1. Introduction

1.1 Antarctica

Antarctica is a cold, barren, windswept continent. The continent contains most of the world's fresh water, locked up in an immense ice sheet miles thick and extending for millions of square kilometers – a space larger than Australia. The coldest temperatures on Earth have been measured there, bottoming out at –89.2°C on July 21, 1983 at the Vostok research station [Guinness 2004]. The continent is almost entirely contained within the Antarctic Circle, the southern hemisphere equivalent of the Arctic Circle, at 66.5°S latitude. Thus, there are times when the sun never shines on the endless miles of snow and ice. At other times, the sun never sets, and merely wanders around the horizon for days on end.

Most of the interior of the continent consists of a large plateau that extends consists of some five million square kilometers (the right half of Figure 1-1), which is roughly the size of the continental United States. The plateau is a huge ice sheet several kilometers thick that contains

Figure 1-1 - A composite satellite image of the Antarctic continent. The vast central plateau makes up most of the right half of this image. Source: US Geological Survey.

most of the world's fresh water. The immense weight of all that ice has gradually pushed the rocky surface beneath it back down into the Earth's crust, so that some of it now sits below sea level. The surface of the plateau is relatively flat and mostly crevasse-free. The primary

surface features on it are *sastrugi*, which are dune-like waves of snow driven by the wind. At the local scale, these features are perhaps 30-40 cm high and occur ever few meters or so. The surface of the plateau, and the *sastrugi*, are quite solid, being packed into a crust by the wind, although softer drifts exist.

The plateau is the coldest part of Antarctica, so far from the relative warmth of the ocean, whose temperature even during the summer months is well below freezing (generally between –20 and –40˚C). It is also a desert, receiving less than 50
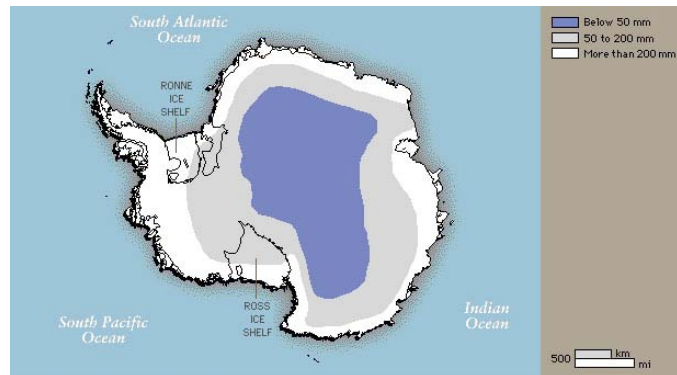

Figure 1-2 - An annual precipitation map of Antarctica. Most of the interior plateau receives less than 50 mm.

mm of precipitation (some 500 mm of snow) each year (Figure 1-2). The cold, dry air above the plateau is relatively calm and stable, averaging only 2 m/s during the austral summer [Valenziano and Dall'Oglio 1999] and with a 5-year maximum wind of 20.5 m/s at the Amundsen-Scott south pole station [CMDL 2004]. The cold air sinks, however, and rolls down the slopes of the plateau towards the coast. As it does do, it accelerates and circulates (due to the Coriolis effect), and occasionally blasts the coast with frigid gales called *katabatic* winds (also known as föhn, bora, or chinook in other parts of the world).

Antarctica is so cold and remote that it is the only land that humans never migrated or settled to. The Antarctic continent was, in fact, a land completely unknown for most of human history. It was not until the age of heroic exploration during the late nineteenth and early twentieth centuries that humans sought it out to try and discover its secrets. The history of this land is inextricably linked to legendary figures like Robert Falcon Scott, Ernst

Shackleton, and Roald Amundsen. Rather than finding indigenous peoples with a culture of their own, rich with adaptations that allowed them to survive, they found no one, and almost nothing but rock, snow, and ice [Mulvaney 2001].

Humans do inhabit the continent today, though only in a limited way. There are numerous research stations on the Antarctic continent today, operated by dozens of countries. Most of these research stations are on the coast , though a few exist in the interior. During the cold, dark winter, only a skeleton crew of technicians and a few researchers inhabit the continent. During the warmer months (from November to March), however, the population of the continent swells to thousands as scientists flock to it. The chief stations of interest are at McMurdo Sound, which is on the coast well south of New Zealand; and the Amundsen-Scott station, which is located almost directly on the south (geographic) pole.

1.2 Research in Antarctica

The researchers and technicians use these stations either to study Antarctica itself, or make use of its unique environment and location to conduct other kinds of science and engineering. Its location at high geographic and magnetic latitudes, and its exceptionally cold, dry, stable air makes Antarctica a perfect location for a host of research concerning the interactions between the atmosphere, ionosphere, magnetosphere, troposphere, thermosphere, and solar wind (a field collectively called "space physics" or "solar-terrestrial physics"). This research is conducted in a variety of ways using different instruments – from stationary sensor packages powered by diesel generators to one-time-use helium balloons. The dry valleys near McMurdo sound and the huge central plateau are excellent testing grounds for technologies that will ultimately be used on the surface of Mars and elsewhere. Research in Antarctica done at U.S. facilities or by U.S. personnel is overseen by the National Science

Foundation (NSF), though there are large number of agencies and institutions that fund or conduct research in Antarctica through them.

Some of the research conducted in Antarctica, especially that related to solar-terrestrial physics, requires a network of distributed sensors in order to capture the "big picture," just as the National Weather Service needs information from California to predict tomorrow's weather in Kansas. Such distributed networks exist throughout the world, but the greatest information and most interesting phenomena are observed at very high magnetic (and geographic) latitudes.

Figure 1-3 - An Automated Geophysical Observatory.

