constant efficiency value.

During the course of this thesis, due to time constraints only the battery model was experimentally obtained before the start of the 2011 WSC. As discussed in Section 2.6.1, lack of a good battery model was a significant shortcoming during the 2009 WSC, and it was important to rectify this. This section presents the work completed on the battery, and suggested methods for modelling the other components, which are left as future work.

3.2.1 Battery

In order for the strategy system to successfully simulate the battery, we need to be able to model the following properties:

- the total energy storage capacity of the battery pack;

- the *equilibrium discharge curve*, or how the 'at-rest' battery voltage varies with respect to the battery state of charge;
- the internal impedance of the battery, or how the battery voltage changes with respect to battery current.

Knowledge of these properties is essential in order to correctly simulate battery discharge and predict the battery current that is required at any point in the race.

Capacity

The nominal capacity of the battery pack can be estimated from the manufacturer's datasheet for the battery cells. An empirical measurement, however, is always preferable. The optimum solar car race strategy implies using up all the stored battery energy by the end of the race, and to do this we need to know exactly where the bottom of the pack lies.

The capacity of a battery is measured in ampere-hours (Ah) and tells us how much electric charge it can store. We expect to be able to draw a current of 1A for 1 hour from a 1Ah battery before it is flat. *Sunswift IV*'s 2011 battery pack is made up of Panasonic *NCA18650A* lithium-ion secondary cells. Panasonic gives their nominal capacity is as 3.1Ah per cell in ideal operating conditions. Due to internal chemical properties of the cells, the capacity is expected to vary with temperature and discharge current, as shown in Figures 3.8(a) and 3.8(b) [Panasonic, 2010].

During the WSC, we do not expect temperature inside the solar car to deviate much from the range of 25-50°C. As can be seen in Figure 3.8(b), the effect of this temperature variation on capacity is negligible. Furthermore, from data collected during the 2009 WSC, we expect the average battery current to be less than 8A, which corresponds to a 620mA current per cell for the 2011 *Sunswift IV* battery configuration[2] [Favaloro, 2011]. Once again, with reference to Figure 3.8(a), we do not expect this to cause significant deviation from nominal capacity.

The fact that our battery is operated in very close to ideal conditions means that we can take its capacity to be constant, greatly simplifying the required battery model. David Favaloro, a *Sunswift* thesis student who was responsible for the design and build of the *Sunswift IV* 2011 battery pack, performed controlled low current charges and discharges of the battery to determine its capacity [Favaloro, 2011]. Using the data from his tests, I calculated the battery capacity to be approximately 39 Ah. This is discussed in more detail in the following section.

[2] The 2011 *Sunswift IV* battery pack is made up of 35 series modules, each containing 13 Panasonic *NCA18650A* 3.1 Ah lithium-ion cells in parallel. An 8 A current drawn out of the battery is shared by 13 cells in parallel, with each cell providing $13/8 = 620\,\mathrm{mA}$.

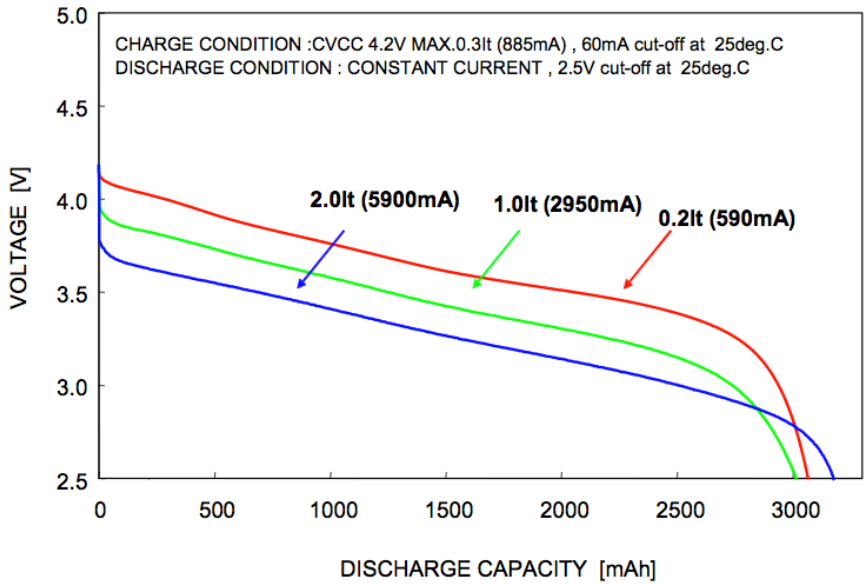

(a) Discharge at different currents

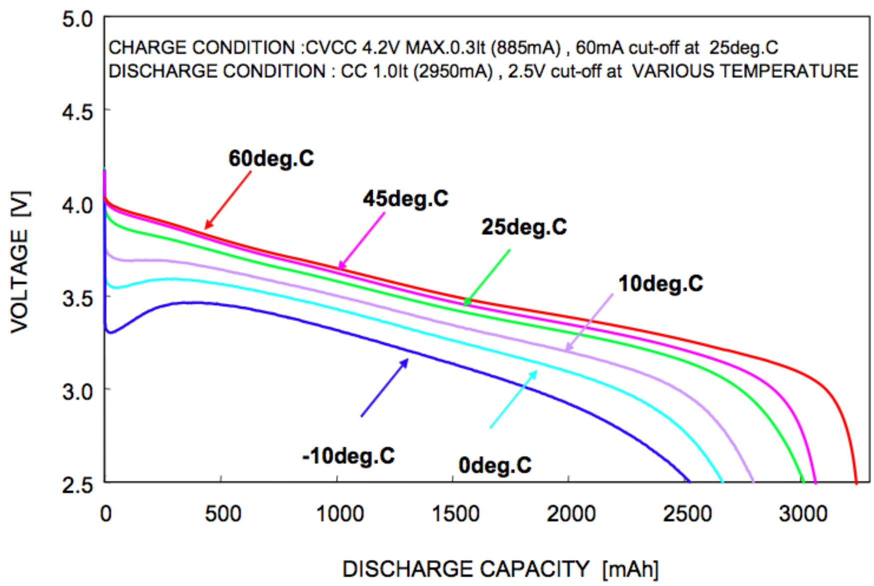

(b) Discharge at different temperatures

Figure 3.8: Discharge properties of Panasonic NCA18650A lithium-ion cells. *Images courtesy of Panasonic.*

Equilibrium discharge curve and internal resistance

At a given temperature, the open circuit voltage of a lithium-ion cell varies as a function of its state of charge, according to an intrinsic *equilibrium discharge curve*, $V_{eq} = f(S)$, where S is the state of charge of the battery. This curve gives the voltage that would be measured across a battery 'at rest', i.e. a battery that had experienced no recent current flow. [Electropaedia, 2005]

Due to the internal impedance of the battery, any current flow will cause the voltage to deviate from this curve. As a good first approximation, we can assume this impedance to be purely resistive [Pudney, 2000]. The voltage across the battery pack then varies as a function of battery state of charge with a current-dependant offset:

$$V = f(S) + IR_i \tag{3.9}$$

where:

V is the voltage in volts
S is the battery state of charge, in Ah
I is the current in amps
R_i is the internal resistance in Ω

We can determine the internal resistance R_i by comparing the raw voltage data from battery charges and discharges performed at different currents. For example, if we have the data for voltage as a function of the battery state of charge (BSOC) at two different currents, then for a given state of charge S:

$$V_1 = f(S) + I_1 R_i \tag{3.10}$$

$$V_2 = f(S) + I_2 R_i \tag{3.11}$$

Subtracting Equation 3.11 from Equation 3.10, we get:

$$V_1 - V_2 = (I_1 - I_2) R_i \tag{3.12}$$

$$R_i = \frac{(I_1 - I_2)}{(V_1 - V_2)} \tag{3.13}$$

This relationship can then be applied to data points for each value of the state of charge, and an average R_i calculated.

Ideally, when time permits, as many sets of data at different currents as possible should be collected, and regression used to find the best value of R_i. Before the start of the 2011 WSC, only two full datasets were available, which David Favaloro collected when testing the *Sunswift IV* battery pack.

With guidance from *Sunswift* mentor David Snowdon, in the days before the race, I used the method outlined above to calculate an internal resistance of $0.205\,\Omega$. Figure 3.9(a) shows the raw battery voltage measured during two discharges at different currents. Figure 3.9(b) shows the same two sets of data, after subtracting the IR_i offset from each point. The two datasets now line up, and the resulting curve is a good approximation for the 'at-rest' equilibrium curve of the battery, $V_{eq} = f(S)$.

I sampled the resulting equilibrium discharge curve at 40 discrete points, and included this information in the *Sunswift* strategy battery model as a look-up table. The simulation software interpolates between the values in the table, and adds the IR_i offset to find the voltage of the battery at a particular current and state of charge.

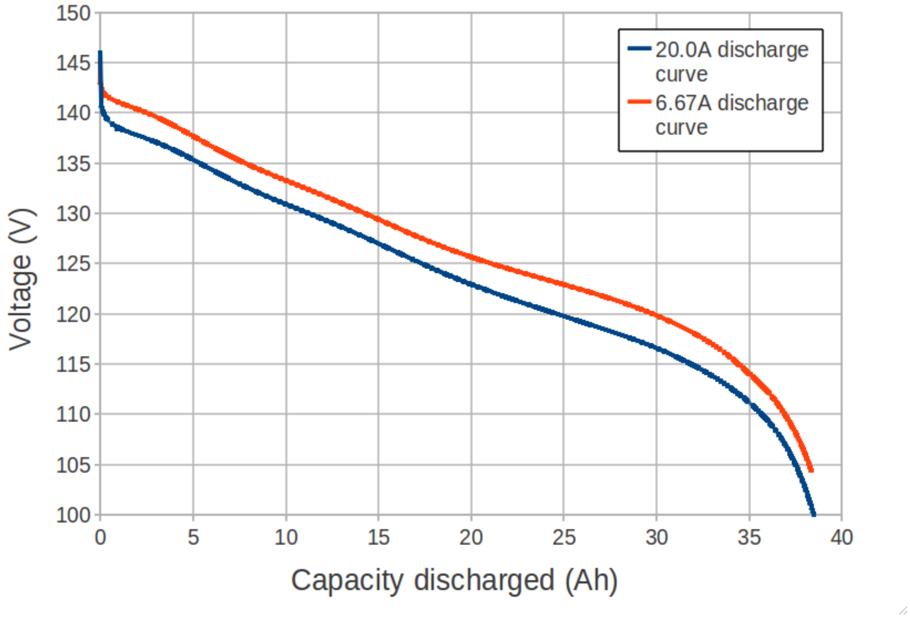

(a) Raw measurement of *Sunswift IV* battery voltage during a 6.67 A (red line) and 20 A (blue line) continuous discharge. *Data courtesy of David Favaloro.*

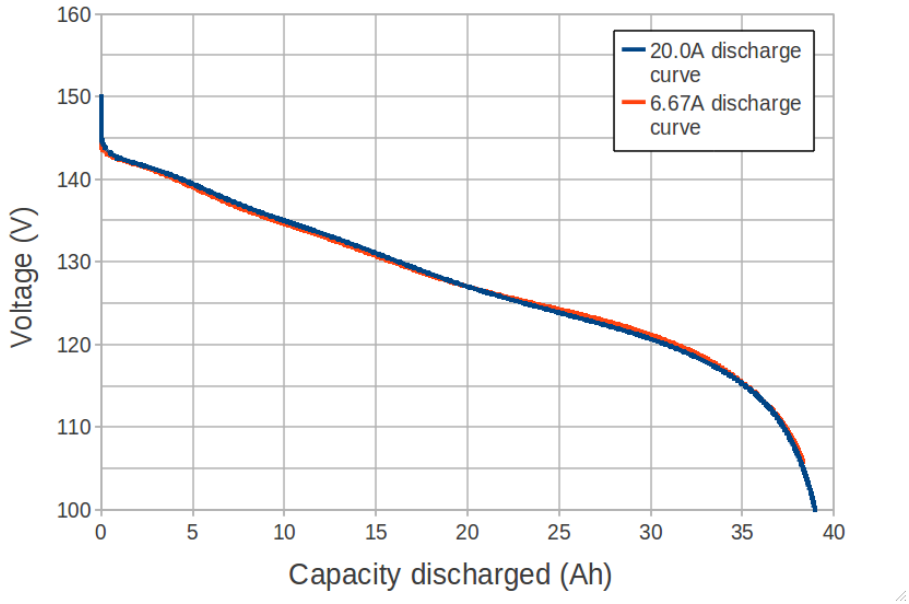

(b) Equilibrium discharge curves, obtained by using the raw data from (a) and adding offset of IR_i to each data point, where I is the discharge current and R_i=0.205 Ω. *Data courtesy of David Favaloro.*

Figure 3.9

From Figure 3.9(b) we can see that the battery reaches a voltage of 100 V after 39 Ah have been drawn from it. As 100 V was chosen as the minimum safe operating voltage for the battery, the operating capacity of the battery is 39 Ah.

Transient behaviour

Although the method outlined above gives a good approximation of battery behaviour for strategy purposes [Pudney, 2000], the internal battery impedance is not purely resistive. One simple model of the lithium-ion battery impedanceis the *parallel RC* equivalent circuit shown in Figure 3.10 [Debert et al., 2008]. This model accounts for the time dependency of the lithium-ion battery voltage, which is ignored in the purely resistive model.

The R_1 parameter in the equivalent circuit is just the internal resistance, and can be obtained in the same way as described before. The values of R_2 and C_1 can be determined by measuring how long the battery takes to return to its equilibrium voltage after being loaded with a high current.

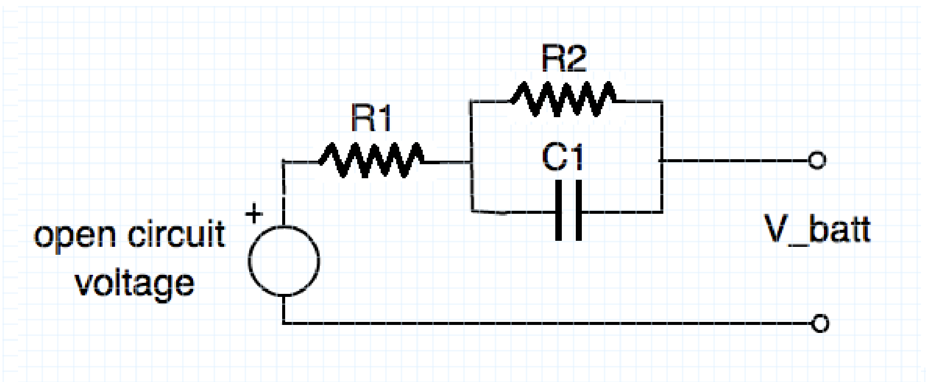

Figure 3.10: The parallel RC lithium-ion battery model [Debert et al., 2008].

A model of this form was not obtained before the 2011 WSC, but is recommended before the next WSC.

3.2.2 Array

As mentioned in Section 2.2.2, the *Sunswift IV* solar array consists of 6 m^2 of silicon solar panels contoured around the aerodynamic top shell of the car. These panels are divided into three series strings, each connected to one maximum-power-point-tracker (MPPT) as shown in Figure 3.11.

To accurately model the power collected by these solar cells under given insolation conditions, we need to take into consideration several of its properties:

- Due to the curvature of the top shell, each solar panel is oriented at a slightly different angle towards the sun, and receives a different amount of effective insolation.

- The maximum current in each of the three series strings is limited by the worst performing panel, i.e. the one least optimally oriented towards the sun.
- The canopy, which sits in the centre of the rear array string, shades the solar panels which immediately surround it when the sun is low above the horizon.
- The efficiency of the solar panels varies with temperature.

Accounting for all these factors amounts to a quite a complex model, and is of lower priority than having good mechanical and battery models. In the past, including during the 2011 WSC, the array model used by the strategy system approximated it as a flat plate covered with 6 m^2 of solar cells. Such a model is especially prone to error when the sun is a low elevation above the horizon. During 2009 and 2010, the *Sunswift* electrical team invested significant effort into implementing such a model, but this project was not completed in time for either the 2009 or 2011 WSC. It is highly recommended that this model should be completed before the next race.

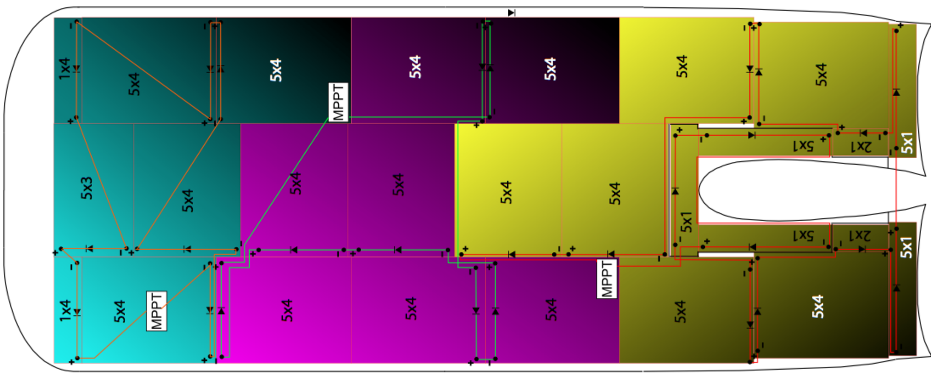

Figure 3.11: The *Sunswift IV* solar array, divided into three series strings of solar panels. *Image courtesy of Etienne Le Sueur.*

3.2.3 Maximum-power-point-trackers

The Drivetek maximum-point-point-trackers used in *Sunswift IV*, which manage the transfer of solar energy from the array to the battery, boast a maximum efficiency of 99%. This efficiency varies slightly with input power, as well as with the boost ratio between the solar panel voltage and the battery voltage [Biel School of Engineering and Architecture].

Within the set-up of the *Sunswift IV* power system, the MPPTs do not operate at their highest efficiency, but are not far from it. From these data sheet values, the approximate efficiency of *Sunswift IV*'s 3 MPPTs is 98% at an input power of 1 kW (see Appendix D for calculations). In the 2011 WSC, as well as in previous races, we have taken this efficiency to be constant.

As with all instruments in the solar car, it would be worthwhile to verify experimentally that the MPPTs follow the performance curves described in the data sheets. The UNSW School of Electrical engineering owns a power diode array, which can be used in conjunction with a high voltage power supply to simulate a solar panel. Using this set up for the input to the MPPTs, and the *Sunswift IV* battery pack as load, we can measure the efficiency of the MPPTs by varying the input power and measuring the output power.