In addition to the instruments attached to the numerous manned stations on the continent, there are currently a half dozen Automated Geophysical Observatories (Figure 1-3) scattered on the Antarctic plateau. These operate year-round on approximately 50 W of power, which is provided through a combination of photovoltaic, diesel generator, and battery electrical power. The information taken from these instruments has already proven very useful to scientists for developing and verifying models describing the Earth's magnetic field and how it changes and evolves [Lanzerotti 1999]. The quality of this ground-based data is mitigated by the difficulty in its collection. Having a larger number of sensors and (ideally) the ability to reconfigure the network in response to evolving phenomena offers the potential for improved research and more accurate modeling. With this in mind, the National Research Council, in its research strategy for solar and space physics over the next decade, emphasized the need for "comprehensive new approaches to the design and maintenance of ground-

based, distributed instrument networks, with proper regard for the severe environments in which they must operate. [NRC 2002]"

Research in Antarctica, though invaluable to increasing our understanding of climate, geologic history, snow and ice properties, extraterrestrial exploration, and space physics, is very expensive and resource intensive. To place a scientist and his or her equipment at the South Pole is a complicated process whose logistics rival that of a small-theater war. Researchers traveling from the United States must first get to Christchurch, New Zealand. Then comes an 8-10 hour flight over the unpredictable South Pacific ocean to McMurdo station where, if they are lucky, their equipment has already arrived by sea.[1] If one is to continue on to the South Pole, another flight must be made along with one's equipment. Supporting these stations is a small army of workers who manage the logistics of moving people and equipment, feed, clothe, and house all of them; provide and maintain vehicles and transportation, etc.

1.3 THE COOL ROBOTS PROJECT

The installation and maintenance of instrument networks in Antarctica, either permanent, year-round stations or seasonal sensor packages, requires many flights of small aircraft to remote locations. Not only is this hazardous during take-off and landing, but it ties up precious human and equipment resources that could be allocated to other tasks, and results in a very high per-instrument deployment cost. The development of small mobile robots that can travel long distances and either carry or deploy various instruments offers a compelling alternative, which would substantially reduce the cost of doing research with distributed

---

[1] In the 2004-2005 research season, a major cog was nearly thrown into the works when an iceberg the size of Long Island threatened to cut off water access to McMurdo sound, which would have necessitated that all equipment and materiel be *flown* in, which would seriously have decreased the amount of research conducted.

networks in Antarctica. Since the fall of 2003, the Cool Robots Project at the Thayer School of Engineering at Dartmouth College has been developing one such robot for use in the harsh conditions of the Antarctic.

Designing a rover to travel hundreds of kilometers on the Antarctic plateau is a challenging task. The cold weather, even during the Antarctic summer, poses a significant challenge to both mechanical and electrical systems. The highly-variable terrain and long distances to cover present difficulties for mobility and navigation. In order to be a viable alternative to sending humans in airplanes, the robot must be relatively inexpensive to build and operate, and needs to be retrievable at the end of the austral summer. In order to be transported about the Antarctic continent (from McMurdo to South Pole stations, for instance), it needs to be relatively compact to fit inside a Twin Otter aircraft (Figure 1-4). Finally, in order to have the endurance

Figure 1-4 - The Twin Otter aircraft used for shuttling people and cargo around Antarctica.

to travel great distances and operate from several months, it needs to have a renewable energy source, with solar being the natural choice, as wind power has difficulties with mechanical breakdown [Plagge 2004], and other renewable sources are unavailable. In order to operate with little or no human intervention during that time, it needs to be able to guide itself autonomously and adapt to a variety of potentially hazardous situations.

The Cool Robots concept is to build a lightweight, solar-powered, four-wheeled vehicle capable of guiding itself over 500 km using Global Positioning System (GPS) waypoint navigation, within a time of two weeks to arrive on station,
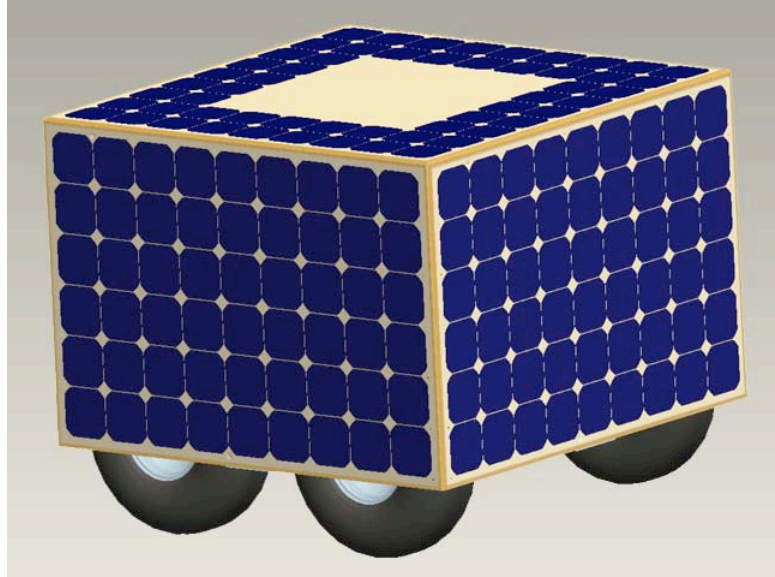

Figure 1-5 – CAD view of the overall Cool Robot design. The robot measures approximately 1.2 x 1.2 x 1 meters.

then collect data for a period of 2-3 months before driving itself back for retrieval at the end of the season. The design utilizes a low center-of-gravity chassis that has a thermally-contained volume for the electronics and batteries, outside of which is a larger box of solar panels for its renewable energy source (see Figure 1-5). The final robot must be able to fit inside of a Twin Otter aircraft's cargo hold and (with payload) weigh less than 90 kg. Table 1 provides a summary of the specifications to which the robot was designed.