Once these empirical efficiency plots are obtained, it may be worth considering including them in the strategy model, instead of a constant efficiency. This is, however, a low priority.

3.2.4 Motor and motor controller

Like the MPPTs, the wheel motor and *Wavesculptor 20* motor controller have previously been modelled as having constant efficiencies. Travelling at typical 80km/h, the *Sunswift IV* solar car consumes about 1050 W of power. According to data sheet values, the motor efficiency under these conditions is about 97.7%. The corresponding motor controller efficiency is about 97% (see Appendix E for calculations).

Once again, it would be useful to verify the performance of these components with real data, but this should not be prioritised over modelling the less efficient systems. It would also be interesting to compare the performance of the new motor built my members of the team in 2011 with the old motor which been used by the team for the last decade. Both motors use the same CSIRO magnet and winding kit (from which datasheet values for efficiency are taken) [Lovatt et al., 1997], but have different casing designs [McLaren and Bycroft, 2011].

3.2.5 Telemetry power

The telemetry system within the solar car continually uses about 30 W of power. While not a large number compared to motor and array powers, it is still a significant quantity and should be included in the strategy calculations. Over the course of an average 40 h race, the telemetry system will use 30 W * 40 h=1200 Wh, or almost a quarter of the battery pack.

Chapter 4

The Current Integrator

Arguably, the most important feedback that the strategist can get about the operation of the car is the battery state of charge (BSOC). As discussed in Section 2.5, even in the presence of inaccuracies in the strategy models, a strategy can be followed relatively successfully by tracking a calculated BSOC curve. Without knowledge of the battery level, however, the strategist can't check whether a calculated strategy is being followed correctly.

4.1 Methods of measuring BSOC

An approximate BSOC can be indirectly inferred from a reading of the battery voltage. The discharge curve for the lithium-ion cells used in the battery pack is fairly linear between 10% and 95% of battery capacity, as shown previously in Figure 3.9(b). This means that the open circuit battery voltage can be roughly matched to the battery's remaining capacity. The battery voltage is measured by multiple devices in the car (such as the motor controller and the MPPTs) and so the strategist is never left *completely* in the dark.

This method of BSOC estimation, however, is not very accurate when current is flowing in or out of the battery - which is generally the case whenever the solar car is in operation. As discussed in Section 3.2.1, the battery is not an ideal voltage source, and its voltage varies with changing current due to its internal impedance. Unless the battery is modelled very accurately to account for these transient changes, there will always be an error when estimating the BSOC of an operating solar car from a single voltage reading. Due to practical limitations in battery modelling (mentioned in Section 3.2.1), a more direct and accurate method of battery level estimation is required.

The current integrator device exists in the solar car the for precisely this purpose - to act as an accurate fuel gauge for the solar car's battery. As the name suggests, the device measures current flow in and out of the battery pack and integrates it over time. In this way, it keeps track of the total number of amp-hours of charge present in the battery (the

BSOC). As this method works by measuring current flow and not voltage, it does not suffer the same problems related to the battery's impedance.

A current integrator was built by the Team and used in the 2009 WSC, but was found to have some serious reliability issues. In order to have a reliable way of measuring the BSOC in the 2011 WSC, a hardware and software redesign of this device was deemed necessary. I began this work as part of a Taste of Research project in the summer of 2010/2011, during which I evaluated the previous device and completed the hardware redesign. The build, software design and testing processes for the device were completed during this thesis.

4.2 Evaluation of previous device

The previous current integrator device generally performed well in bench tests in Sydney [Au and Zhang, 2009] but did not perform as required during the 2009 WSC. In order to guide the design of the new version of the device, I analysed the previous design and its shortcomings. These are summarised in this section.

The device used an *MSP430F149* micro-controller as the heart of the system, which was the standard for all *Sunswift IV* telemetry nodes at the time of design. The device interfaced with the CAN bus of the car using the *MCP2515* controller area network (CAN) controller via a *s*erial peripheral interface (SPI). Current measurement was done by measuring the ground current of the battery (i.e. low-side measurement), using a precision resistive shunt and the MCP3909 ADC. Integration was implemented in software, and integrated current (BSOC) values were saved to the on-board flash of the micro-controller every 10 minutes [Au and Zhang, 2009].

The design had some known issues, which the Team knew could affect its performance, but did not have time to fix before the start of the 2009 race, as follows:

- The design did not follow many standard design procedures to protect its sensitive analogue measurement circuitry from electromagnetic noise and coupling. These include shielding of the analogue circuitry and use of analogue and digital power planes. When used in the solar car, the battery current readings from the device appeared to be very noisy. As a result, the strategist could not trust these readings to more than about ± 0.5 A, and had to assume that the integrated current value was similarly inaccurate as a result [Snowdon, November 2011].
- It was difficult to accurately calibrate the device, as the response of the analog-to-digital converter (ADC) used was not purely linear [Au and Zhang, 2009].
- The 10 minute store interval for the integrated value was too long[1].This time interval was chosen to ensure that the on-board flash memory of the micro-controller, which

[1]Assume the car is in cruise on a flat road, drawing a typical 8A current out of the battery. If the device lost power just before a store operation, the last saved integrated value could differ from the actual battery state of charge by up to (10 min * 8 A) = 0.13 Ah, or 3% of the full battery charge of 40 Ah.

is rated to 10,000 write-erase cycles, did not malfunction from overuse.

- The device connected to the battery using wires soldered directly to the PCB, both for the high current and high voltage connections. This was done to reduce the contact resistance. However, this meant that once the device was wired into the battery, removing it was a non-trivial task.

Additionally, the device failed to perform to the required standard in unexpected ways the Team could not accurately diagnose:

- During bench testing in Darwin, the current measurements appeared to be very sensitive to noise injected through the device power supplies. The current measurement fluctuated significantly ($\pm 100\,\text{mA}$) when the grounded chassis of the CAN bus connector to the device was disturbed in any way.

- Operating in the solar car during the race, the device appeared to corrupt its recorded integrated current value every 20 minutes or so. All in all, the strategist could not rely on the device, and had to use the battery voltage discharge profile to roughly estimate the BSOC [Snowdon, November 2011].

Unfortunately, the causes of both these problems were not found and could not be reproduced on the bench when the Team returned back to Sydney. The Team decided, however, that a redesign of the device was needed as a priority for the next WSC, with emphasis placed electromagnetic noise considerations and on fixing the other known issues discussed above.

4.3 Requirements

In light of the problems experienced in the 2009 race, the following requirements were used to guide the design of the new current integrator:

1. **Compatibility with *Sunswift IV*'s CAN bus:** The device must operate as a regular node on the CAN bus of the car. It must use the *scandal* protocol, developed by Dave Snowdon for his 2002 thesis and standard to all *Sunswift* telemetry nodes [Snowdon, 2002]. The CAN circuitry on the node must be isolated from high-voltage circuitry.

2. **Accurate current measurement:** The device must be able to measure current in the range of -50 to +50 A. These are the peak battery currents during maximum regenerative braking and acceleration. These limits are user-defined at the motor controller and chosen with consideration to the safe operation of the power train as well as expected power demands from the motor.

   For the majority of race time, however, the battery current is within the range of -20 to +20 amps. In order to minimise drift errors during integration, it is important to have highly accurate measurements in this range. An absolute error $<10\,\text{mA}$ in this range is acceptable. For measurements outside this range, which do not occur as often, an absolute error <100 mA is acceptable.

Current should be sampled at least 10 times a second.

3. **Accurate current integration and frequent back up:** The device must perform accurate integration of the current over time. The integrated current value must be saved regularly, and stored even when power to the device is disconnected. A store frequency of at least once every minute is desirable. This would limit the typical error due to a sudden shut down of the device to about 0.13 Ah (assuming a typical current of 8 A).

4. **Accurate voltage measurement:** The device must be able to measure the battery voltage, which ranges between 100-150 VDC, with an absolute error <100 mV for this range. The battery voltage and current can be used to calculate battery power, as an additional function of the device. Voltage should be sampled at least 10 times a second.

5. **Sound electrical connection to battery:** The device must be able to be securely and safely connected to the battery for both current and voltage measurements. It should be easily removable for maintenance.

6. **Good performance in electromagnetically noisy environment:** The device must be able to operate within the specified error margins inside the solar car as well as on the bench. As the solar car is an electrically noisy environment, the necessary precautions must be taken to ensure that the measurements are not affected by the noise[2].

7. **Low power consumption:** The total power used by the device should be no more than an average of 0.5 W during the course of the race. Over a typical race-time of 40 h, it would dissipate 0.5 W * 40 h = 20 Wh, or less than 0.5% of the total battery pack energy.

8. **Small size and low weight:** The device should be approximately the same size and weight as all other telemetry nodes in the car.

4.4 Hardware design

This section outlines a careful assessment of the main hardware related decisions for the new device.

[2]Although the level of noise has not been quantified, we know that the presence of multiple high current switching circuits in the car (such as those in the MPPTs and the motor controller) does give rise to significant noise. For example, one source of evidence of this is the observed increase in static in the radio transmission between the driver and the chase vehicle during periods of high motor current.

4.4.1 Current measurement method

Resistive shunt vs Hall effect sensor

The simplest method of measuring current is to use a resistive shunt, as was done in the old current integrator device. A resistor of known value is placed in the current path, and the voltage drop across the resistor is measured with an ADC. The resistor typically has to be very small to minimise power dissipation, and to ensure that the voltage drop across it does affect the operation of the rest of the circuit. This small voltage drop often needs to be amplified before being measured [Friedrich and Lemme, May 2000].

Another common solution, in particular for high current applications, is to use a Hall effect current sensor. These are readily available in integrated circuit form. Such a device works by measuring the voltage developed across a Hall element due to the magnetic field induced around the current carrying conductor. Unlike the resistive shunt, this solution provides galvanic isolation between the high current circuit being measured and the sensing electronics. However, these sensors are typically less accurate than the shunt method due to magnetic hysteresis in the sensor, and a larger temperature dependency [Friedrich and Lemme, May 2000].

After considering both methods, I chose to stay with the resistive shunt solution due to its relative simplicity and typically higher accuracy. I chose a $0.5\,\mathrm{m\Omega}$, low temperature coefficient (20 ppm) precision resistive shunt from manufacturer Isabellenhuette. This gives a typical power dissipation of $0.5\,\mathrm{m\Omega} * (8\,\mathrm{A})^2 = 0.013\,\mathrm{W}$ and a peak power dissipation of $0.5\,\mathrm{m\Omega} * (50\,\mathrm{A})^2 = 1.25\,\mathrm{W}$, both of which are acceptable.

Low side vs high side measurement

Another choice which I needed to make was whether to measure the current at the high voltage terminal of the battery (high-side) or at the ground terminal (low-side). Low-side measurement was used in the old device, and I chose to use this topology once more.

The main disadvantage of low-side measurement is that the voltage drop across the resistive shunt creates a variable potential difference between the battery ground and the ground of the rest of the system [Mehta, 2009], which is a problem in some applications. However, given that the voltage drop across the chosen shunt at maximum battery current was only $0.5\,\mathrm{m\Omega} * 50\,\mathrm{A} = 25\,\mathrm{mV}$, this was not a significant concern for our high voltage system. High-side current measurement usually requires a special IC designed to handle large common mode voltages [Mehta, 2009]. A low side measurement definitely appeared to be the simpler choice. The added advantage of a low-side measurement is that connecting and disconnecting the device does not require interfering with the high voltage battery line, making the procedure inherently more safe.

4.4.2 Choice of analog-to-digital converter

An analog-to-digital converter (ADC) was required to measure the voltage drop across the shunt resistor. The ADC had to have the following properties:

- Ability to measure positive and negative differential voltages across the shunt, since battery current can flow both ways.
- At least two measurement channels, so that both current and voltage can be measured by the same device.
- At least 16-bit resolution

The *MCP3909* ADC chip which was used in the previous current integrator had all the above properties, and had the advantage of existing and tested interface code. It also incorporates an internal programmable gain amplifier, making it possible to measure the tiny voltage drop across the shunt without additional amplification [Microchip, 2006].

I initially hesitated about using this ADC again in light of the calibration issues seen in 2009. However, after investigating the nature of these problems, and comparing them with the published performance characteristics of other similar ICs, I established that these issues were part of normal operation. In light of this, I used the same ADC in the revised current integrator. Device calibration, and how I compensated for non-linearities, is discussed in detail in Section 4.6.1.

4.4.3 Voltage measurement

To measure the high battery voltage, we have to first step it down from the 100–150 V range to the measurement range of the ADC. I achieved this with a high resistance voltage divider, using two 330 kΩ resistors in series and a 1 kΩ resistor. The input to the voltage measurement circuit is fused, to prevent possible high voltage short-circuit faults.

4.4.4 Safe battery connections

The positive battery terminal is connected to the current integrator device using a 25A *Power Pole* connector from Anderson Power Solutions. The battery current connections use the Anderson 55A *Power Claw* connector [Anderson Power Products, 2008]. These allows the current integrator to be easily and safely connected and disconnected from the battery, while providing adequate insulation when the device is operation.

4.4.5 Isolation of high voltage circuitry

The analog section of the device, which measures the battery voltage and current directly, as discussed above, is electrically isolated from the logic circuitry that communicates with the rest of the solar car telemetry system. This is achieved in the same way as in the old current integrator device, using an isolated DC/DC converter (Murata *NMK1209*) to power the analog circuitry, and optical isolators (NVE *IL711* and *IL717*).

4.4.6 Migration to *LPC11C14* microcontroller

This year the *Sunswift* electrical team decided to start using the NXP *LPC11C14* microcontroller for any new telemetry devices, including this revision of the current integrator[3]. Previously, all our custom made telemetry nodes used the Texas Instruments *MSP430F149* controller. The main catalyst for this change was the fact that the *LPC11C14* has an integrated CAN controller. The change simplifies both the hardware and the software required by our telemetry nodes to communicate on the CAN bus.

4.4.7 External memory

In order to make it possible for the integrated current values to be saved more frequently than in the old device, I added an external memory chip. I chose the Microchip *24AA256* external EEPROM chip, which can store 256 Kb and is rated to 1 million write-erase cycles [Microchip, 2004]. Although I did not have time to implement the software interface to this chip in time for the 2011 WSC, its presence will make it possible to save values multiple times a minute without risking over cycling the memory.

4.4.8 Electromagnetic interference considerations

During the hardware redesign process, I placed a lot of emphasis on ensuring the correct design procedures for minimising the effect of electromagnetic interference (EMI) were followed. I added a metal shield around the analog measurement section of the circuit, as well as a common-mode choke filter on the CAN bus communication lines to reduce the effects of noise [Zumbahlen, 2007].

The design of the printed circuit board (PCB) was also greatly affected by EMI considerations, as is discussed below.

4.4.9 Printed circuit board design

- **Segmentation:** The noisy digital circuitry and analog measurement circuitry on the board are placed physically as far away from each other as possible, to reduce coupling between them [O'Hara, 2001].
- **Ground planes and power planes:** I used 4 layers in the PCB design, unlike the old device and most other *Sunswift* telemetry devices, which only use two. I used

[3]Some new devices built this year, such as the SION onboard computer and battery monitoring master device, had special requirements and used the more powerful *LPC1768* micro-controller instead.

the two central layers exclusively for ground and power planes. Ground and power planes in mixed-signal designs provide low impedance paths for currents on the board, which is important for minimising transients [O'Hara, 2001]. The ground and power planes are partitioned according to mixed-signal design conventions, ensuring the digital currents do not have return paths that can interfere with the analog circuitry [Ott, 2001]. Figure 4.1 shows the partitioned group in the current integrator design.

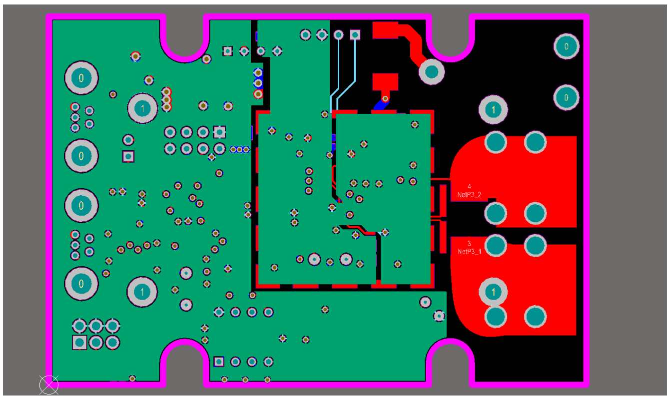

Figure 4.1: Current integrator ground plane, showing separation into analog and digital sections

4.5 Software design

4.5.1 High level functionality

The micro-controller software performs the following high-level functions:

- Upon initialisation, the saved integrated current value is read from flash memory
- Current and voltage samples are read every 100 ms from the *MCP3909* ADC chip through the I2C bus
- The integrated current value between each sample is calculated and added to the cumulative integral variable
- The latest current and voltage samples, as well as the current value of the integral are broadcast onto the CAN bus every 1000 ms
- The cumulative integrated current value is saved onto the device flash every minute

Current integration

The sampled current is integrated in software using the *rectangle rule* principle, illustrated in Figure 4.2.

I attempted to implement integration using the more sophisticated *trapezoidal rule* algorithm, but was unable to get it running correctly in time to the WSC, and chose to revert

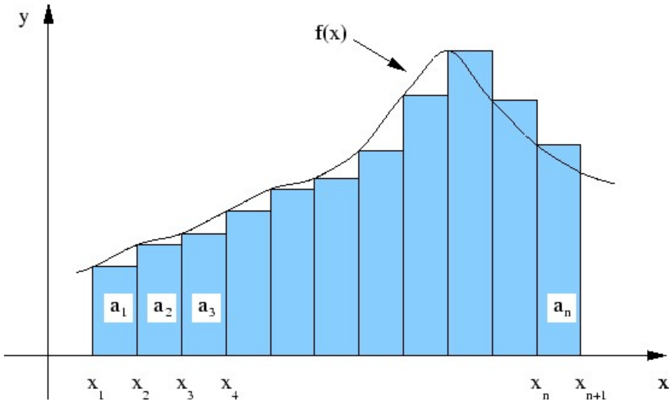

Figure 4.2: Numerical integration using the rectangle rule. *Image courtesy of Charles Dyer [Dyer, 2002]*

to the simpler rectangle rule method. The current is sampled and integrated at a fast rate relative to its temporal variation, so this approximate algorithm is still highly accurate.

4.5.2 Low level functionality and migration to *LPC11C14* architecture

As mentioned in Section 4.4.6, the *Sunswift* electrical team made the choice this year to migrate from the *MSP430F149* to the *LPC11C14* micro-controller. The *Scandal* protocol, used in the telemetry nodes for CAN bus communication, had to be ported to this new platform to ensure compatibility of the new device with the rest of the network. The *Scandal* protocol manages many aspects of the network and node functionality, from the communication format, to a standardised method of saving calibration and configuration information for each node [Snowdon, 2002].

Due to differences in the architectures and build environments for the two devices, significant changes needed to be made to the low-level *Scandal* software to make it run on the new device. The majority of these changes were carried out by Etienne Le Sueur, the current leader of the *Sunswift* electrical team, with some help from other members, including the author. For completeness and future reference, the main areas which needed to be addressed in porting *Scandal* to the new platform are listed here:

- Flash read and write routines
- Timer routines, including the Watch Dog Timer (WDT)
- Linker instructions

4.6 Bench testing and performance

I built and tested the new current integrator device (as well as a spare) in the weeks preceding the 2011 WSC. The completed device is shown in Figure 4.3.

I populated and implemented all the main features of the device, as outlined in the requirements. I ran out of time to implement some of the secondary features in time for the race, including the external EEPROM memory.

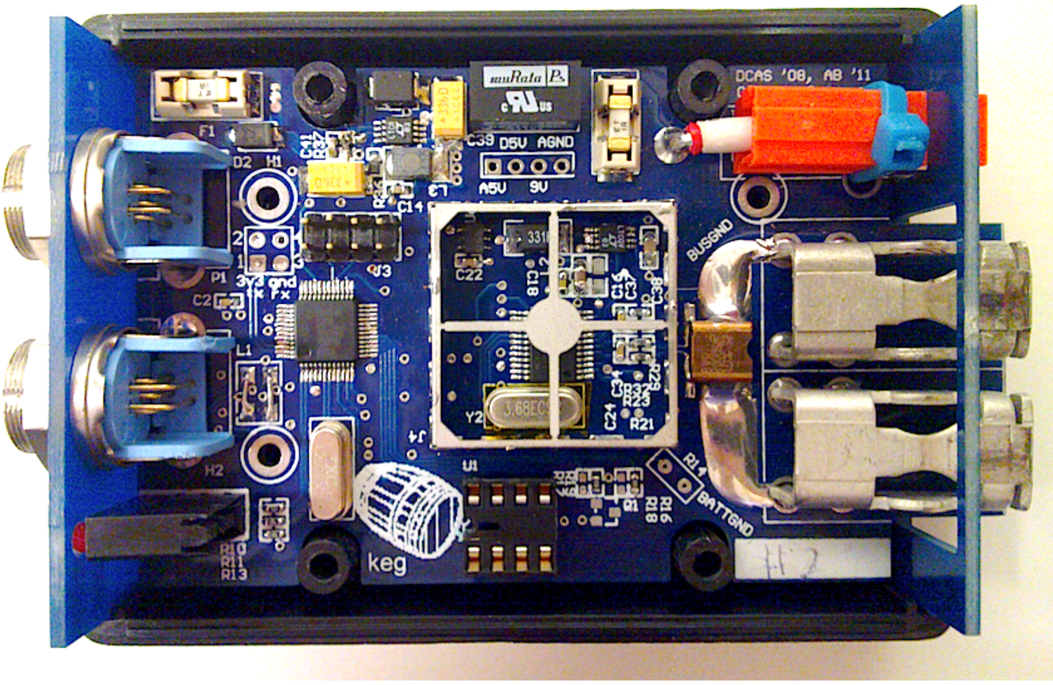

Figure 4.3: The populated new current integrator device.

4.6.1 Calibration

Before the device could be used in the car, it required careful calibration over the full range of operating current and voltage (-50 to +50 A and 100 to 150 V).

While voltage calibration over the required range could be trivially achieved with a variable high voltage power supply and an accurate multimeter (both of which were available at the *Sunswift* workshop), current calibration posed more of a challenge. I had access to a high current power supply from the UNSW Department of Astrophysics, but I did not have a device which could accurately measure current over such a wide range. To overcome this issue, I used a secondary shunt of known resistance, and measured the voltage drop across it to infer the reference current. When taking each calibration measurement, current was only allowed to flow in the circuit for a brief moment while the measurement was taken. In this way, the resistive shunts did not heat up significantly, and errors due to temperature drift were mitigated.

The resulting calibration plots for current and voltage are shown in Figure 4.4(a) and Figure 4.5(a) respectively. At first glance, the ADC response seemed to be highly linear for both measurement channels. Non-linear behaviour, which was expected and discussed in Section 4.4.2 only becomes apparent in the post-calibration residual error plots.

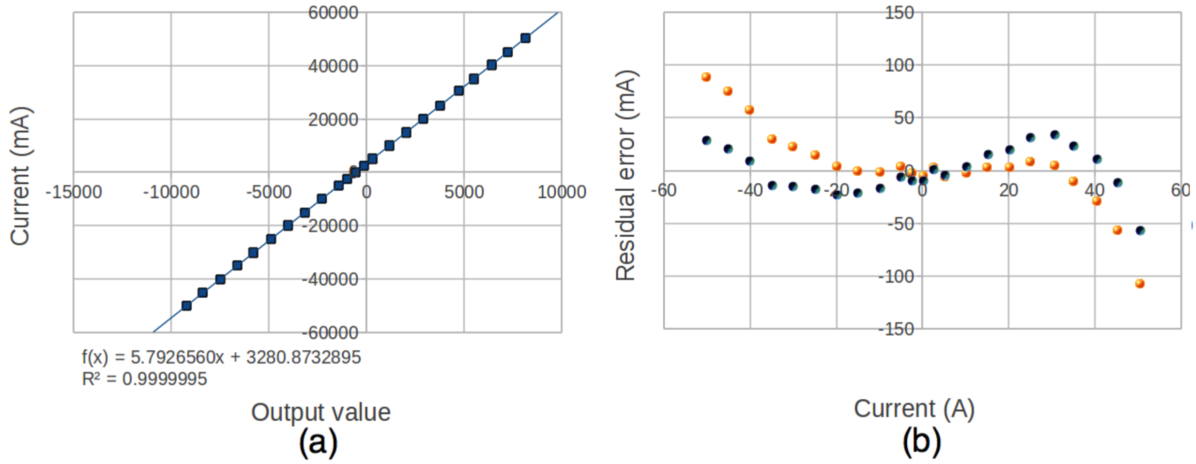

Figure 4.4: Current integration current channel calibration

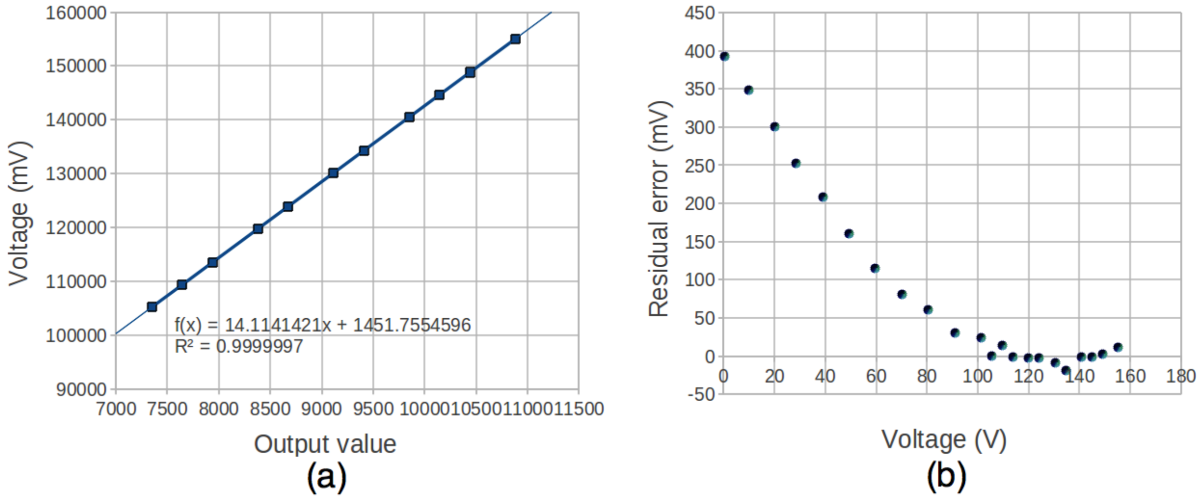

Figure 4.5: Current integrator voltage channel calibration

The residual plot for voltage (Figure 4.5(b)) shows negligible measurement error within the calibrated range of 100 – 150 V, and significant error at low voltages (up to 400 mV). This is acceptable, since the device is used to measure the battery voltage, which never falls below 100V.

The residual error for current, after calibrating over the entire -50 to +50 A measurement range, is shown as the dark blue plot in Figure 4.4(b). The response is clearly non-linear, although the errors at currents below 20 A are only slightly above the 10mA maximum established in the design requirements (Section 4.3).

To minimise these errors further, I re-calibrated the device using only data from the -10 to +10 A range. The residual error for this calibration is plotted orange in Figure 4.4(b). It shows considerable improvement over the -20 to +20 A range, and increased error at values outside it. Since the battery current rarely deviates outside this range during the race, minimising the error here is most important to ensure accuracy in the integrated current. The higher error outside this range is still well within the stated requirements.

4.6.2 Battery discharge

After calibration, the current integrator was connected to the new *Sunswift IV* battery on the bench. I logged the battery voltage, current and integrated current measured by the device while the battery was discharged using a constant current load. Figure 4.6(a) shows a plot of this collected data, which demonstrates the current integrator's correct operation. The log data gives the average discharge current as 6.696 A, and the total discharge time as 5.75 h. The total integrated current value should then be 6.696 A * 5.75 h = 38.495 mAh. The device recorded a total integrated current value of 38.440 Ah, corresponding to an error of just 0.15% after almost 6 h of operation.

Figure 4.6(b) shows a plot of the error in integrated current over the course of the logged discharge (dark blue line). Although the error is very small, it appears to be cumulative over time. It does not seem to have any significant correlation with the deviation of the battery current from its average (red line). This leads me to believe that the source of error is rounding or approximations in the software integration algorithms. Nevertheless, this performance was well within requirements. It would be worthwhile to test whether a different software integration algorithm would lead to improved results.

4.7 On-road performance

Despite performing very well on the bench, the current integrator experienced some initial problems when operating within the solar car's telemetry network. Most of these issues were not identified until after the start of the 2011 WSC. Due to this, for first two days of the race until some of these problems were addressed, the team did not receive data from the device.

The initial symptoms of the problem appeared the first time that the current integrator device was connected to the rest of the solar car network at the *Sunswift* workshop. The device operated normally as long as the the network bus was undisturbed, but would stop sending data whenever another device was unplugged or plugged into the network. We presumed that the *LPC11C14* microcontroller used in the device had some in-built protection capability, which caused the device to pause its operation when a disturbance was sensed on its power lines. This could not be verified from the device data sheet, and we could not yet identify the cause of the problem. We hoped that this was not going to be a issue on the road, since no network devices are connected or disconnected during normal operation of the solar car.

However, just off the starting line on the first day of the race, the current integrator stopped transmitting data. During the course of that day, the GPS telemetry node in the car, which also used the new *LPC11C14* micro-controller, stopped transmitting data as well. We assumed that it must be a low-level problem associated with this new architecture.

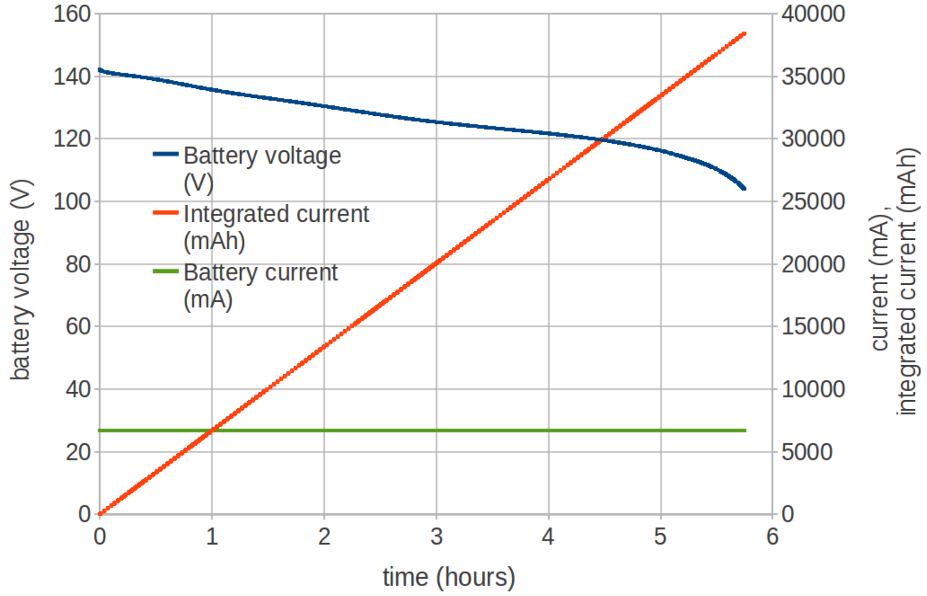

(a) Voltage, current and integrated current measured by the current integrator during a battery discharge.

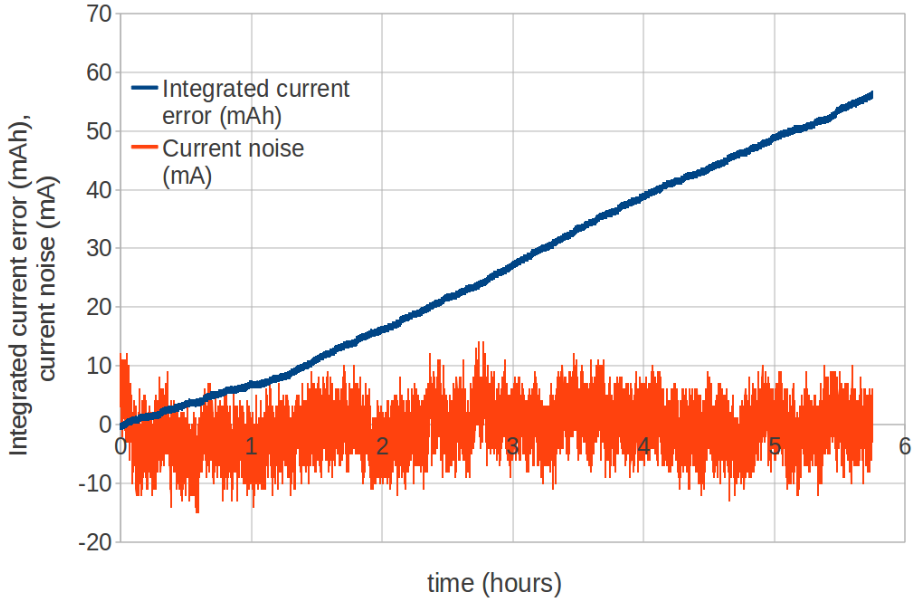

(b) Error in integrated current during the course of the discharge.

Figure 4.6

After analysing the low-level code, we found that there were some instances where the device was held in an infinite loop waiting for a flag register to be cleared. This happened in the driver code for the CAN communication peripheral as well as the SPI peripheral. The *Sunswift* electrical team corrected these errors and after this, the telemetry nodes using the *LPC11C14* micro-controller were much more reliable.

From the 3rd day of the race, the current integrator transmitted current, voltage and integrated current values with good reliability. It did not appear to suffer from the noise problems which affected the previous device in the 2009 WSC. At end of Day 5 the device went offline again briefly, although the cause of this was not identified. Some very thorough testing of this new *LPC11C14* architecture is required before we can confidently use it in all of our telemetry nodes.

Chapter 5

The Race

The work completed and knowledge gained throughout the year was all put to the test during one week in October, in the 2011 World Solar Challenge.

This Chapter presents a logbook of this race, with annotated telemetry data from each day to help tell the story. Our team's race experience demonstrates the way that changes in weather, inaccuracy in the car's models, and technical complications can affect strategy decisions.

5.1 Race logbook

5.1.1 Day 1, Sunday

We started the race in 4th pole position, at 8:34am on Sunday the 16th of October.

The morning's calculated strategy speed for the day is 91km/h, and we went faster than this to get out of Darwin. We were eager to get out of Darwin as quickly as possible to minimise time spent in high traffic and poor insolation due to trees.

The current integrator did not come online at the start of the day, so we did not have an accurate measure of the BSOC.

We soon realised that our model parameters had been too optimistic, and our starting speed too fast, and we slowed down gradually over the course of the day to conserve battery.

In Katherine, we noticed that the rear tyre was damaged and performed a tyre change. This would have contributed to the car using more power than expected.

We ended the day at Dunmurra, with a worryingly flat battery pack. We had less than 5Ah remaining when we arrived, where as we expected to have almost 20Ah.