The project began in the fall of 2003, under the direction of Professor Laura Ray (PI) at the Thayer School of Engineering at Dartmouth College, and in collaboration with Dr. James Lever at the U.S Army Corps of Engineers' Cold Regions Research

Project Specifications	
Vehicle Empty Mass	< 75 kg
Payload Capacity	> 15 kg
Ground Pressure	< 20 kPa
Materials Cost	< $20,000
Travel Range	> 500 km
Travel Time	< 2 weeks
Max Speed	> 0.8 m/s
Temperature Range	± 40°C
Max Dimensions (L, W, H)	1.5 x 1.2 x 1.2 m

Table 1 - Specifications for the Cool Robot.

and Engineering Laboratory (CRREL). Since the project's initiation, there have been a number of students who have made major contributions to the project. Beginning in the fall of 2003 and continuing through the winter of 2004, Guido Gravenkötter performed most of the mechanical paper design, and selected and tested the major components of the drivetrain. At the same time, Gunnar Hamman concerned himself with the selection of the microcontrollers, GPS unit, and laid the programming groundwork for operating the motors and accessing the GPS data. Meanwhile, Ben Kasdon, Ian Kahn, and Xianghui Weng began development of a model for predicting the amount of solar power available in a snowy environment like Antarctica, and selected a few of the components for the power system. In the winter through summer of 2004, Alex Price completed and refined the mechanical design, implemented it in the CAD software package Pro/Engineer, designed and manufactured the lightweight wheels, and developed the honeycomb construction technique. Dan Denton joined the project late in the spring of 2004, and focused on developing a software framework for driving the robot and laying the groundwork for the robot's communications. During this same timeframe, I was assisting Alex Price in fabricating many of the components for the drivetrain, developing the solar panel construction method, and constructing many of the circuits for the power system. In the fall of 2004 and through the winter of 2005, Götz Dietrich completed and tested the GPS navigation and waypoint following algorithm, while Toni Zettl focused on the communication systems and data acquisition. Finishing the chassis construction and installing the drivetrain assemblies was the responsibility of Götz, Toni, and myself.

The Cool Robots project will continue at least through the summer of 2005, culminating in a test field deployment at Summit Camp in Greenland. Plans are currently in

the works to extend the project through the winter of 2006, so that the robot can be tested for a season in Antarctica. Future plans include a multi-year second phase to the project, during which a handful of second generation of robots would be built and deployed, to demonstrate their use and effectiveness in creating distributed instrument networks in Antarctica.

CHASSIS CONSTRUCTION

In order to be lightweight, stiff, and provide an enclosed space, the chassis is made using honeycomb composite panels, rather than a metal space frame. Like corrugated cardboard, these panels derive their strength and lightweight characteristics by separating the two faces by some inner material. In the case of our honeycomb panels, the two skins are a thin fiberglass/epoxy matrix, and the inner separator is a vertically-oriented honeycomb of Nomex (see Figure 1-6). For a comparable thickness, these panels are as stiff as solid aluminum, yet weight just $1/10^{th}$ as much [Gravenkötter and Hamman 2004].
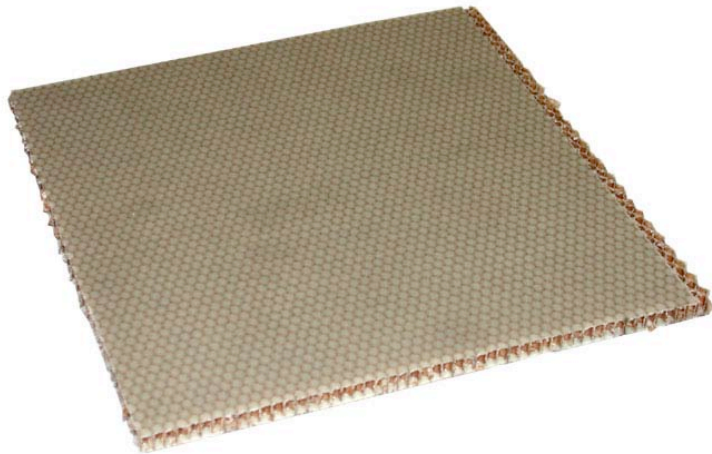

Figure 1-6 - A sample of 1/4" thick Teklam honeycomb panel material, which was used throughout the robot construction.

The chassis consists of a low-profile box made of 3/8" honeycomb panels, and measures approximately 24 x 26 x 10 inches. The chassis consists of a bottom and top section, like a large shoebox. The bottom and top sections were each made from a single continuous sheet of honeycomb, and had their sides folded up and epoxied in place, rather than being assembled from many smaller panels. The folding is accomplished by removing a width of the face that forms the inside of the corner, leaving the outer face intact. When the

joint is folded, the Nomex honeycomb collapses like an accordion. Using an epoxy and (optionally) reinforcing the joint with angle
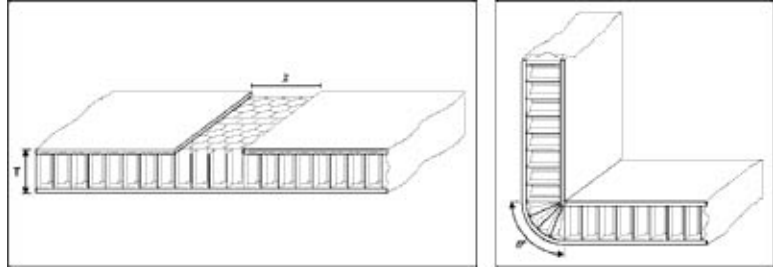

Figure 1-7 - "Slit and Fold" method used to create strong corners with the honeycomb panels. [Hexcel I]

aluminum, an incredibly strong and stiff corner can be formed. Figure 1-7 outlines these steps. The bottom section of the chassis also has several partition walls across its width that provide torsional stiffness and mounting locations for the four drivetrain assemblies. Figure 1-8 shows several Pro/Engineer CAD images of the chassis design. Further torsional stiffness is achieved when the chassis lid is in place and screwed to the chassis bottom.
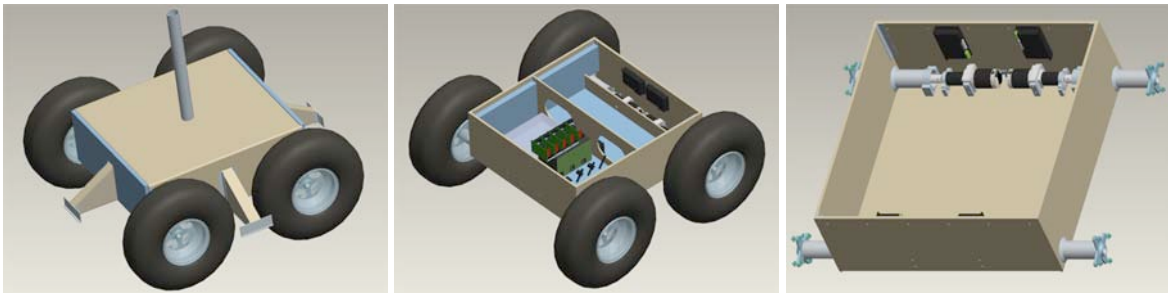

Figure 1-8 - Several Pro/E views of the chassis design.

The honeycomb panel material is not particularly resistant to crushing loads. As a result, one must specially prepare locations where screws and other mounting hardware will go. The team used a variety of bobbin-like inserts made of anodized aluminum that allow one to create a reinforced hole. The inserts come in standard-sized screw holes (#6, #8, 1/4", etc.), either tapped or clearance. These inserts consist of two halves that are press-fit together during installation. During the press-fit, the inserts slightly crush the honeycomb matrix between the flanges, so that the finished insert is flush with the surface of the panel. For

10

strength and to prevent tear-out, an epoxy is applied between the flanges and the skin of the panel during installation. Please refer to Figure 1-9.
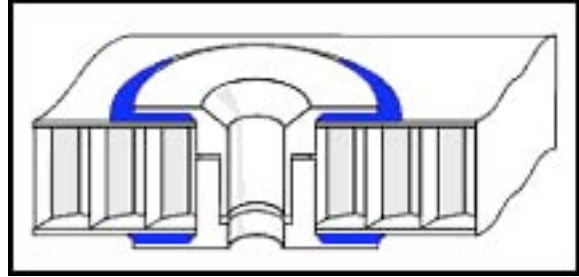

Figure 1-9 - Cross-sectional view of a hole-reinforcing insert in a honeycomb panel. [Hexcell II].

The selection of drivetrain components was an important task during the design of the Cool Robot. Calculations and approximations had to be made about the necessary speed and torque required for the robot to travel over the rugged terrain and cover the distance required. The tradeoffs between wheeled and tracked vehicles, active, passive, and articulated steering; etc. had to be considered thoroughly very early in the robot's design. Finally, the specific drivetrain components to meet the requirements had to be selected. For a full discussion of these design steps, the reader is directed to [Gravenkötter and Hamman 2004] and [Price 2004].

DRIVETRAIN

The final design utilizes four independently driven wheels, which steer the robot using skid steering. Each wheel is driven by an EAD DA23-GBB-M300 brushless DC motor coupled to a Neugart PLE40-100 planetary gearhead with a 100:1 ratio. These components have been outfitted for –50˚C operation. The motors have six permanent magnetic poles on the rotor and four wire-sound poles on the stator. As a result, these motors need to be electronically commutated, which is generally done with a separate motor controller. The motors are run by an Advanced Motion Controls BE15A8-H servo motor controller, which has a variety of operating modes (principally torque or speed control) and a variety of feedback outputs (principally motor current and shaft speed). The motors are designed to be operated from a 48 V (max) power supply with a maximum continuous current of +/- 4.2 A,
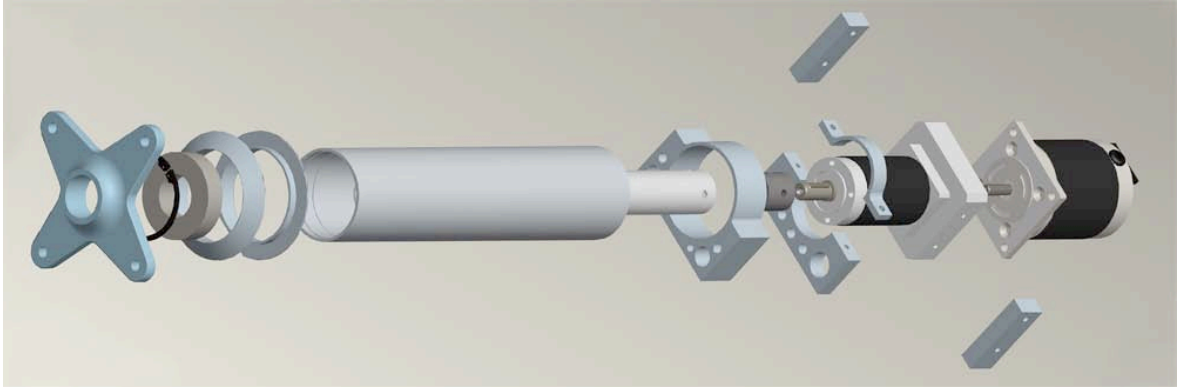
Figure 1-10 - Exploded view of the drivetrain assembly for the Cool Robot, minus the wheel.

while the motor controllers can operate over a far wider range. The gearhead has a 90% power transfer efficiency. Each drivetrain assembly has a peak electrical-to-mechanical efficiency of 75%, can provide on the order of 100 Nm of torque, and spin at up to 50 rpm. With 20" tires, that results in a maximum forward speed of 1.25 m/s, well in excess of the maximum required speed of 0.8 m/s.

The mechanical power is transferred to the wheels via an aluminum shaft that consists of a $1.0\varnothing$ x 1/16" tube, at the end of which is welded a hub. So as to minimize the load that must be borne by the gearhead, a bearing holds up the drive shaft very near to the hub. The bearing is in turn held in a larger, $2.125\varnothing$ x 1/16" aluminum tube, which transfers the load of the wheels to the chassis by means of a cradle at the inboard end and rings where the support tube passes through the chassis wall. The motor and gearhead are bolted together and form a single unit, which is in turn mounted to the chassis with several screws. Figure 1-10 shows an exploded view of this drivetrain assembly.