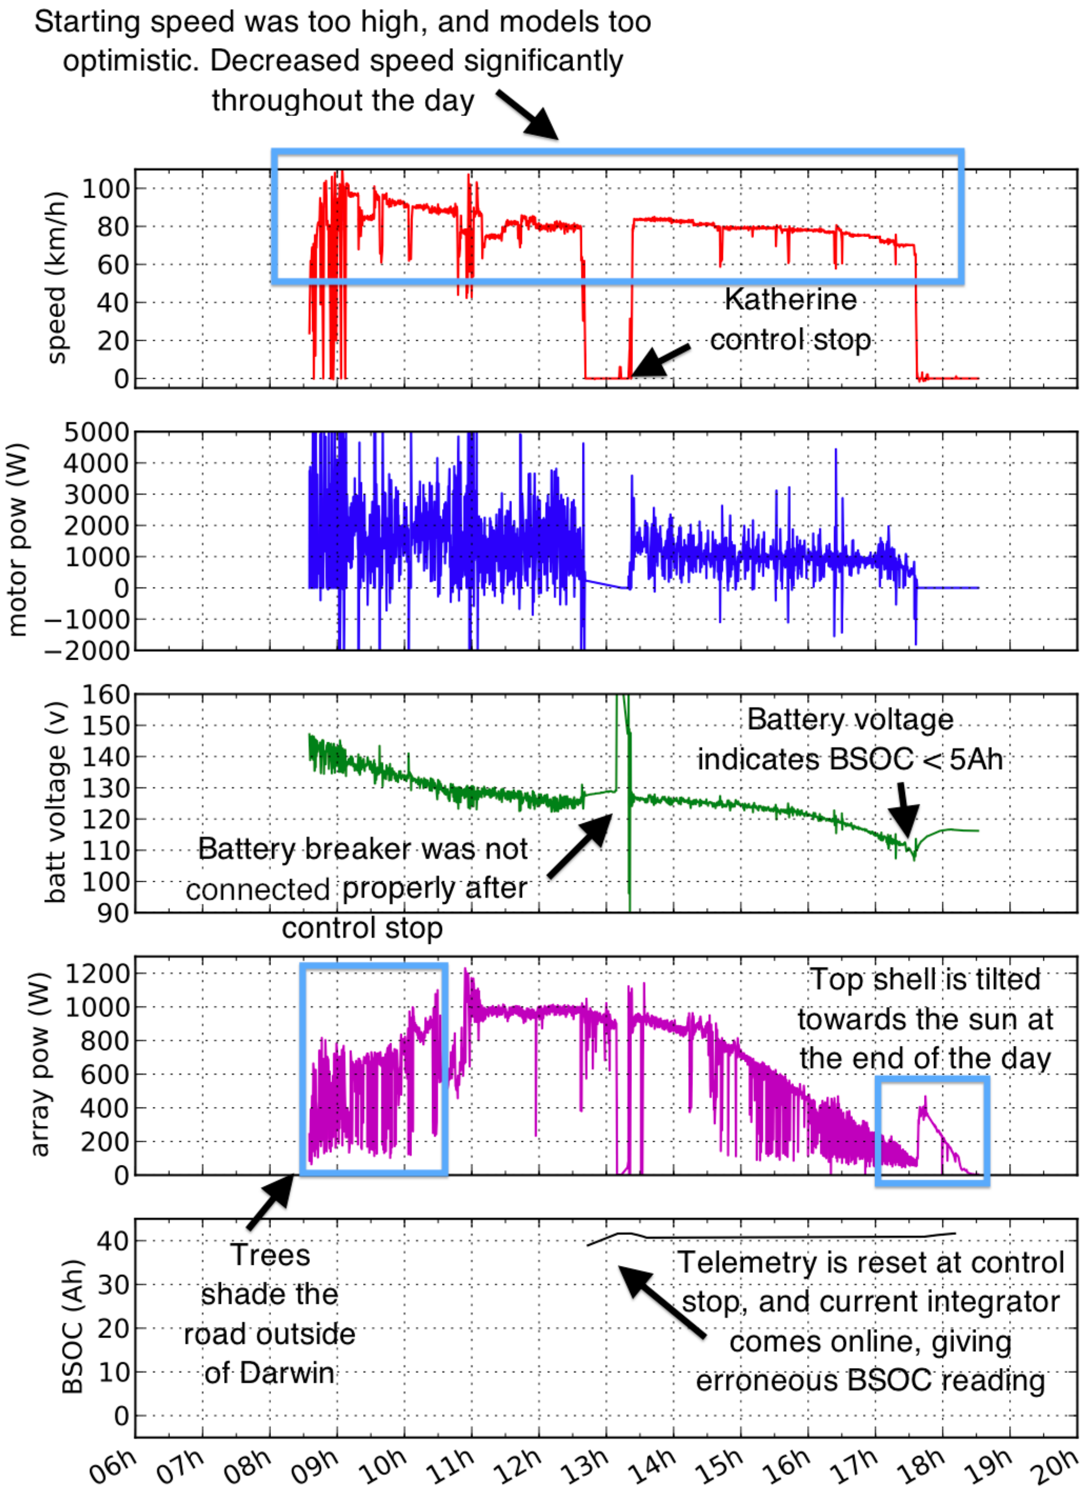

Figure 5.1: Race day 1

5.1.2 Day 2, Monday

In the morning, we carefully planned our next strategy, to avoid making the same mistakes as during the previous day. The weather forecast told us that we could expect good weather until Thursday, when we were due to hit clouds and rain along the route in South Australia.

We manually created speed profiles which used slightly different speeds to travel during good and during bad weather. We then used the strategy simulator to find the best of the strategies we suggested. We found that travelling at 65km/h and charging the battery pack slowly for the next three days would make us have a full pack by the time the bad weather hit. After that, we would cruise into Adelaide at about 60km/h, using the energy in the battery pack through the bad weather.

We successfully followed this strategy to Tennant Creek, adjusting our model parameters slightly when the car performed slightly better than the simulation.

In Tennank Creek, however, were informed no cars could proceed due to bush fires along the highway ahead. We were held there until the following morning, and charged our battery pack fully. Had we known this in advance, we would have travelled faster in the first half of the day.

That night we had time to address some of the telemetry problems experienced during the first two days of the race. The electrical team found issues in the low-level code for the current integrator and GPS and fixed them. We hoped these devices would now be reliable.

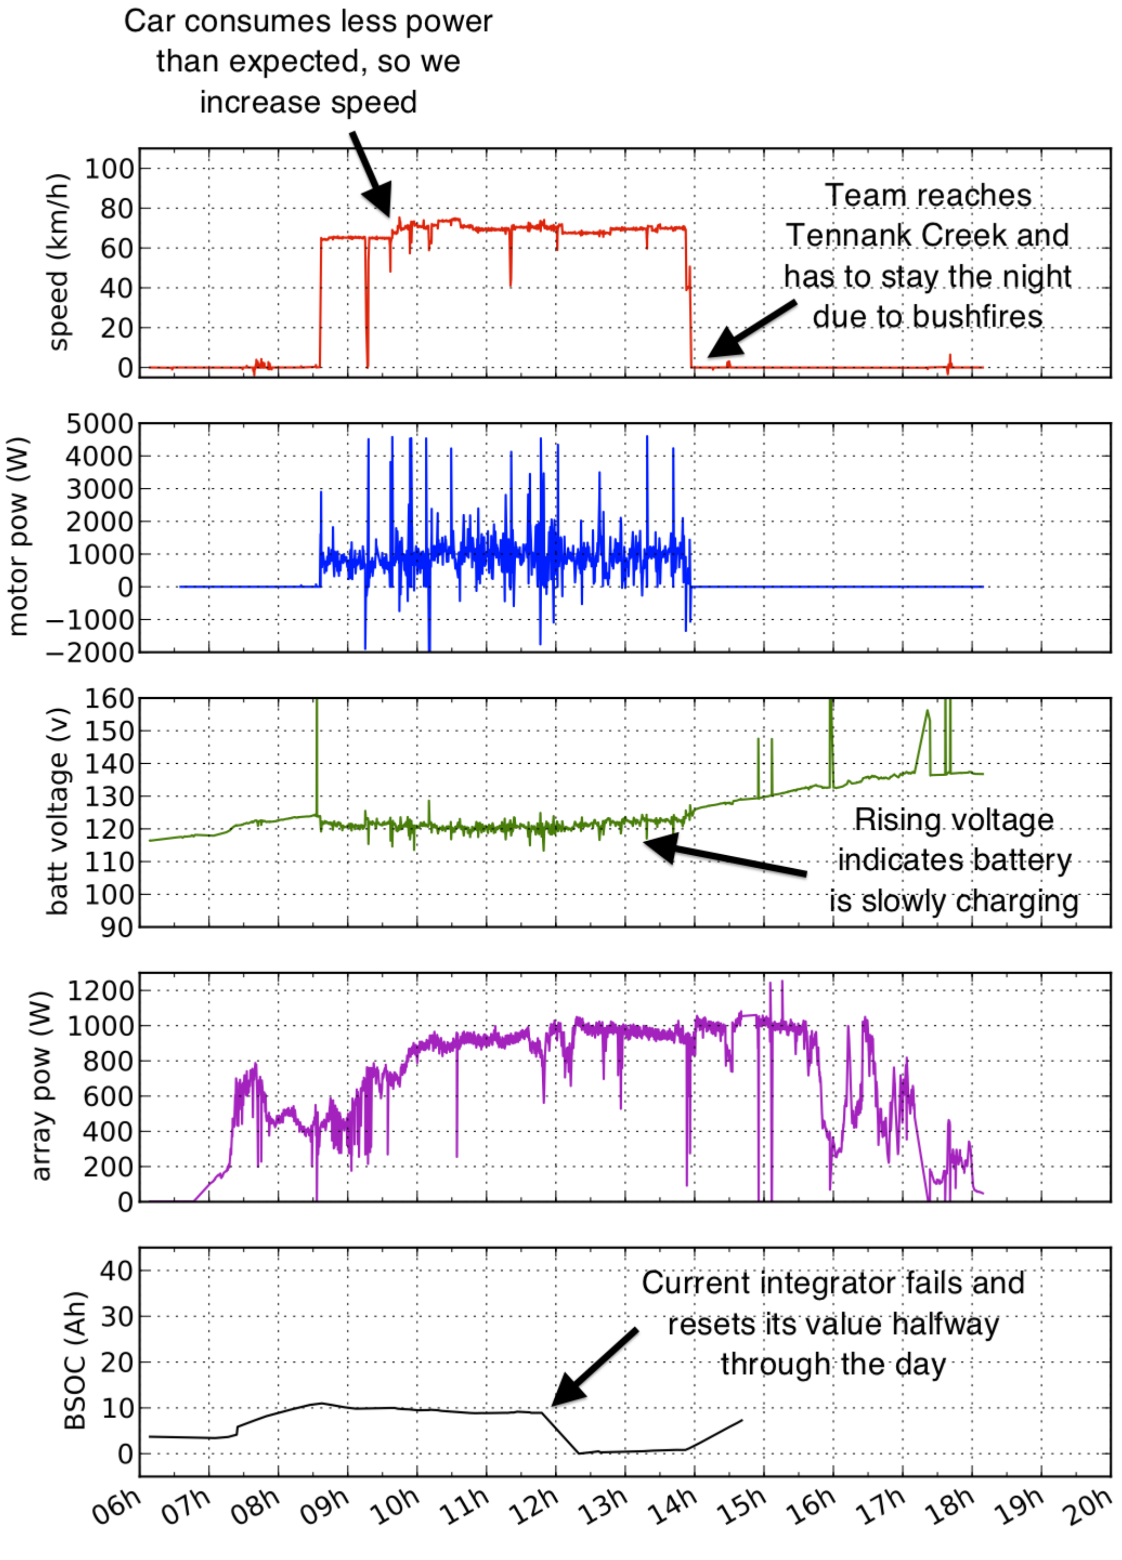

Figure 5.2: Race day 2

5.1.3 Day 3, Tuesday

The teams were allowed to leave Tennant Creek in the morning, maintaining the time separation of their arrivals the previous day. We left at 9:20am with a practically full battery pack.

Due to the bushfire delay, we now expected the bad weather ahead to affect us more severely. We would be further from the finish line when it reached us, and we would have to travel slowly through cloudy weather for longer to get there. Nevertheless, the best long term strategy we could employ was to make sure our battery pack was charged when the rain hit.

We were quite conservative with our models on this day, not wanting to risk running the battery pack close to flat again like we did on the first day of the race. The battery level was consistently above what was expected by the strategy system. We increased our speed gradually to compensate for this, but it was not enough to perfectly track the expected battery state of charge curve. We stopped for the night just after Alice Springs, with a battery pack that was fuller than it should have been. This unfortunately meant that the following morning, we would not be able to store all the available solar charge in our full battery pack.

Another problem that day was a flat tyre just after 11am, which the mechanical team changed within a few minutes. Since we seemed to be consistently having problems with the untested Schwalbe *Explorer S* tyres, the mechanical team decided to change the two front tyres to the Sava *MC2*.

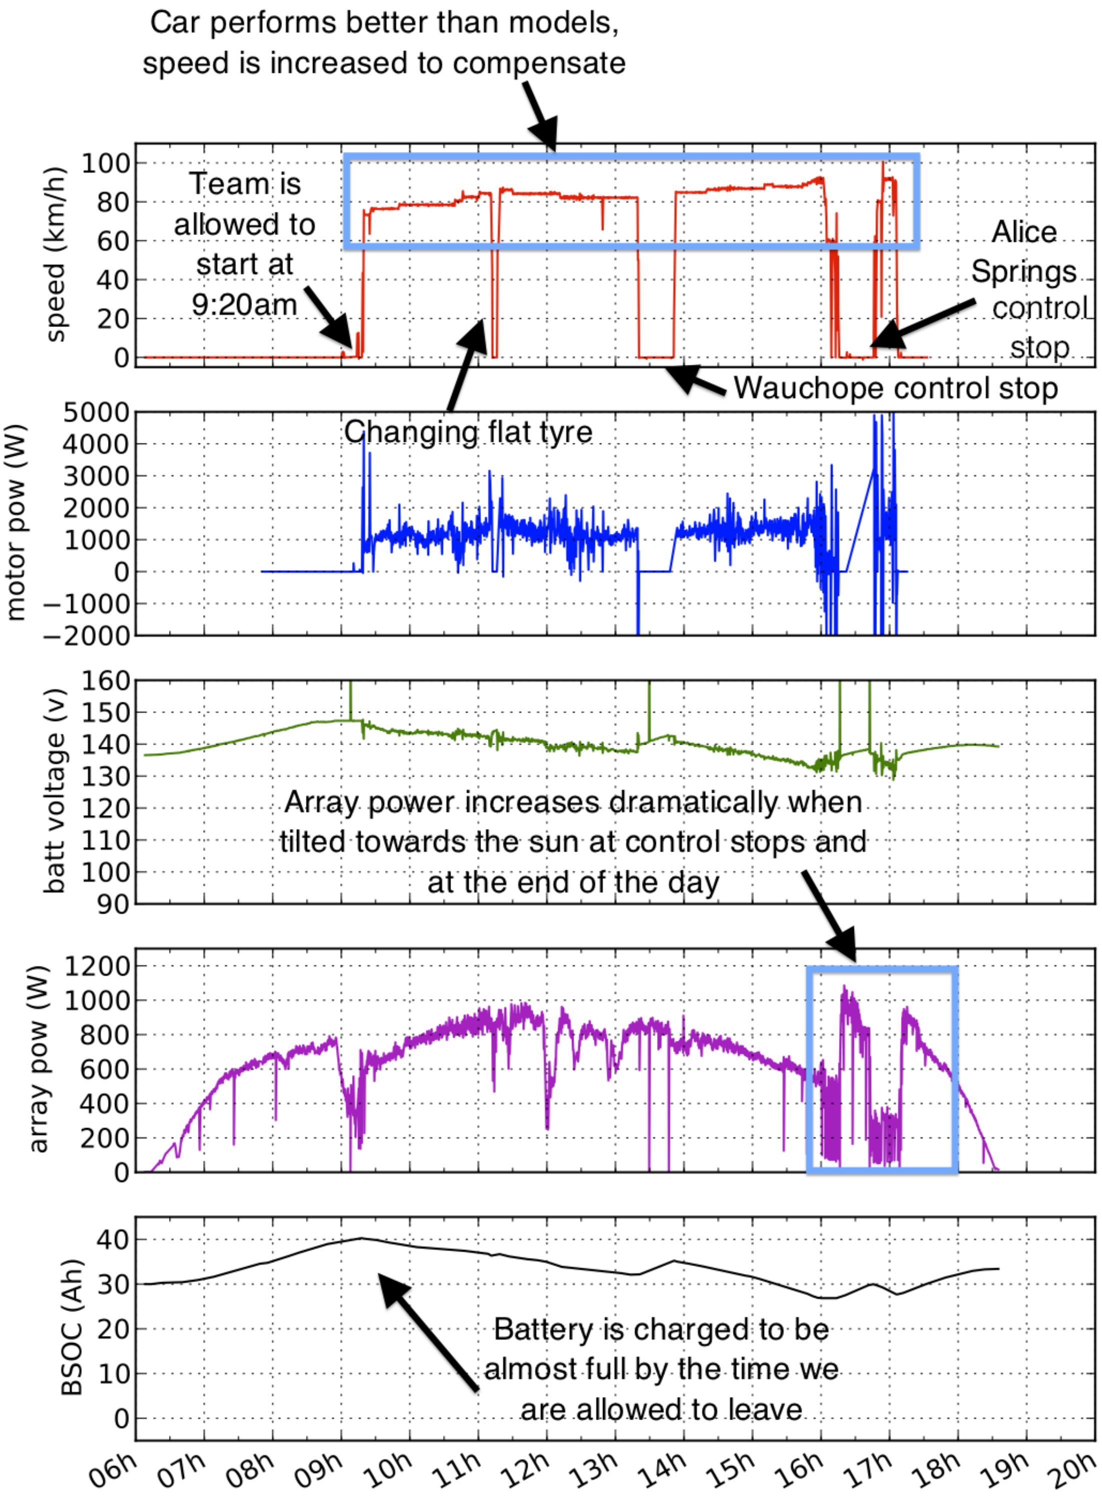

Figure 5.3: Race day 3

5.1.4 Day 4, Wednesday

The morning weather forecast showed that the weather outlook was improving significantly. The rain that we expected ahead was no longer supposed to be as severe. We had high confidence in this new forecast, which was only predicting a couple of days ahead.

We recalculated strategy to account for this, and it gave us an optimal speed of 90 km/h for the day. We started the day at this speed, slowly gaining on the Aurora and Twente solar car teams 30 minutes in front of us. Umicore was only 15 minutes behind us.

We did not get to stick to this strategy for long, however, because within a few hours, we started to see strange behaviour in the car's power consumption. The power use would sporadically jump up from a stable 1500 W to over 3500 W for periods of a few seconds. At first, we thought it was the effect of sidewinds on the car, but that was not the problem.

This increased and intermittent power use was quickly flattening our battery pack, and we were forced to pull over to try to fix the problem. During this time, the Umicore team overtook us.

Over the course of the next three hours, we started and stopped the car multiple times, swapping out almost every component in the car's power train to try and fix the issue. Nothing seemed to help, and in the end it ended up being a much simpler problem.

The *MC2* tyres, which were fitted the night before, left only a very small clearance to the aerodynamic fairing which covered the wheels. When not adjusted precisely, they were prone to rub against the fairing, and this is exactly what had been happening during the course of the day. 2 kW of power was being dumped into heat and deformation whenever the tyre and fairing touched.

We finally identified and fixed the problem by the end of the day, but had severely fallen behind in the pack, and wasted about a third of our battery pack.