The robot uses commercially available ATV tires, which come in appropriate sizes for the robot, can be used with or without inner tubes, and have aggressive tread patterns for maximum traction over difficult terrain. The wheels these tires usually are mounted to, however, are made of cast or stamped steel, and are designed to carry far greater loads than

our robot will encounter. As a result, custom wheels were designed and manufactured in-house from 6061-T6 aluminum, thereby saving approximately 9 kg of weight (see Figure 1-11). These wheels were designed to withstand a 220 N static load (1/4 of the robot's weight), and an 880 N dynamic load, with appropriate margins of safety. Although fatigue failure is inevitable with aluminum (unlike steels, aluminum has no endurance limit), the wheels (and all the aluminum components) were designed to have an appropriately long cycle life. Although the wheels originally had 16x6-8 knobby turf tires mounted to them, the robot

currently uses the larger 20x6-8 chevron tires shown in Figure 1-11 for greater ground clearance, and possible better traction over snow. Because the drivetrain is mounted to the bottom of the chassis, the ground clearance is quite large for the robot's size – a full eight inches with the 20" tires. The tires are used without inner tubes (another substantial weight savings), and operate at an inflation pressure of 3-5 psi, although the robot is light enough (compared to an ATV) that they could be run flat.

Figure 1-11 - The custom machined aluminum wheel with a 20x8 ATV tire mounted to it.

This drivetrain has proven to be quite rugged and powerful during the testing phases of the Cool Robot. While running on battery power, the robot can climb stairs, up 30-45° inclines of packed snow, or over obstacles larger and less firm than the *sastrugi* expected in Antarctica. At the same time, it has a very low internal rolling resistance. Measurements made of the motor current under no-load and also during driving indicate that the internal

rolling resistance coefficient is approximately 5% [Dietrich and Zettl 2005], although complete data on this value over the whole range of the robot's operation is not yet available.

POWER

One must also contend with the rolling resistance of the robot as it is driving over snow. Ideally, with rigid wheels on a rigid surface, no power would be necessary to maintain a constant forward speed. However, since the tires deform a small amount, and the snow compresses under the weight of the robot, a continuous power must be applied to the wheels in order to maintain a constant speed. The rolling resistance of a vehicle in snow is attributed mostly to sinkage, or the compaction of snow immediately in front of the rolling element. The amount of sinkage depends on the size of the tire, the ground pressure, and the strength of the snow [CRREL 1995]. Based on empirical studies conducted by the U.S. Army Corps of Engineer's Cold Regions Research and Engineering Laboratory (CRREL), a conservative estimate of this figure, given the (maximum) 20 kPa ground pressure of our robot, is 0.25. This means that the 90 kg robot has a traction force requirement of 221 N over flat terrain (and greater if it is climbing). In order to traverse 500 km within two weeks, the robot needs to maintain an average forward speed of 0.41 m/s. Therefore, to maintain its forward velocity, our robot needs to supply at least 91.25 W of mechanical power to the wheels. However, the robot's maximum specified speed is 0.8 m/s, necessitating 180 W of mechanical power.[2] Taking into consideration the motor and gearhead efficiency, the electrical power requirement becomes 240 W. Adding an additional 40 W for the scientific payload and internal electronics, the power budget for the robot becomes 280 W. A more typical power requirement, based on field testing of the prototype, is closer to 200 W. The

---

[2] It is likely that the robot will drive at close to its maximum speed when it is driving, in order to make up for drift and dither around its intended path, and for the fact that it may have to stop from time to time, either to call home or to recharge its batteries. In this way, we hope to achieve the *average* speed needed to travel 500 km in two weeks.

large decrease in power is due partly to the fact that the estimated rolling resistance coefficient appears, based on preliminary rolling chassis data, to have been conservative [Dietrich and Zettl 2005]. Secondly, it is unlikely that we will always have the payload activated while the robot is traveling, or that it will continuously draw its full allotted power.

To satisfy the power requirement, the robot has a photovoltaic power system, augmented with lithium ion (Li-ion) batteries. Chapter 2 will describe in much greater detail the specifics of solar power in a snowy environment such as Antarctica, and describe the architecture of the Cool Robot's power system. In brief, the robot has a five-sided box attached outside of the chassis. The box has four vertical side panels, each measuring 47 x 32 inches, and a top panel measuring 47 x 47 inches. Each side panel is covered with an array of 54 solar cells, wired in series, arranged in a 9 x 6 grid (as shown in Figure 1-5). The top panel has 56 cells, also wired in series, arranged as a double-ring around its perimeter. Each cell is a SunPower A-300 single junction monocrystalline, single-junction silicon cell, which has a rated efficiency of at least 20%, and a typical output in full sunlight of 3 W. After being soldered together encased in a silicone encapsulant, each of the solar panels on the box has an efficiency of 18-20% (depending on operating conditions).

I custom designed the solar panels and manufactured them in-house in order to save particularly on weight but also on cost. As a substrate, the panels use a 1/4" thickness of the same composite honeycomb material as the chassis, rather than a solid aluminum or steel plate, which is common in commercial panels. Onto this substrate is spread a layer of silicone, onto which the cells are laid. After the final electrical connections are made, another layer of this same silicone is poured to cover the cells. A final top-coat of a more durable, enamel-like silicone is later added, which is textured to reduce the reflectivity of the

panel. Full and illustrated details of this assembly method are provided in section 3.4 and in Appendix C. This assembly method creates a panel that weighs less than 70 g per finished cell (3.5 kg/m^2), which is less than half the weight of commercially-manufactured panels.

At any given time of day during the austral summer, some of the solar panels are exposed to the direct sunshine of the sun. The other panels, although they receive no direct sunlight, still receive a considerable amount of insolation due to the reflection and scattering of sunlight from the snow, which has a very high albedo. The modeling and light experimentation performed during the design phase of the project estimated that, under typical Antarctic conditions, the panel box configuration could produce 330 W of electricity, which is more than enough needed to continuously drive the robot, even at its specified full speed of 0.8 m/s. Details of this modeling are provided in Appendix A.

Electrical Systems

The robot is crammed with electronics. A full discussion of some of the robot's electrical components can be found in [Gravenkötter and Hamman 2004] and [Dietrich and Zettl 2005]. However, and overview of various important pieces will be given here.

Microcontrollers – The brains of the robot consist of two Jackrabbit RCM3100 8-bit microcontrollers from the company Z-World (Figure 1-12). One microcontroller is responsible for commanding the motors, utilizing the GPS to navigate, communicating with the outside world via radio or satellite modem, and storing data to the datalogger.