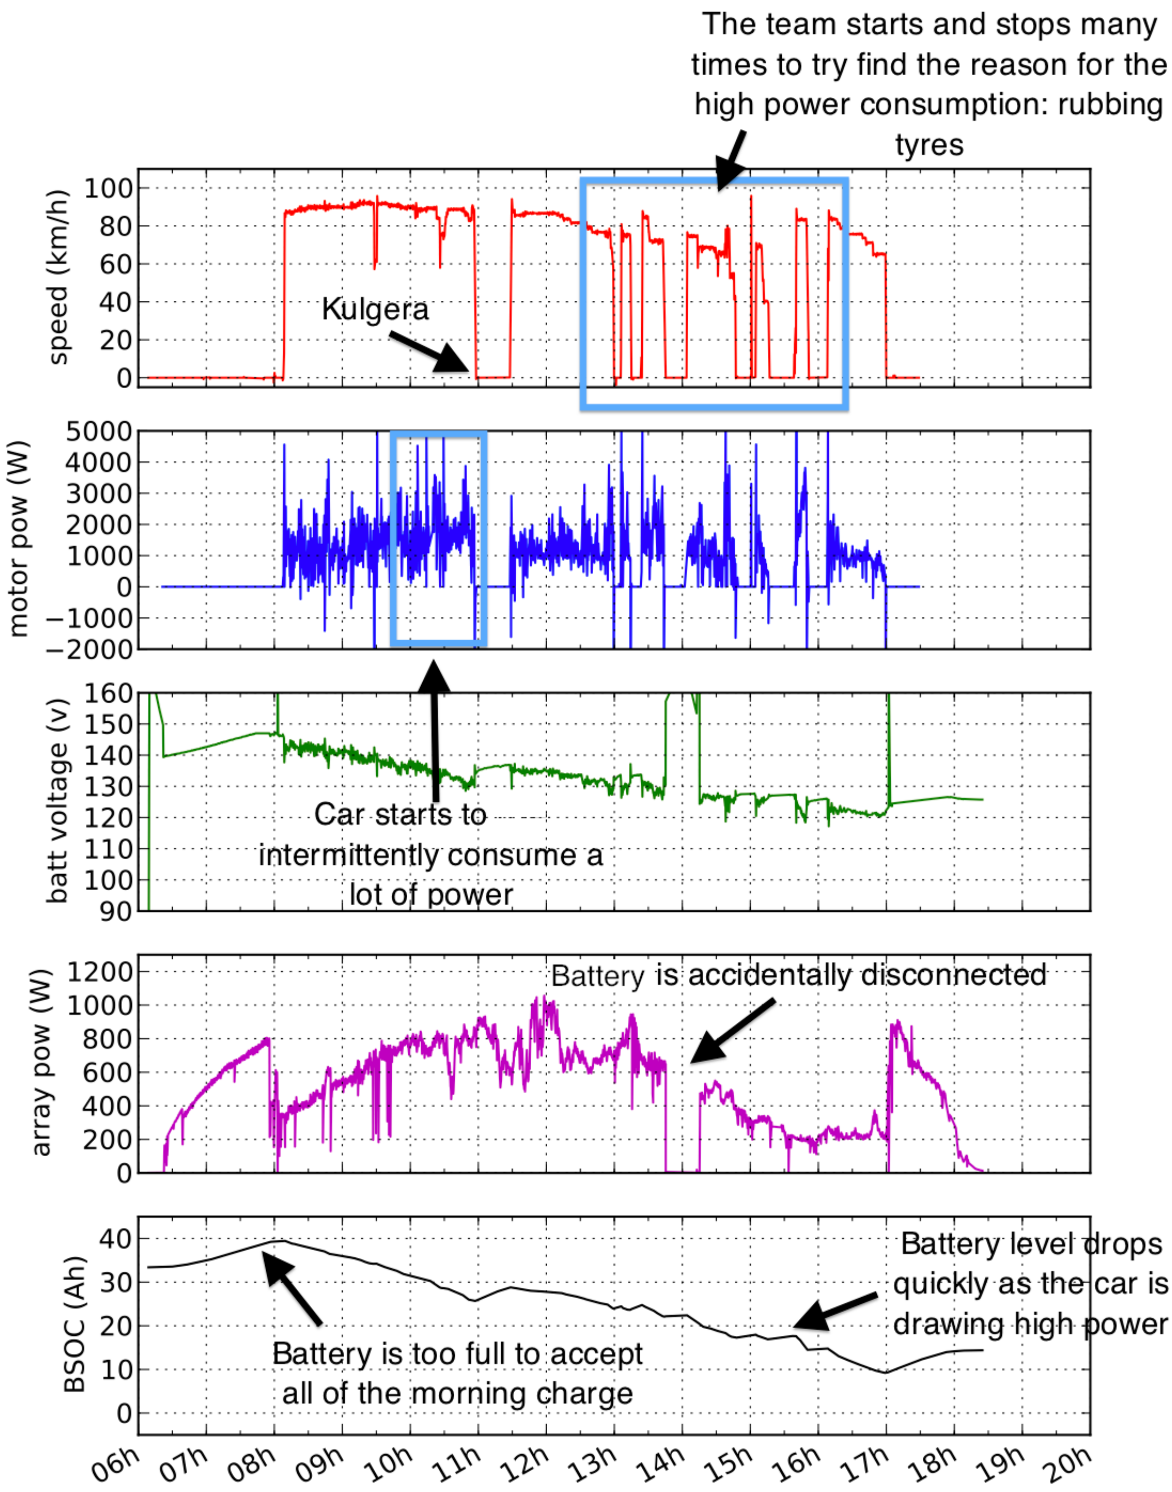

Figure 5.4: Race day 4

5.1.5 Day 5, Thursday

By this time we felt that we had no hope of catching the rest of the teams, and all we could do was travel to our own strategy for the rest of the race. The weather forecast still looked optimistic, and with a calculated set speed of 80km/h we were hoping to reach the finish line by the end of Friday.

However, within a few hours into the day, we realised that we were not receiving as much array power as the weather forecast promised. Just after 12:30 pm we received an updated weather forecast, which had changed dramatically since the previous night. We found that we were heading straight into the cloudy and rainy weather which we had expected days earlier.

In response to this, we were forced to drastically change our strategy. A strategy recalculation gave the new set speed at just over 50 km/h, and we slowed down to this just as we started to enter the thick cloud which would surround us for the next 24 hours.

At this point, we sent our scout car 100 km ahead of the fleet, carrying a calibrated reference solar cell to measure the insolation along the route. They reported back periodically, confirming that the weather in front was just as dire as the latest forecast predicted.

In addition to clouds, this weather system brought high wind. Towards the end of the day, we were forced to slow down even more, to 40 km/h, as the driver was struggling to control the solar car against 80 km/h wind gusts.

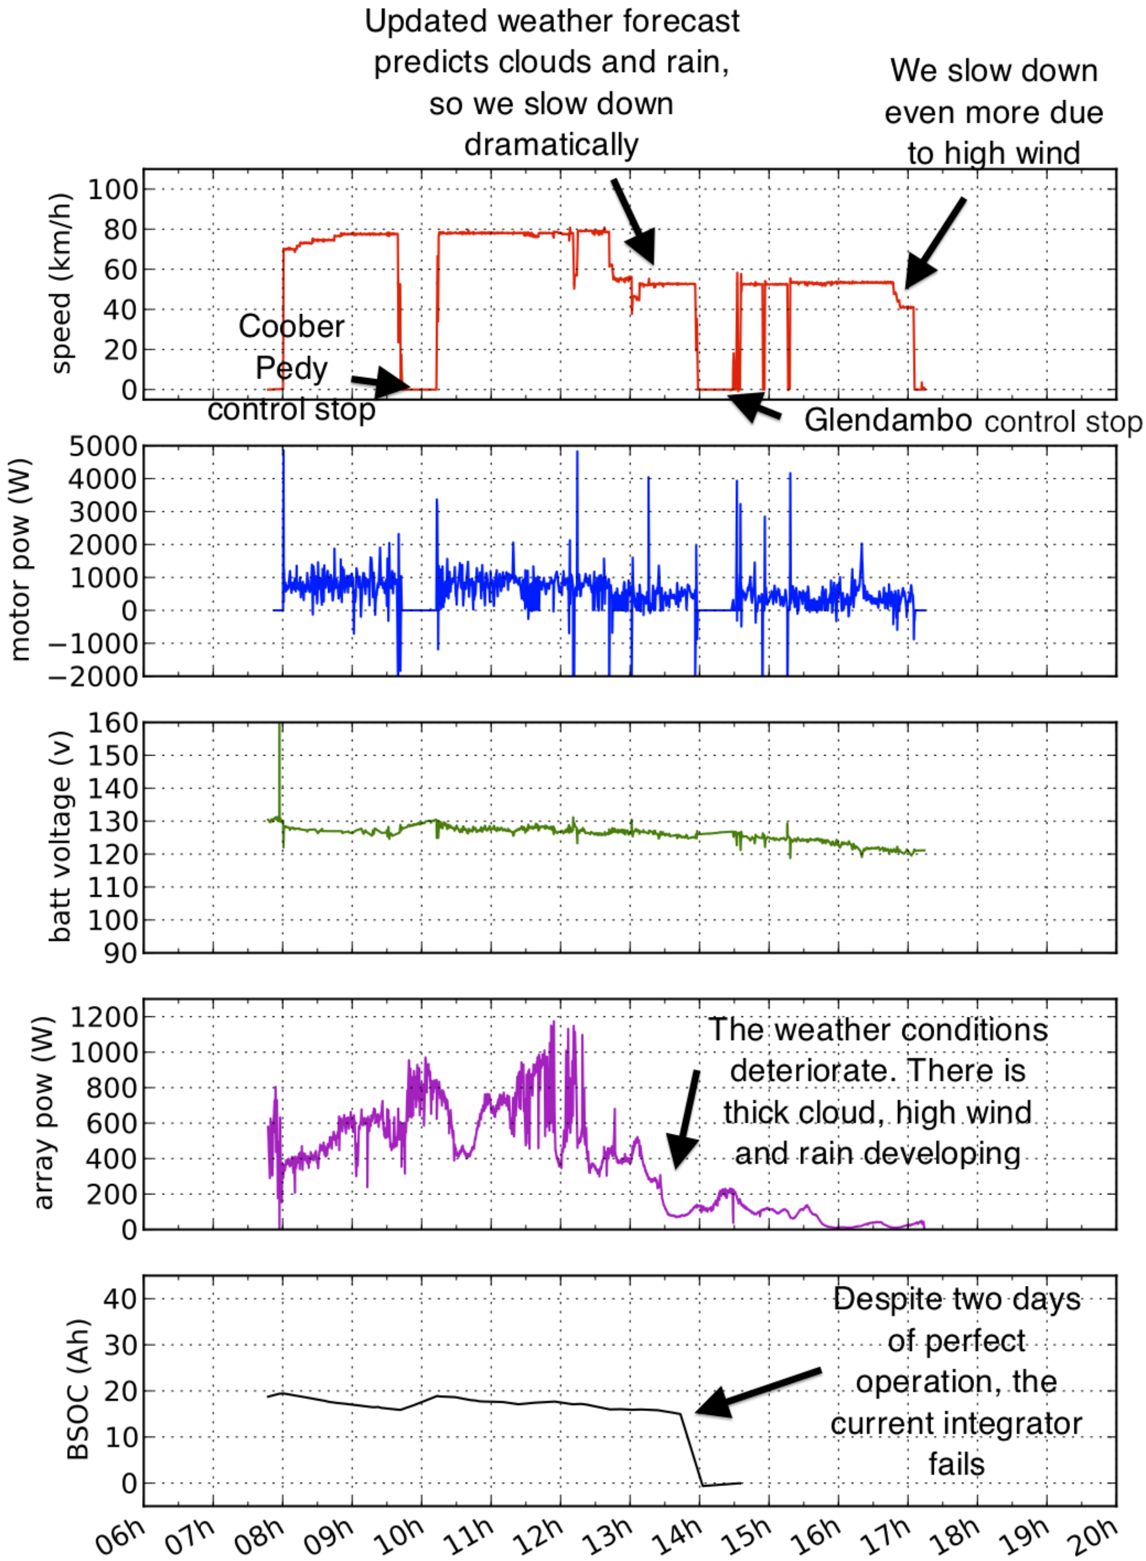

Figure 5.5: Race day 5

5.1.6 Day 6, Friday

By the morning of the 6th day, the thick clouds had developed into rain. We recovered almost no battery energy during the afternoon and morning charges, and could not run any faster than 40 km/h. At this rate, we could not make it into Adelaide until Saturday.

We slowly rolled into the last control stop at Port Augusta. By that time, the top shell of the car was completely wet. Some members of the electrical team had invested a lot of time waterproofing the top shell the night before, but it was inevitable that some water was also leaking into the solar car.

Just outside of Port Augusta, we stopped receiving all telemetry data. As we would find out later that night, this was because the on-board computer had failed due a short circuit caused by the water inside the chassis.

We continued on despite this, and soon after leaving Port Augusta we caught up to the Aurora solar car, which was travelling at under 20km/h, also crippled by the bad weather.

We did not have enough battery energy remaining to overtake them. Eventually the two teams agreed to pull over early, at 4pm, as having both fleets travel slowly in a line on the busy highway out of Port Augusta was quite stressful and dangerous. We stopped for the night in Port Germain, just 240km from the finish line.

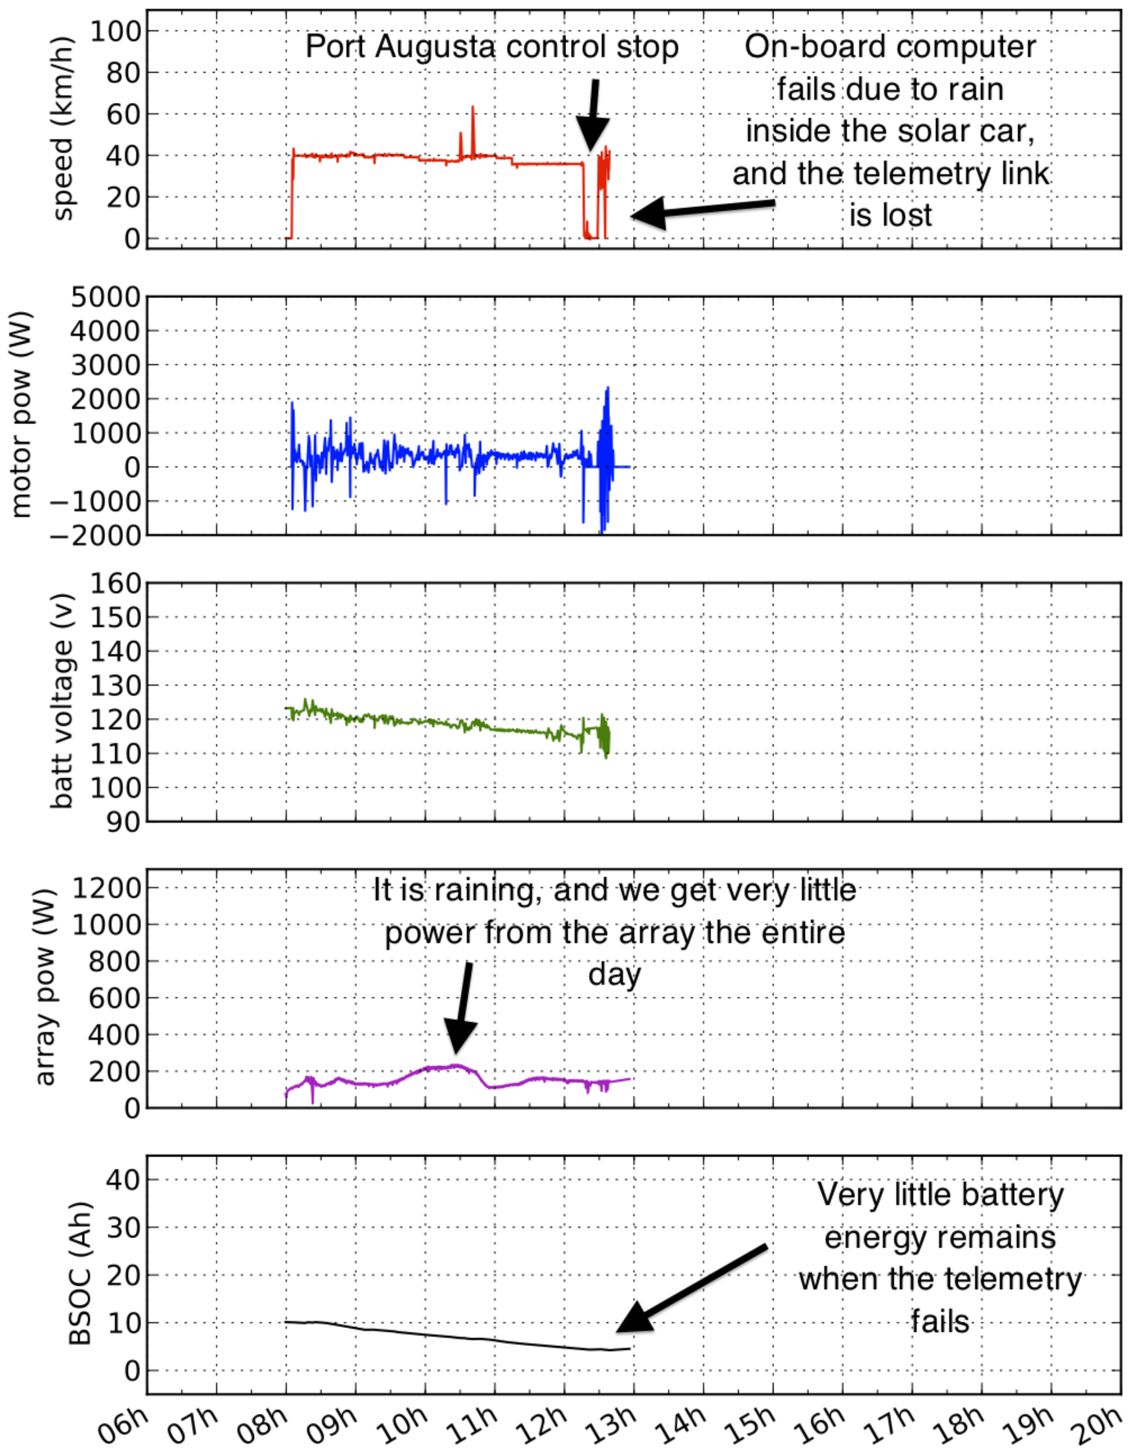

Figure 5.6: Race day 6

5.1.7 Day 7, Saturday

Thanks to a lucky break in the weather, we managed to get some charge back into our almost flat battery pack during the morning of the 7th day. This was enough to get us into Adelaide.

We overtook Aurora almost immediately after leaving the campsite, and continued to make good progress.

The weather forecast was still quite uncertain, and quite cloudy, so we kept to a conservative speed to avoid flattening the battery pack within an arm's reach of the finish line.

Fortunately, the weather held out, and we arrived at the end of timing in Angle Vale at exactly 11am that morning, just 7 minutes ahead of Aurora.

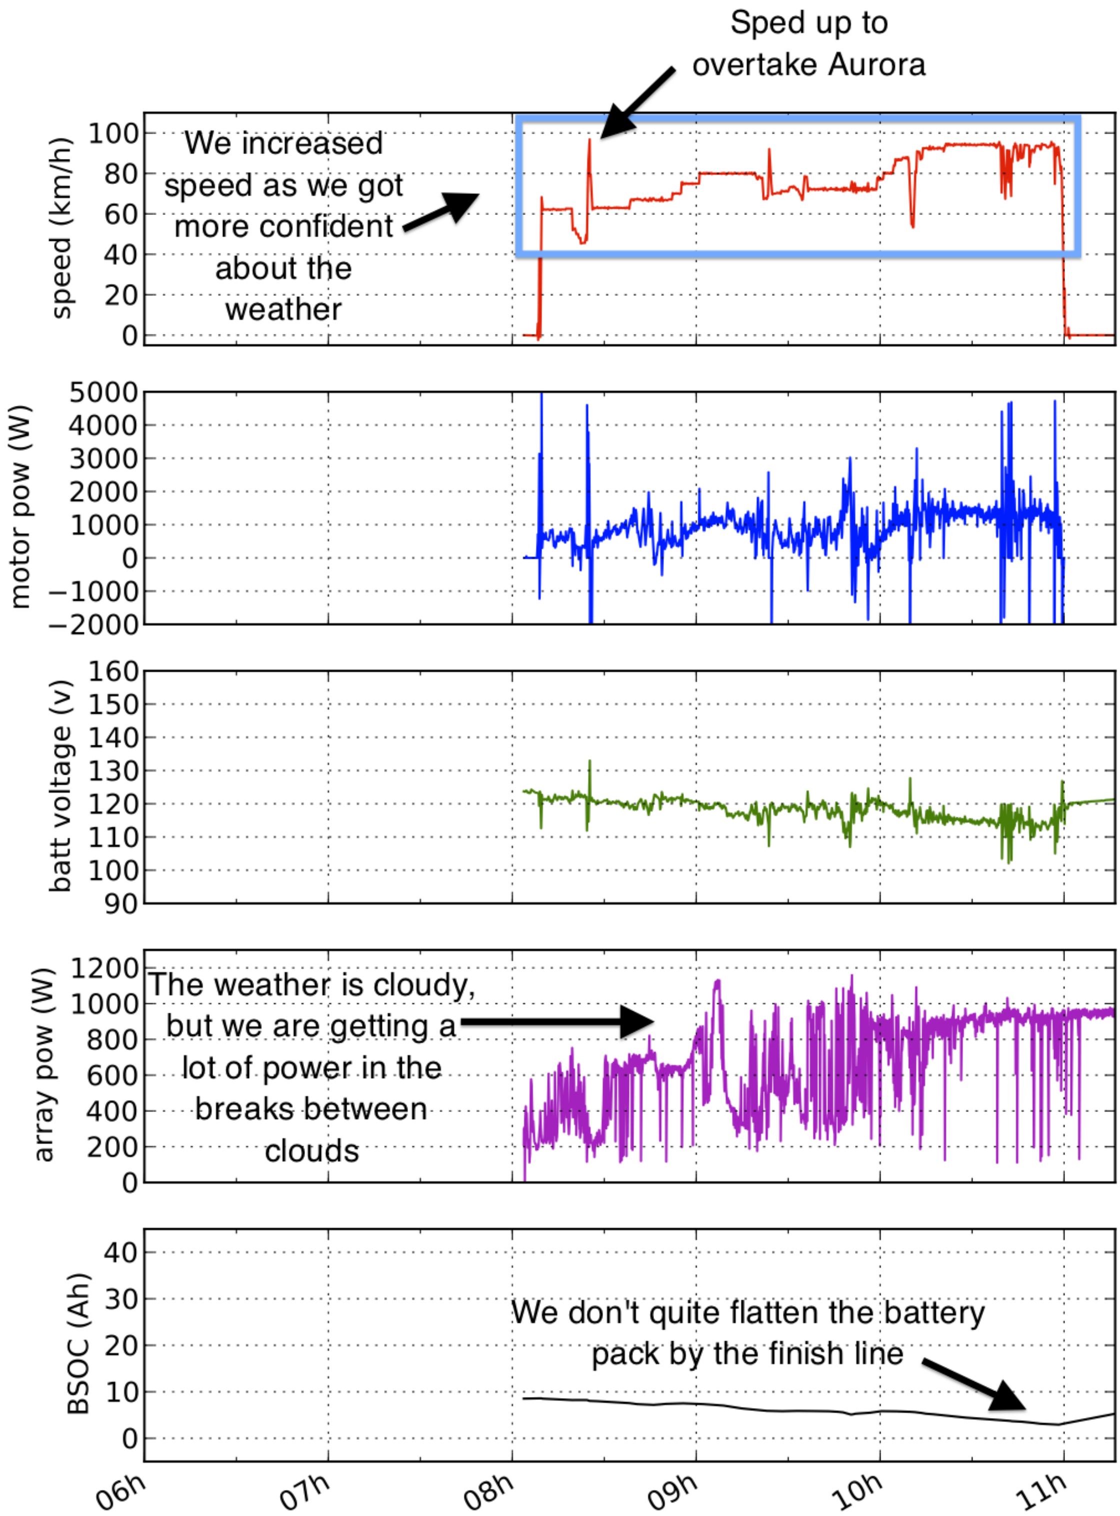

Figure 5.7: Race day 7

5.2 Discussion

This year's WSC competition was an usually eventful one, with extreme and unpredictable weather and very close competition among the top ten teams. The *Sunswift* team was faced with many challenges, and plenty of difficult strategic decisions. Many important observations can be made from this experience regarding practical solar racing strategy. The following stand out in particular:

- The strategy software was very useful in helping to make informed decisions regarding the approximate speed that the solar car should run at every day. However, a lot of manual adjustment and recalculation was required to compensate for inaccurate models and for the highly variable weather.
- Telemetry feedback was extremely important during the race. In particular, the ability to compare the expected power consumption of the car to the real power consumption was very useful. This helped with adjusting model parameters, as well with the diagnosis of faults, such as tyre rubbing problem experienced during day 4.
- A particular calculated strategy could never be followed for too long without having to recalculate either as a result of changing weather, inaccurate models, or unexpected events.
- The current integrator provided a practical way of following a calculated strategy, by tracking the actual battery state of charge to the predicted one. We could follow a calculated strategy with a lot more confidence on the days when the current integrator worked properly (Days 3 and the beginning of Day 4 in particular).
- When the car's performance deviated significantly from the expected performance, we did not always respond quickly enough, resulting in unnecessarily wasted energy. During day 1, we progressively slowed down over the course of the day, after realising that the car was performing much worse than expected. We could have saved a significant amount of energy if we had reacted quicker, and stuck to one average speed over the course of the day. Similarly, during day 3, when the car was performing better than expected, we should have adjusted the models right at the start and followed a constant speed that would get us to the desired battery state of charge by the end of the day. This way, we would have avoided having a battery that was too full the following day to accept all of the morning solar charge.
- During cloudy weather, the array power was very difficult to predict. Clouds diffuse solar radiation, and silicon solar cells do not work well in diffuse light. This is something we don't account for in our model of the array. As a result, the array power during the cloudy and rainy days (day 5 and day 6) was consistently overestimated.

Chapter 6

Conclusions and Future Work

The goal of this thesis was to improve upon the previous *Sunswift IV* strategy system by obtaining accurate models of the car's mechanical and electrical systems, and to implement a reliable current integrating device to measuring the battery state of charge during the 2011 World Solar Challenge.

During the course of the thesis I carried out multiple controlled roll-down tests on the *Sunswift IV* solar car, and developed an effective technique for analysis the data to extract the model parameters. Although I was not able to obtain the highly accurate results which were required by the strategy models, the analysis method developed will be very useful for future characterisation of the team's solar cars.

The results of the roll-down tests were also very useful in helping the *Sunswift* solar racing team make informed decisions regarding the tyres used in the 2011 World Solar Challenge.

An effective model of the *Sunswift IV* lithium ion battery was obtained, which takes into account its equilibrium discharge curve and internal resistance. This model was used in the strategy simulation software during the 2011 WSC.

A redesign of the current integrator device was also completed, and thoroughly tested on the bench. The device was used in the 2011 WSC by the team, and despite some initial problems, it successfully measured and transmitted the solar car's battery state of charge during many days of the race.

Using the solar racing strategy knowledge gained throughout this year, I raced *Sunswift IV* in the 2011 World Solar Challenge with the UNSW Solar Racing Team. I contributed to all of the important strategic decisions during the race. The team performed very well, arriving 6th of the 37 competitors.

6.1 Further Work

6.1.1 Car models

Mechanical models

The *Sunswift IV* roll-down testing technique used during the course of this thesis should be used again to test future *Sunswift* cars. In particular, this method should be used to systematically evaluate the effect of different changes to the car, in the way that was done this year with tyres.

The following additional work can be done to improve the roll-down method:

- Wind measurements should be recorded during the tests, and accounted for when analysing the data.
- The parameters extracted from the roll-down data should be verified using constant speed tests.
- Simple analysis software could be developed to quickly and consistently extract parameters from the roll down data.

Electrical models

- The current battery model could be improved upon by including transient effects. It would be worthwhile to develop a technique for determining the parameters of the *parallel RC* battery model, which was mentioned in this thesis but not implemented.
- The array model should be improved to take into account the curvature of the top shell. As mentioned previously, most of the work for this has already been done by members of the *Sunswift* electrical team, but needs to be completed.
- The MPPT, motor and motor controller efficiencies should be verified experimentally

6.1.2 Instrumentation

Low-level problems with the *LPC11C14*

It is of very high priority for the *Sunswift* electrical team to diagnose and fix any remaining problems with the *LPC11C14* micro-controller within the telemetry network as soon as possible. We made the transition to this architecture from the over a relatively short period of time during 2011, and significantly more testing needs to be done to verify its robustness within the *Sunswift* telemetry system.

Additional current integrator features

The external EEPROM memory which was included in the design of the new current integrator, but was not made to work before the race should be implemented. In addition, the integration algorithm used in the software could be improved.

Measuring local weather conditions

The effectiveness of the feedback obtained about the car's operation during the course of the race could have been improved greatly with the addition of windspeed measurements, as well as more regular measurements of the local insolation. For the next race, the strategy system should incorporate a reliable windspeed sensor and reference solar cell.

6.1.3 Race data analysis

A lot of valuable information can be extracted from the logged telemetry and strategy data collected during the 2011 World Solar Challenge. This information will give the next *Sunswift* team important feedback to guide the further development of the *Sunswift* strategy system as well as the design of next *Sunswift* car.

I consider the following analyses to be the most important for further work:

Retrospective evaluation of different strategies

Using the latest weather forecasts received during the 2011 WSC, different possible strategies should be evaluated against each other, and against *Sunswift IV*'s actual performance. Possible strategies include:

- one constant speed over the entire race,
- one constant speed for each day,
- *hill anticipation*,
- *sun chasing*.

Retrospective car models

It would also be useful to verify how well the car's performance during the course of the race fits to the models which we used. Possible analysis includes:

- Obtaining the power consumption of the car during the race as a function of speed. This involves plotting the power at each point during the race against the speed, compensated for gradient, and fitting the mechanical model to the data to extract the model parameters.

- Refining the above model further by compensating for the effect of wind at each point, using the latest weather forecast received for the 2011 race. The spread of the data should be narrower after the inclusion of the wind term.
- Obtaining the battery equilibrium curve and the battery internal resistance from the battery voltage, current and state of charge data during the race.
- Obtaining a model of the performance of the solar array in diffuse light, using insolation and array power data from days 5 and 6 of the race, which were very cloudy.

Effect of driver technique

Something that has been brought up often in discussion about solar racing strategy is how significantly the ability of the driver affects the car's power consumption. During the 2011 World Solar Challenge, *Sunswift IV* had three drivers. Two of them were student members of the team with no previous racing experience, while the third was Barton Mawer, a professional racing driver.

It would be interesting to see if whether the power consumption model of the car is different for these three different drivers, and whether any conclusions can be drawn about the relative importance of driver technique in the race.

6.1.4 Optimisation

Finally, once the team has obtained accurate and tested models of the car, further strategic benefits can be achieved through improving the optimisation algorithms used by the strategy system. The improvements made by David Snowdon to the strategy software over the last year make it really easy to incorporate new optimisation methods, and this definitely something which should be attempted before the next race.

Bibliography

Anderson Power Products. Powerpole connectors website. http://www.andersonpower.com/products/singlepole-connectors.html, 2008.

Carson Au and Teng Zhang. Current integrator user's guide. 2009.

Chris Beeves and Graham Doig. The aerodynamic design and performance of Sunswift IV. 2009.

Biel School of Engineering and Architecture. Mppt new generation user's manual. Technical report, University of Applied Science.

Michael Daniels and P Kumar. Racing with the sun: The optimal use of the solar power automobile. *Decision and Control, 1997., Proceedings of the 36th IEEE Conference on*, 1:571–576, 1997.

M Debert, G Colin, M Mensler, Y Chamaillard, and L Guzzella. Li-ion battery models for HEV simulator. Technical report, Renault, 2008.

Sam Diasinos and Graham Doig. Personal communication regarding roll-down data analysis, 22nd August 2011.

Charles Dyer. Rectange rule. http://pathfinder.scar.utoronto.ca/~dyer/csca57/book_P/node40.html, 2002.

Electropaedia. Battery performance characteristics. http://www.mpoweruk.com/performance.htm, 2005.

David Favaloro. Battery and battery monitoring system for Sunswift IV. BE thesis, School of Electrical Engineering, University of NSW, Sydney 2052, Australia, December 2011.

Andreas P. Friedrich and Helmuth Lemme. The universal current sensor. http://www.sensorsmag.com/sensors/electric-magnetic/the-universal-current-sensor-1029, May 2000.

Ruiyuan Li and Sheng Qiu. Personal communication and Sydney Internation Regatta Centre survey data, 23rd September 2011.

Lovatt, Ramsden, and Mecrow. In-wheel motor for solar-powered electric vehicles: technical details. Technical report, CSIRO, 1997.

Campbell McLaren. Structural design of SSIV. BE thesis, University of NSW, Sydney 2052, Australia, 2009.

Campbell McLaren and Luke Bycroft. 2011 wheel motor design report. Technical report, Sunswift, 2011.

Arpit Mehta. Low-side vs. high-side current sensing. Technical report, Maxim Integrated Products Inc, 2009.

Microchip. 24AA256 datasheet. http://ww1.microchip.com/downloads/en/devicedoc/21203m.pdf, 2004.

Microchip. MCP3909 datasheet. http://ww1.microchip.com/downloads/en/DeviceDoc/22025a.pdf, 2006.

Ceriel Mocking. Optimal design and strategy for the solutra. Master's thesis, University of Twente, 2006.

Martin O'Hara. Grounding strategies for printed circuit boards. *Compliance Engineering*, March/April:Article 3, 2001.

Optimum G Vehicle Dynamics Solutions. Aerodynamic analysis: Estimation of aerodynamic performance from constant speed and coast down testing. http://www.optimumg.com/OptimumGWebSite/Documents/Aero_BertaReport.pdf, year = 2011.

Henry N Ott. Partitioning and layout of a mixed-signal PCB. *Printed Circuit Design*, June:8–11, 2001.

Panasonic. NCR18650A datasheet. http://industrial.panasonic.com/www-data/pdf2/ACA4000/ACA4000CE254.pdf, 2010.

Peter Pudney. *Optimal energy management for solar-powered cars*. PhD thesis, University of South Australia, 2000.

D Roche, A Schinker, J Storey, C Humphris, and M Guelden. *The Speed of Light: The 1996 World Solar Challenge*. Photovoltaics Special Research Centre, University of New South Wales, 1997.

Jeremy Rutman. How to do a roll-down test. http://physics.technion.ac.il/~rutman/car/Roll-downtest.pdf, 2007.

Sharp-World. http://sharp-world.com/corporate/solar_car2009/, 2011.

Y Shimizu, Y Komatsu, M Torii, and M Takamuro. Solar car cruising strategy and its supporting system. *JSAE Review*, 19:143–149, 1998.

David Snowdon. Spreadsheet analysis of roll-down test data collected for SSIV in nowra, January 2011.

David Snowdon. Personal communication regarding performance of current integrator device during 2009 WSC, November 2011.

David Snowdon. Personal communication regarding performance of Sunswift strategy in 2009 WSC, October 2011.

David C. Snowdon. Hardware and software infrastructure for the optimisation of Sunswift-II. BE thesis, School of Computer Science and Engineering, University of NSW, Sydney 2052, Australia, November 2002. Available from publications page at http://www.disy.cse.unsw.edu.au/.

Goro Tamai. *The Leading Edge: Aerodynamic Design of Ultra-Streamlined Land Vehicles*. Bentley Publishers, 1999.

Irving S. Tjiptowarsono. Advanced data communications for Sunswift IV. BE thesis, School of Mechanical and Manufacturing Engineering, University of NSW, Sydney 2052, Australia, June 2011. Available from publications page at http://www.disy.cse.unsw.edu.au/.

Tritium Power Engineering Group. Wavesculptor 20 datasheet. http://www.tritium.com.au/products/TRI50/TRI50.016_Datasheet_v6.pdf, 2008.

Kenneth Wong and Chris Rizos. Surveying the Stuart Highway with GPS for the 1996 World Solar Challenge. *The Australian Surveyor*, 41(3):194–203, 1996.

World Solar Challenge. World Solar Challenge technical regulations for the 2011 event. http://www.worldsolarchallenge.org/files/78_paticipants_regulations_technical-regulations.pdf, 2010a.

World Solar Challenge. World Solar Challenge general regulations for the 2011 event. http://www.worldsolarchallenge.org/files/76_paticipants_regulations_general-regulations.pdf, 2010b.

World Solar Challenge. Wsc 2011 results. http://www.worldsolarchallenge.org/files/320_2011_veolia_world_solar_challenge_final_results.pdf, 2011.