Figure 1-12 - The Jackrabbit RCM3100 core module, which is the microcontroller used for the Cool Robot. [Z-World 2005]

This device will henceforth be referred to as the master microcontroller. The other

microcontroller is responsible for controlling the solar power system. Hence, I refer to it as the power microcontroller. The latter is slaved to the former, and communications take place between them using a glueless 8-bit parallel bus. These microcontrollers operate at 29.4 MHz and have a host of on-chip peripherals, including:

> *6 Serial Ports. Some of these ports are hardwired for a particular type of serial communication (RS232, for example), but mostly they can be programmed at will for a variety of different communication protocols (RS232, IrDA, SPI, and other synchronous/asynchronous modes). One serial port is generally reserved for programming, but can be reconfigured for different use.
> *Flash Memory, each RCM3100 has 512k, half of which is used for program storage, the other half is used for data storage with a file system.
> *Static Ram, each RCM3100 has 512k, used for run-time memory
> *54 General Purpose I/O (GPIO) pins, some of which are shared with the serial and other ports
> *Real Time Clock and ten 8-bit timers.

The Jackrabbits are rated to operate in a temperature range between -40 and +85˚C, and operate off a 3.3V supply. The two microcontrollers are mounted to a development kit that has breakout connectors for all the pins of interest, is prewired for master/slave connections, provides a battery backup for the Real Time Clock and SRAM, and provides both 5 V and 3.3 V power from a 8-24 V input. For bench testing, a 9 V "wall wart" power supply is used to power the evaluation board. When the microcontrollers are running in the robot, power is provided by the housekeeping supply, which is described in section 3.4.

Communications – While the robots are designed to be able to operate for long periods of time without human intervention, a means of communication with an operator is essential. The communications link is invaluable during debugging and testing, but is also necessary for transmitting a series of waypoints to the robot (the route it is to follow), sending back limited amounts of scientific data, sending back regular reports on the health and location of the robot, alerting the user to serious problems the robot has encountered and

17

is unable to correct itself, and allowing the user to assume manual control of the robot. Full details of the communications system are given in [Dietrich and Zettl 2005].
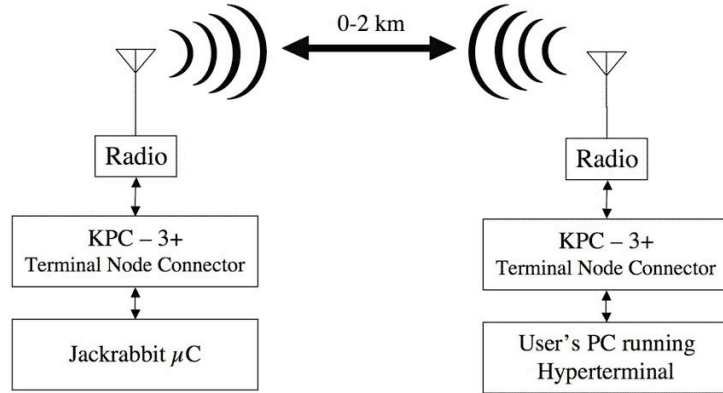

Figure 1-13 - Local communications using handheld radios.

Local communications are carried out using handheld radios: the master Jackrabbit communicates with a packet radio modem (a Kantronics KPC-3+) via an RS232 serial port, while the user's computer does the same with its own radio modem through the program Hyperterminal. The two modems talk to one another at 1200 baud, full-duplex, over a handheld radio link. The handheld radios currently used by the Cool Robots team have a reliable range of over one kilometer. This short-range communications link, given in Figure 1-13, is continuous, allowing the user to always monitor the robot and send it data or commands at any time.

Over long distances, the only suitable means of communications is over satellite phone using the Iridium system, which is shown in Figure 1-14. The robot will have an A3LA Iridium modem, which the
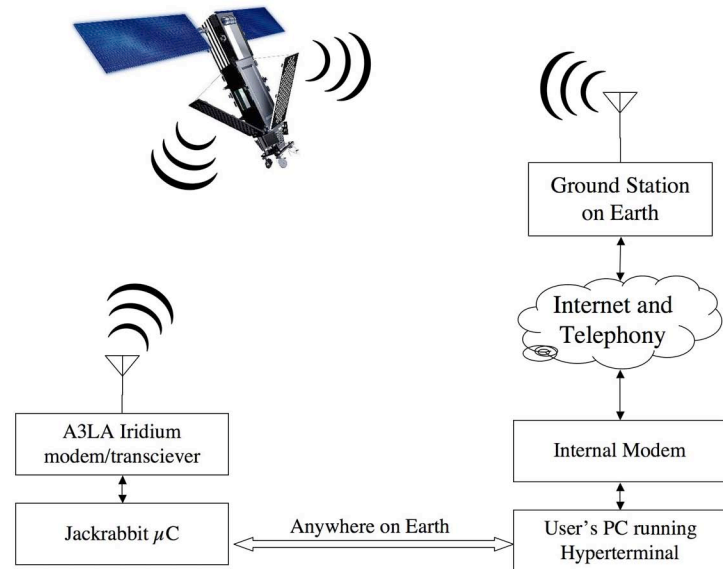

Figure 1-14 - Long-distance communications using the Iridium satellite telephone system.

Jackrabbit communicates with through an RS232 serial port, which would place/receive a call via satellite to/from the world's wired telephone system, and through that establish a peer-to-peer communication with the modem in the user's computer. The data communications take place at 2400 baud (quite low bandwidth by today's standards). The Iridium system is quite expensive, however, costing about $2/min within the United States and up to $7/min elsewhere. Therefore, this link will be used sparingly, and only when the robot requires the user's intervention, or when the user wishes to upload new commands or check on the robot's position and health. From the user's point of view, there is almost no visible difference between the two methods of communications.