Hank Zumbahlen. Basic linear design. Technical report, Analog Devices Inc., 2007.

Appendix A

Outline of existing *Sunswift* strategy software system

As mentioned in Section 2.6.1, the Sunswift team has developed a sophisticated strategy software system over the course of the last decade, and it was used once again during the 2011 WSC race. Its functionality is outlined in this Appendix, as a reference for understanding how it was used over the course of this thesis.

This description of the system draws from my own understanding of it after using it over the last year, as well as from personal communication with David Snowdon.

The front-end user interface to the strategy system is the *Scanalysis* program, which is implemented in Python. It manages communications with the solar car, and displays many streams of telemetry and strategy data.

The core strategy software is implemented in C to maximise its performance speed. It consists of two parts, a comprehensive simulator, and a basic optimiser.

The strategy software works by first reading in a speed profile file, which contains a potential speed which the solar car should follow at every location during the race. It then runs a simulation, and uses model data to calculate the following parameters for the modelled solar car along every point in the race:

- the power consumed by the motor;
- the power received from the array;
- the battery power, current, voltage and state of charge;
- the total time to complete the race.

These are printed to an output file, which can be read and displayed by *Scanalysis*.

The optimiser part of the software also has access to the output file, and checks to see whether the calculated performance is satisfactory. If it finds that the speed profile that was the input to the simulation could be improved, it edits the profile file, and the simulation is run again. The optimiser makes use of user defined optimisation algorithms to improve the speed profile.

Once the optimiser can make no more improvements to the speed profile, *Scanalysis* can be used to display the calculated results, which are now the strategy to be followed, along useful feedback from the solar car telemetry network.

Some of the most useful graphs displayed by *Scanalysis* for the purpose of strategy feedback are:

- the actual and calculated power consumption of the solar car,
- the actual and calculated array power;
- the actual and calculated battery state of charge.

This system also allows the speed profile file to be edited manually by the user, if they want to simulate the effect of running at a particular speed profile.

Appendix B

Current Integrator schematics

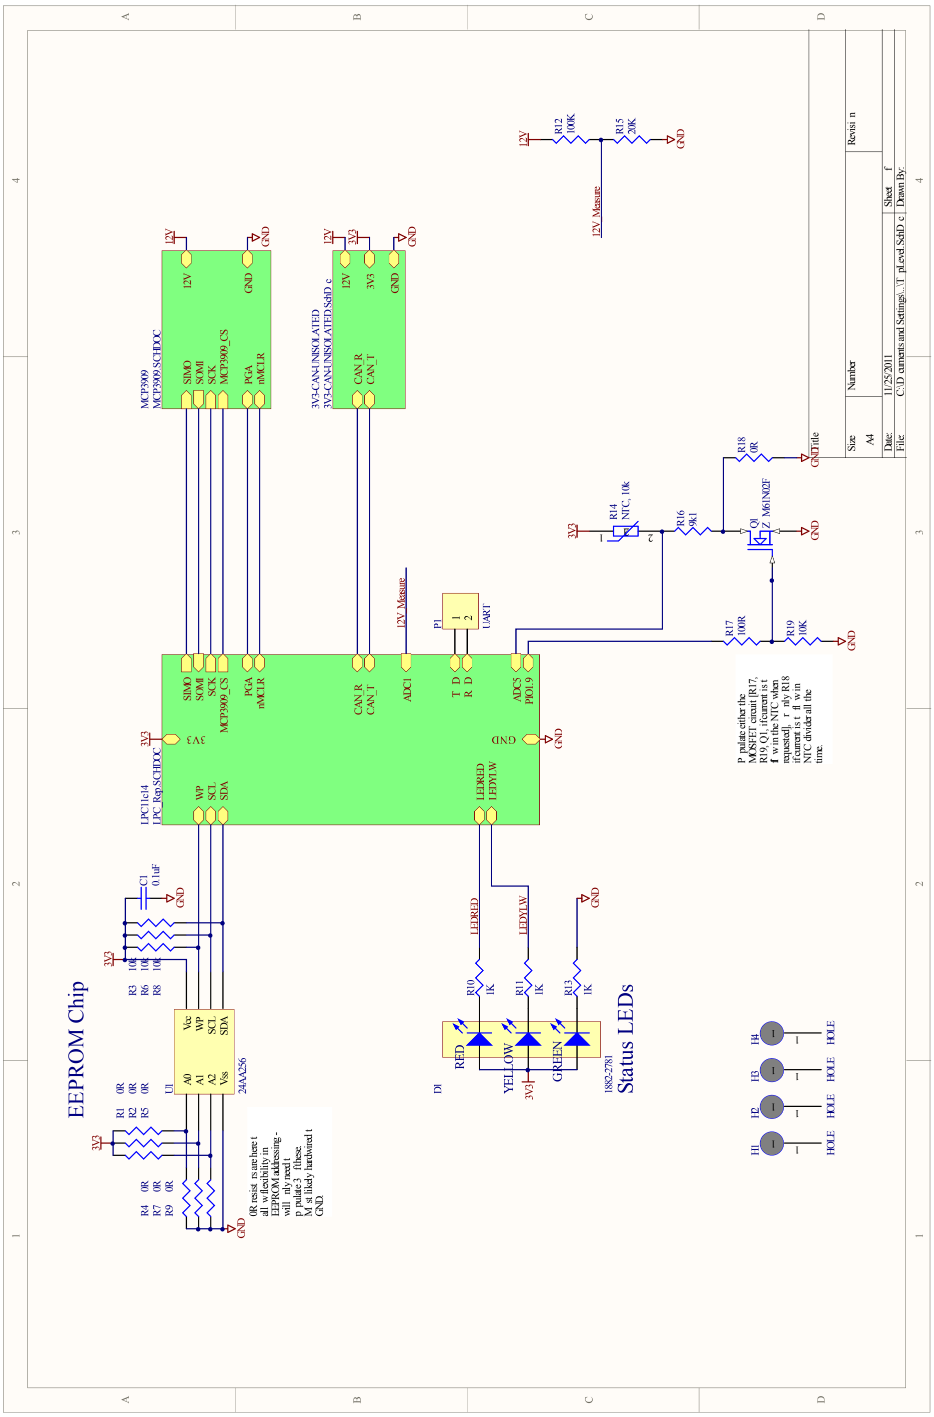
EEPROM Chip
Status LEDs

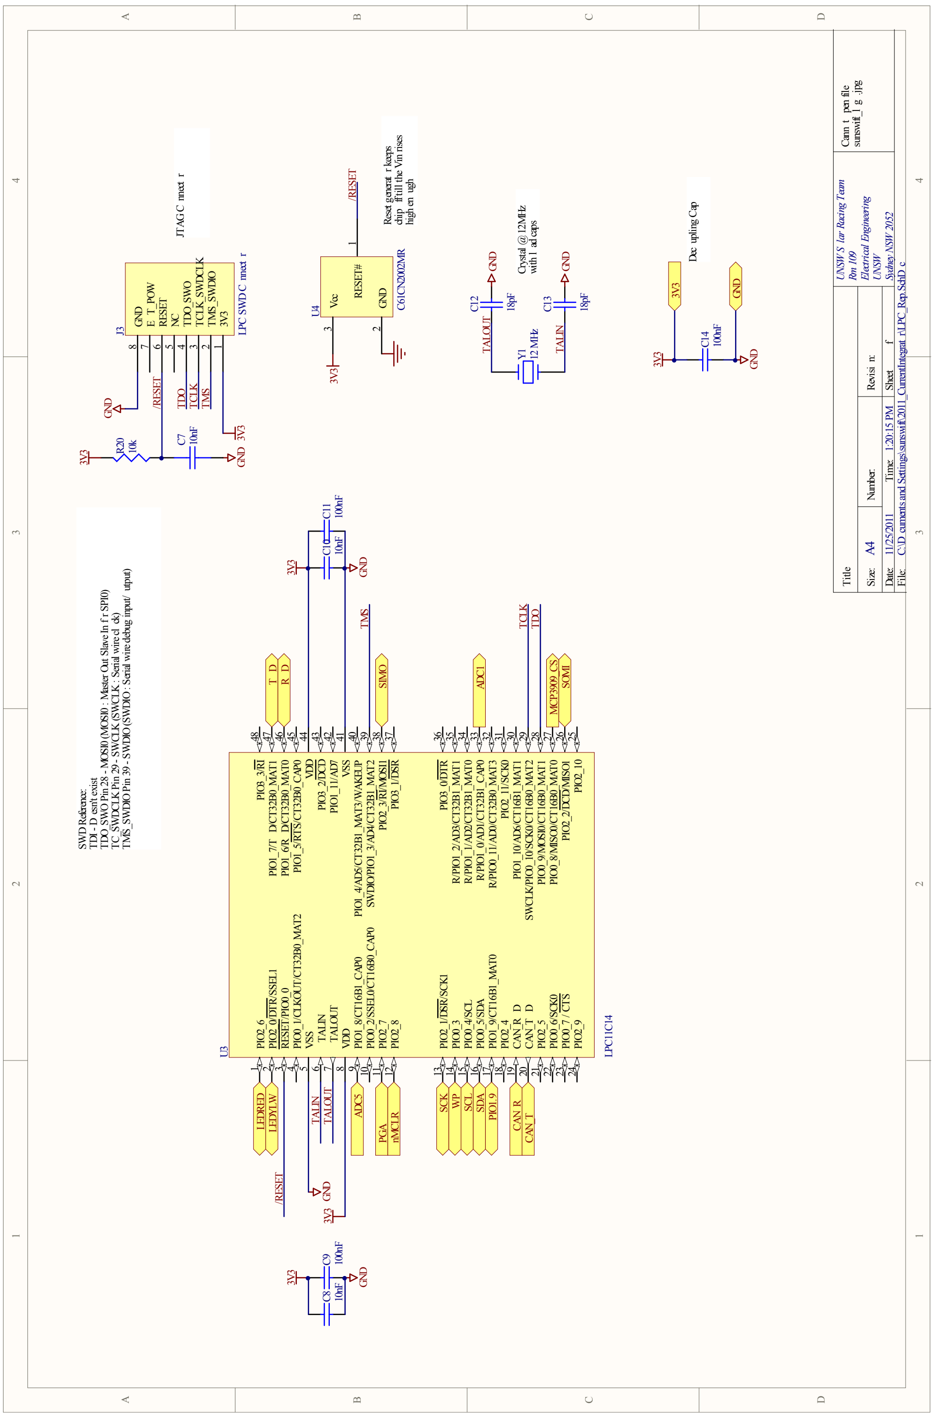

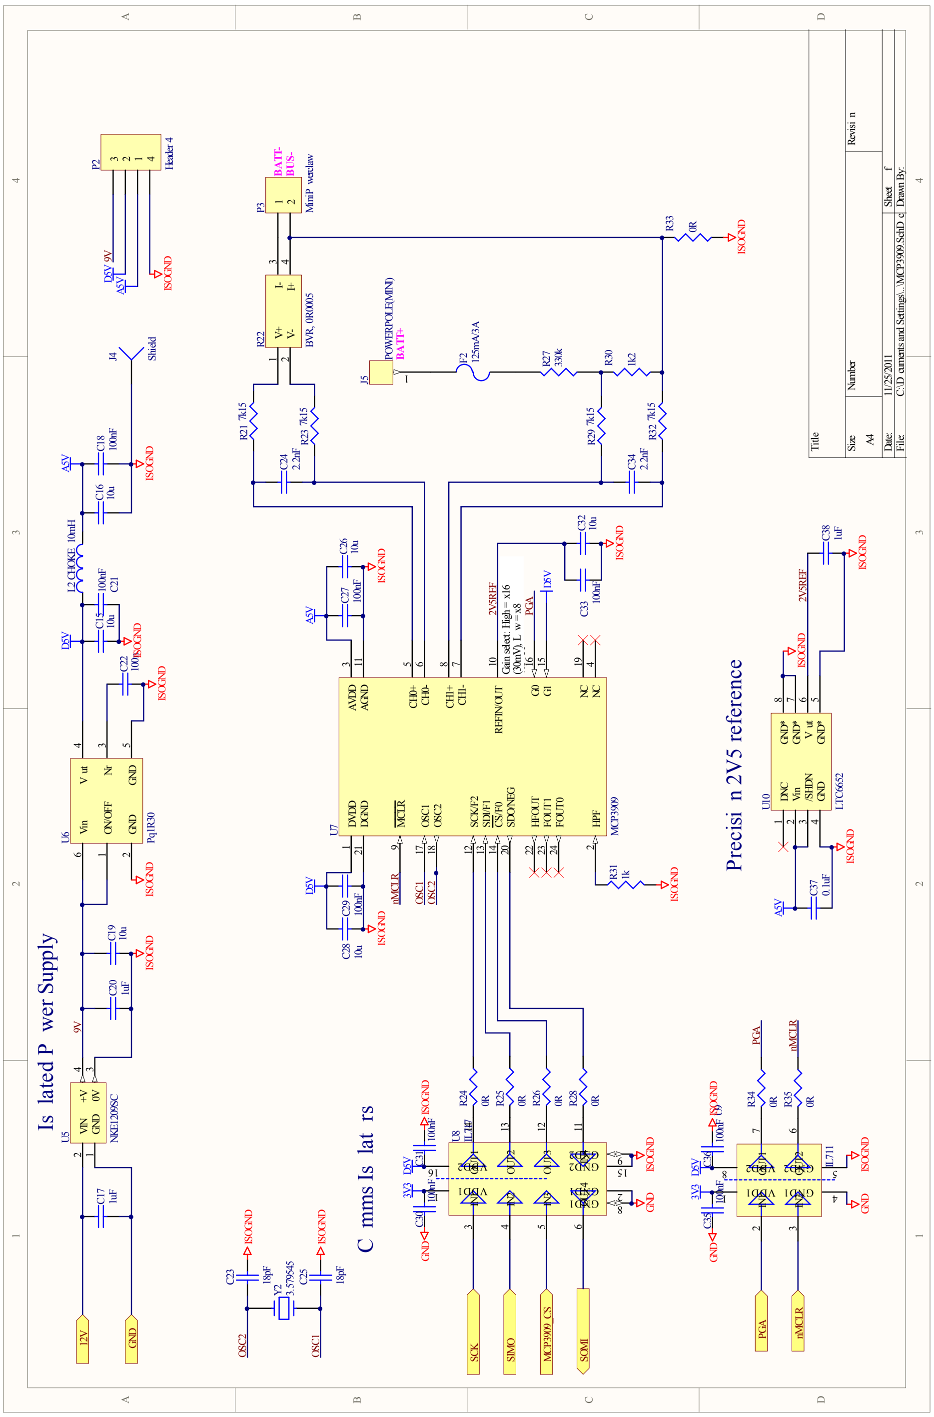
Is lated P wer Supply
C mms Is lat rs
Precisi n 2V5 reference
MCP3909
LTC6652
NKE1209SC
IL717
IL711
Title
Size
A4
Number
Revisi n
Date: 11/25/2011
File: C:\D cuments and Settings\...\MCP3909.SchD c
Sheet f
Drawn By:

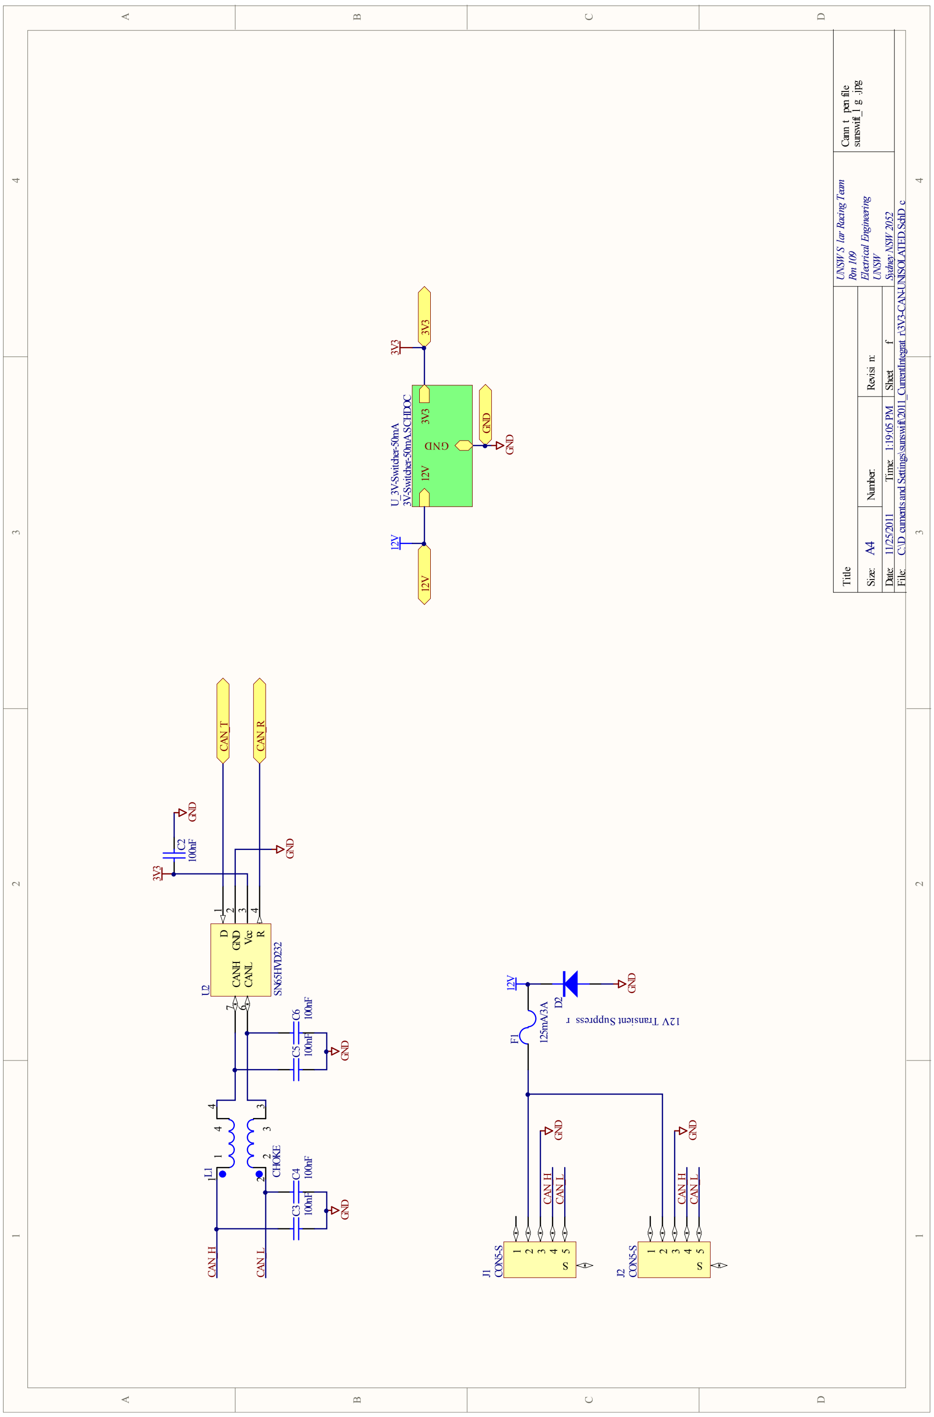
CAN_T
CAN_R
SN65HVD232
CHOKE
CON5-S
GND
3V3
12V
100nF
125mA/3A
A4
11/25/2011
1:19:05 PM