NAVIGATION

One of the defining features of the Cool Robots project is the robot's ability to navigate autonomously over very long distances. The Antarctic plateau consists of large areas of flat terrain, interspersed with patches of rougher, *sastrugi*-strewn terrain. The interior of the plateau is essentially free of crevasses or other large obstacles. The low center-of-gravity of the robot, combined with its high ground clearance and powerful drivetrain, should allow it to overcome any reasonable obstacle one could anticipate finding on the plateau. That is to say, the team is confident of the robot's mobility.

Navigation is accomplished using GPS and a low-bandwidth path planning algorithm, which is described in full detail in [Dietrich and Zettl 2005]. The robot has a stand-alone GPS unit with an external antenna, which is configured on a computer using a special utility.[3] Once configured, the Jackrabbit can access the position data via an RS232 serial port. The data comes in the forms of long strings, and contains the position data in the form of latitude-

---

[3] The configuration information is retained using a battery backup, so it only needs to be done once.

longitude coordinates, as well as day and time, the current bearing, and how many satellites the unit is using. The robot already has a list of latitude-longitude waypoints provided by the user (which, from satellite image data, can be planned to avoid major obstacles). To avoid having to send lots of data to the robot, the waypoints are as little as several hundred meters (for local testing) and as much as tens or hundreds of kilometers apart (while in Antarctica). The navigation algorithm fills in the space in between waypoints with a number of basing points that are spaced at regular intervals along the route. The navigation algorithm then figures its current position and bearing in relation to the position and bearing to the next basing point – the difference between the two, the "off track" is an error signal (Figure 1-15).
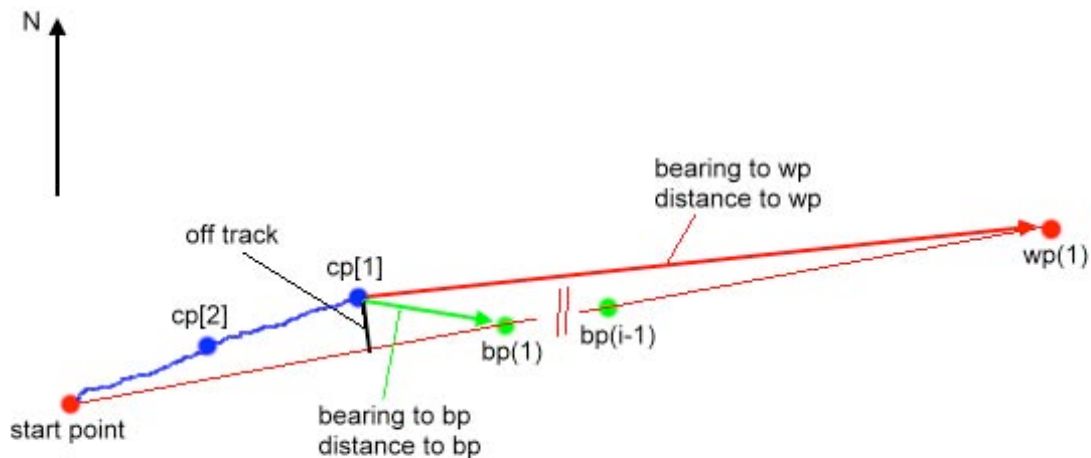

Figure 1-15 - Terminology and arrangement of the GPS navigation system. [Dietrich and Zettl 2005]

The bearing information from the GPS unit is not very accurate at the slow speeds the robot is traveling, and can vary by as much as ± 20˚. This error is caused by the fact that the bearing is determined by comparing the current position to the position one second ago. If one is driving in a car at 65 m.p.h., the difference between positions is large enough to determine an accurate bearing. However, when one moving at only 1 m/s (or worse, standing still) the relative error between two close position readings make an accurate determination of bearing difficult. A magnetic compass would likewise be ineffective at

such high magnetic latitudes, as the deviation of magnetic north from true north varies widely over the Antarctic plateau. Therefore, the bearing is instead calculated by comparing where the robot currently is to where it was 15 or 30 seconds ago.[4] The relative speeds of the motors are adjusted according to the error signal to try and keep the robot heading in the right direction. While the possibility exists to formally implement a compensator (using classical or modern control techniques) based on that bearing error signal, the team has found that low-gain, proportional, closed loop control is sufficient for the requirements of the project. If the bearing error is determined to be X number of degrees, then the robot will reduce the speed of the inboard wheels by 10%, while increasing the speed of the outboard wheels by 10%, and remain in that state for αX seconds, where α is the gain of the closed loop controller. Once the turning time has elapsed, the robot resumes a straight heading (i.e., the four motors turn at the same speed once again). The controller does not check to verify that the desired turn of X degrees has actually taken place, nor does it try to ensure that there is zero steady-state error between the robot's heading and the desired heading. Over the long distances involved, the speed (i.e., bandwidth) of this algorithm need not be great. After all, in Antarctica, a position error or tens or even hundreds of meters is of little consequence over the hundreds of kilometers to traverse. For local testing, one iteration is executed every 15 seconds. In Antarctica, the navigation algorithm may be executed as rarely as once every few minutes.

The accuracy of the autonomous navigation system was tested on a small scale on the frozen and snow-covered surface of Lake Mascoma in March of 2005. A short course of four waypoints spaced several hundred meters apart was created, with a total course length of