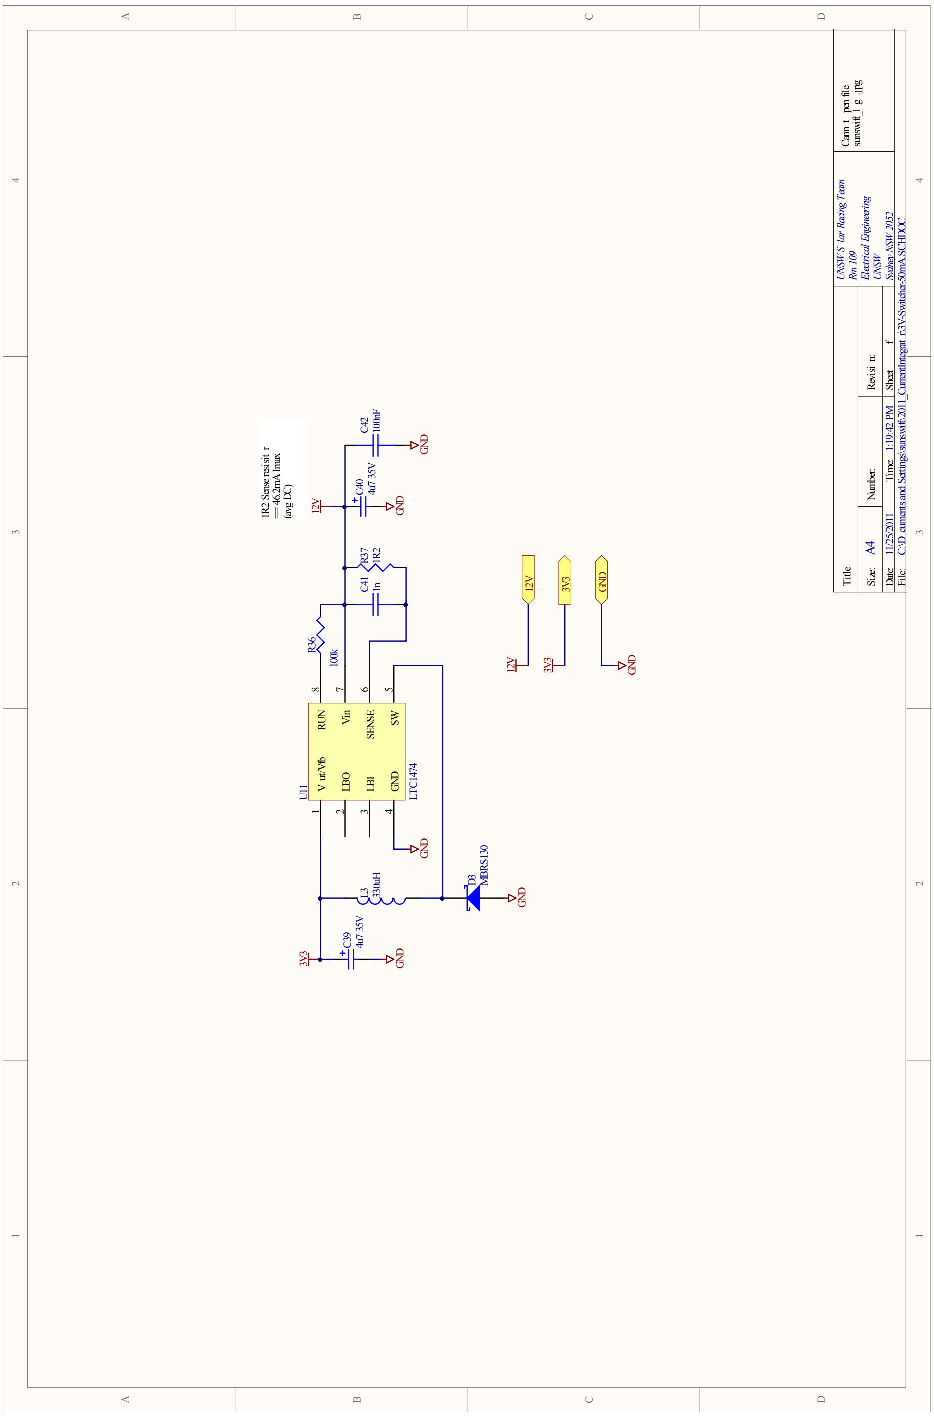
U11
LTC1474
RUN
Vin
SENSE
SW
LBO
LBI
GND
R36
100k
C41
1n
R37
1R2
C40
4u7 35V
C42
100nF
C39
4u7 35V
L3
330uH
D3
MBRS130
12V
3V3
GND
Title
Size:
A4
Number:
Date:
11/25/2011
Time:
1:19:42 PM
Sheet
File:
UNSW

Appendix C

Current Integrator PCB design

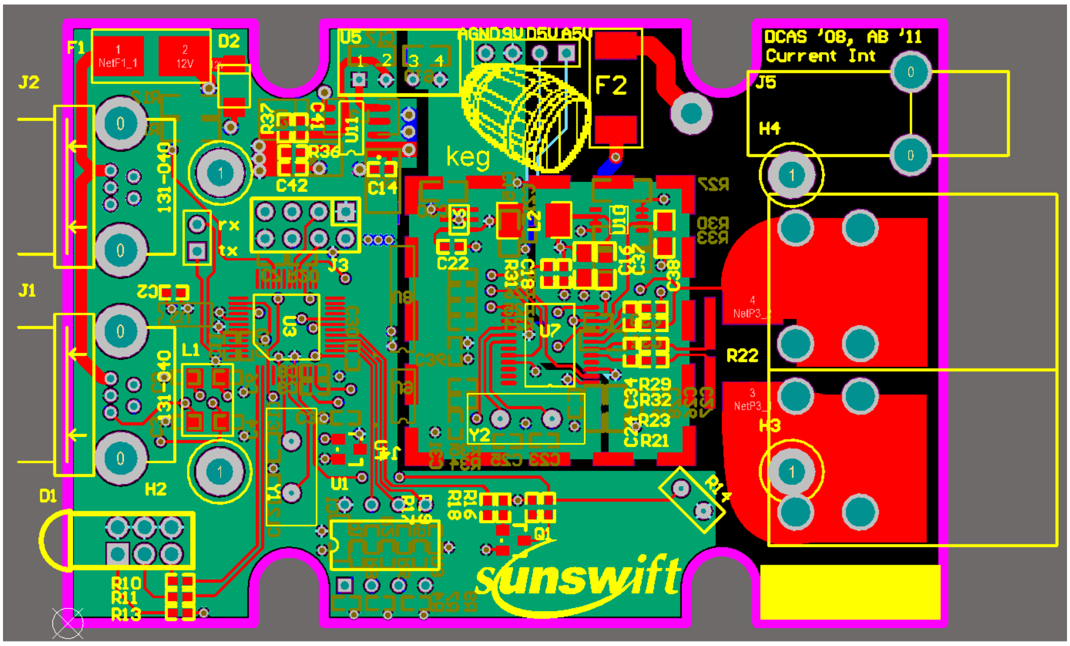

Appendix D

MPPT efficiency

Figures D.1 and D.2 show the efficiency maps for the Drivetek maximum-power-point-trackers (MPPTs) used in *Sunswift IV*.

The *Sunswift IV* solar array is split into three strings, which have open circuit voltages of approximately 75 V. The nominal battery voltage is 126 V. This gives a voltage transmission ratio of 126/75 = 1.7. From Figure D.2, this corresponds to an efficiency of 98% at an input power of 300 W per tracker. This corresponds to a total array power of about 900 W.

At the full rated array power of 1.3 kW, each tracker will be processing approximately 430 W, and from Figure D.1 we can see that the efficiency stays the same.

The theoretical efficiency of the Drivetek MPPTs used on *Sunswift IV* is about 98% at full array power.

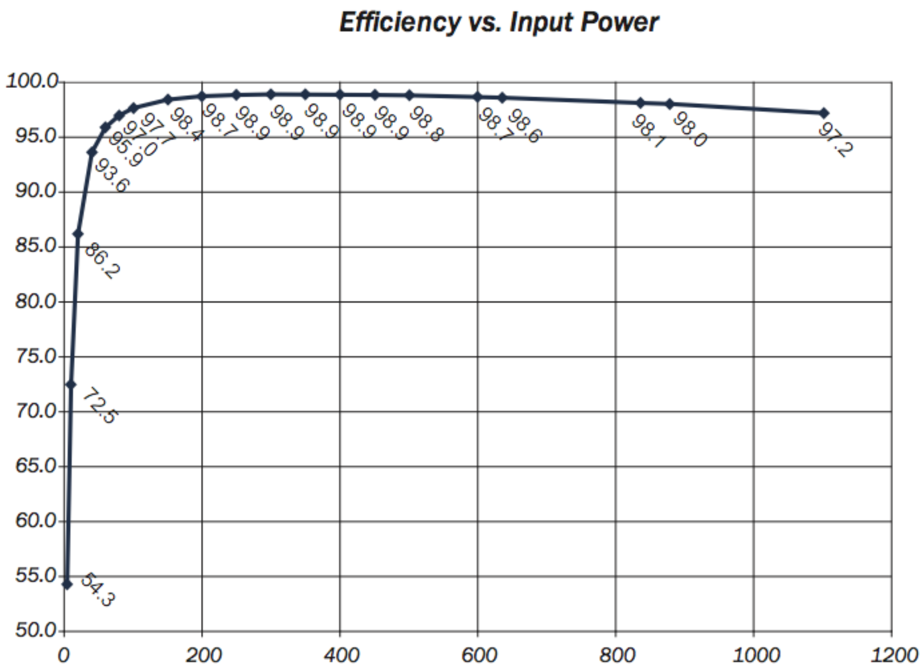

Figure D.1: MPPT efficiency (%) as a function of input power. *Images courtesy of Drivetek [Biel School of Engineering and Architecture]*

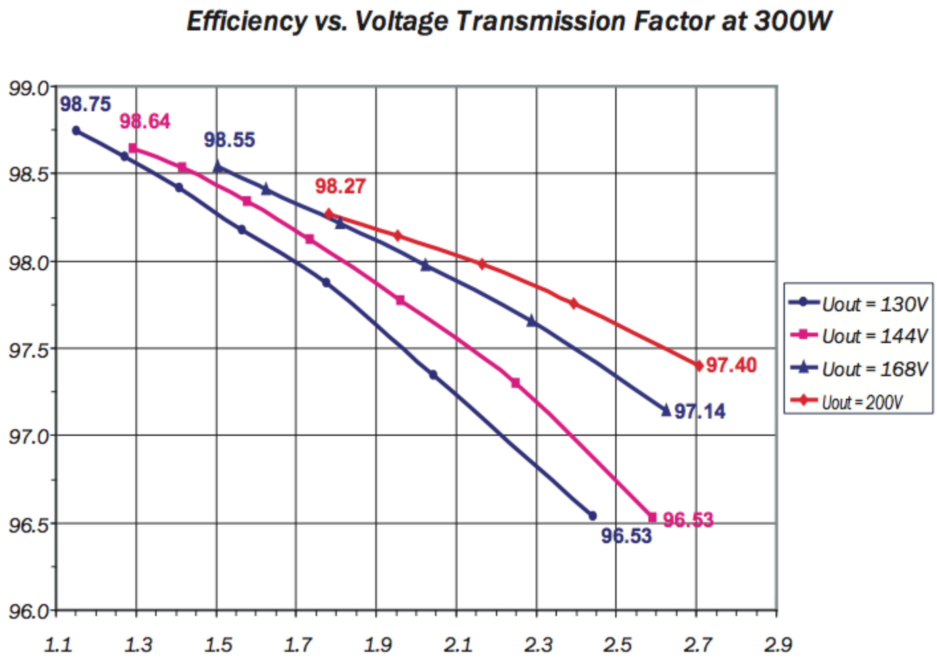

Figure D.2: MPPT efficiency (%) as a function of voltage transmission factor between the solar panel and the battery. *Images courtesy of Drivetek [Biel School of Engineering and Architecture]*

Appendix E

Drivetrain efficiency

Assume that *Sunswift IV* is travelling at 80km/h and using 1050 W of power (see Figure 2.5). From Figure E.1 we can see that the CSIRO motor efficiency under these conditions is 97.7%.

A vehicle speed of 80 km/h corresponds to an angular speed of 89 rpm, if a wheel diameter of 250 mm is used. From Figure E.2, this corresponds to an efficiency of about 97%.

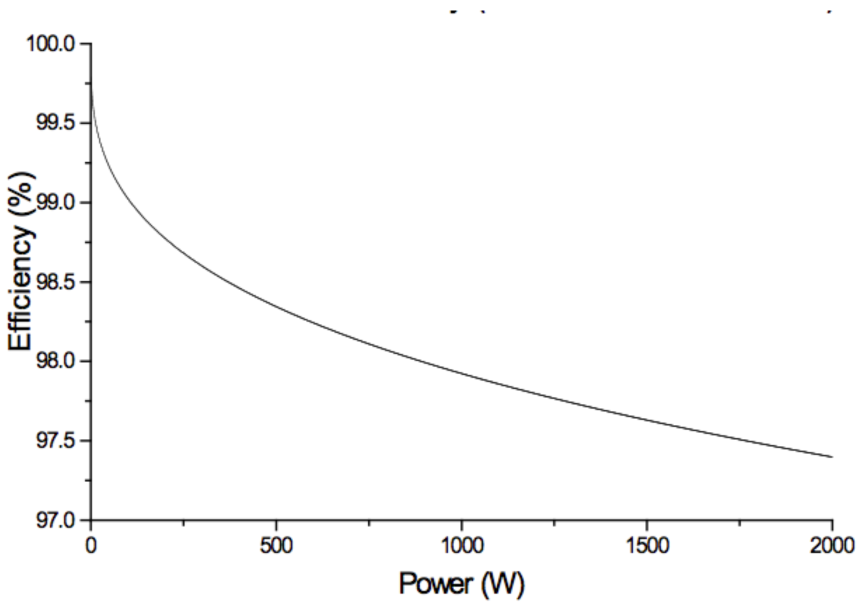

Figure E.1: CSIRO surface magnet wheel motor efficiency [Lovatt et al., 1997].

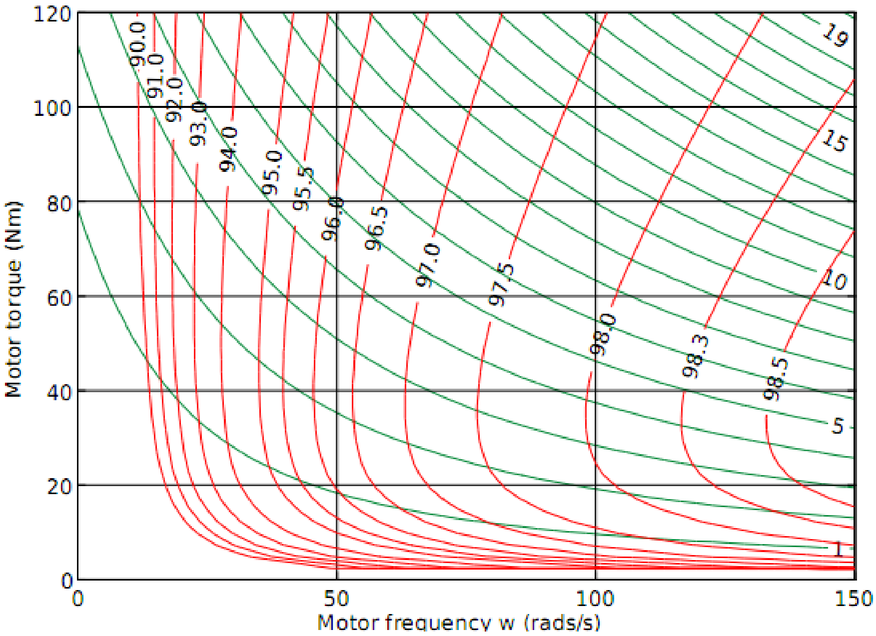

Figure E.2: Efficiency map of *Wavesculptor 20* motor controller using a 160 V input voltage and CSIRO surface magnet wheel motor. The red lines correspond to efficiency (%) and the green lines to power being processed by the motor controller. This plot also assumes a solar car set up with a 250 mm wheel radius and 275 kg vehicle mass. *Image courtesy of Tritium Power Electronics Engineering*

Appendix F

Weight distribution

The weight distribution on each of the wheels of the *Sunswift IV* solar car was measured as follows on the 30/08/2011 (view from top):

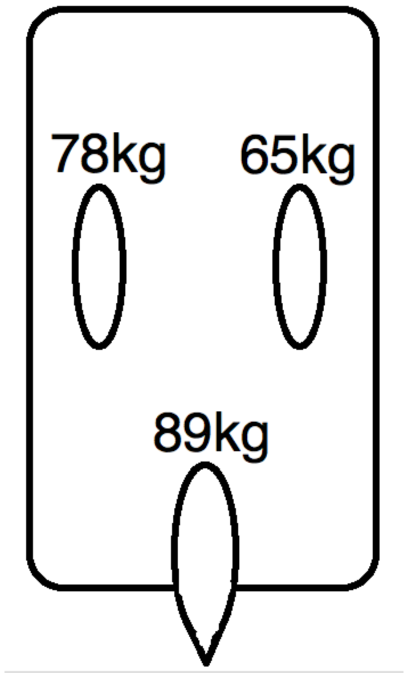

Driver: Tommy Heyser (70 kg)

Battery pack: 2009