---

[4] This is one of the reasons why the algorithm needs to be low bandwidth – there must be a sufficiently large distance between the two points used to calculate bearing, otherwise the GPS errors in figuring position would likewise make this bearing information inaccurate.
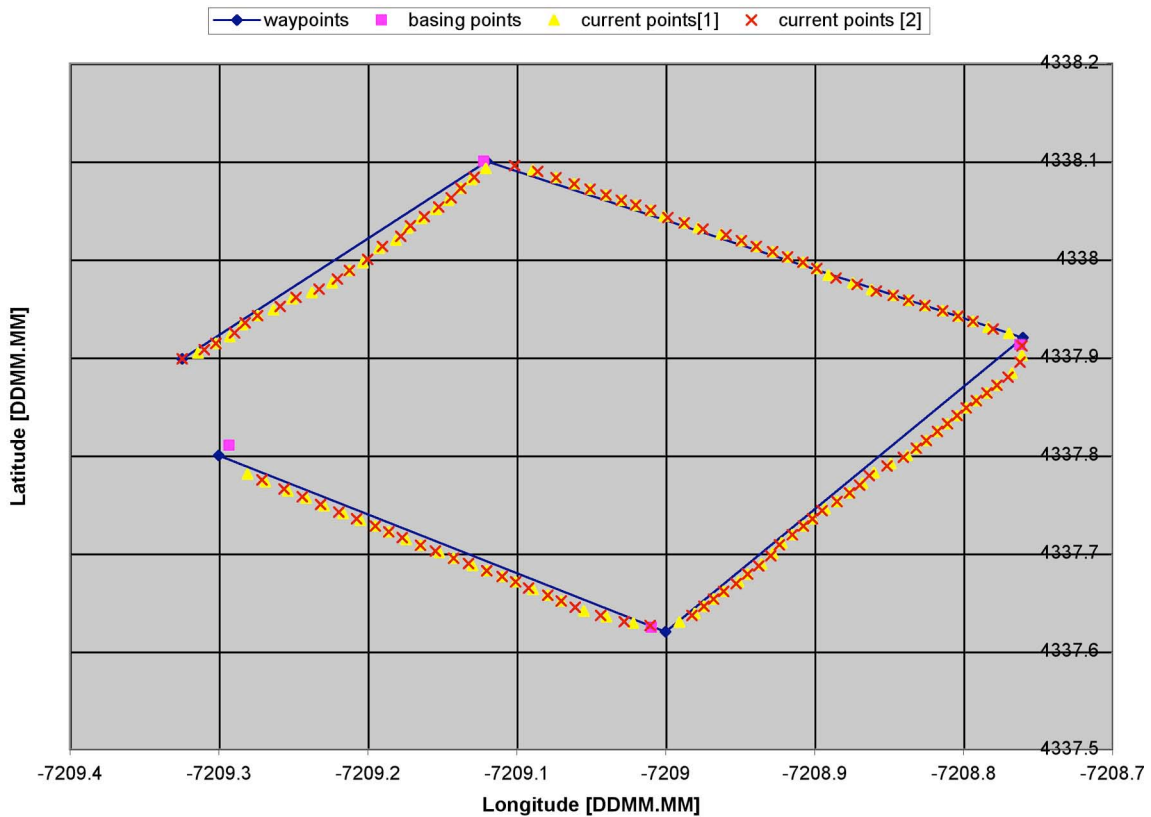
Mascoma Lake Short Course, 24 March 2005

Figure 1-16 – The short, 3 km circuit driven by the robot to demonstrate its waypoint following capabilities. The four waypoints were spaced several hundred meters apart. The data points on this graph are taken from the on-board GPS at 15 second intervals for two different runs through the course.

approximately 3 km. As one can see from the track shown in Figure 1-16, the robot was able to follow its course quite accurately, to within several meters. The greatest deviation from the desired track occurred, as one might expect, after the course makes a sharp turn.

The normal navigation algorithm can be interrupted by a variety of fault modes. One such fault mode could come from the power system's microcontroller, which is slaved to the master microcontroller, indicating that there is insufficient power available from the sun to continue driving at the given speed. The master microcontroller could reduce speed or choose to use the batteries to augment the solar power. In an extreme case, where the batteries are close to exhausted, the power microcontroller could cut power to the motors

entirely. Another fault mode would be indicated by a tilt sensor installed on the robot chassis, which monitors the pitch and roll angles of the robot. If the navigation algorithm detects that the robot is in danger of tipping over, probably due to running over a particularly large obstacle or parallel to a steep slope, then the algorithm will stop the robot, back it up, and change its route to try and circle around whatever obstacle there may be. A third fault mode could come from the robot's anemometer. Very rough calculations show that the robot, with its angular shape, and large cross-section, could be tipped in a 20-24 m/s (72-86 km/hour) wind. In the case of dangerously high winds, the robot will attempt to turn itself to place a corner of the panel box into the wind, thereby reducing its drag coefficient, and wait until the wind dies back down.

It would be possible to augment this navigation system and perform more rigorous obstacle-avoidance by adding a laser rangefinder and/or vision-based system. However, experience in Antarctica has shown that vision-based systems work poorly, due to the low visual contrast of the terrain [Apostolopoulos 2000]. In addition, the plateau itself is nearly free of obstacles, and careful route planning by the user can avoid the largest ones. What's more, the robot is traveling so slowly that it is unlikely to damage itself by running into a large obstacle. Finally, the added complexity and required computational capacity make this a prohibitively complicated choice for this modest robot. For instance, a system capable of using low-resolution stereo camera data to perform visual odometry or path planning would be able to run in realtime (at several frames per second) only with the computing resources of a full desktop computer, such as a multi-GHz Pentium IV processor. Due to the low bandwidth of the communications available to the robot, sending back image data from the robot would be time-consuming and, in the case of the Iridium communications, quite costly.

Based on the observations of Dr. Lever, who has been to Antarctica several times, most images would not even be worth viewing: a tedious expanse of flat white.

INSTRUMENTATION

The whole purpose of the Cool Robot is to be a general-purpose instrument delivery platform. As such, it has been designed to accommodate a variety of payloads, depending on a researcher's requirements. There is ample power available; 20 W is currently budgeted, but more could be available, especially when the robot is stationary. Also, the weight budget allocates up to 15 kg for scientific payloads. For the first field deployment of the Cool Robot, all of the on-board sensors that monitor the solar power input, weather, and motor output will be of great value in validating the design assumptions and formulating a better understanding of the environment the Cool Robot is operating in.

In addition to this wealth of data internal to the health and well-being of the robot, several test payloads will be carried. The first is a triaxial fluxgate magnetometer. This is a device that measures the absolute strength of the magnetic field around it, and breaks it down into x, y, and z components. There will naturally be a large DC component to this data, which represents the steady state of the Earth's magnetic field. This is the field that a magnetic compass will orient itself to. What is of interest to a space physicist, however, are the small fluctuations in this field, which are caused by a variety of phenomena.

A second test payload will be a dual-frequency GPS unit. Such an instrument, in addition to utilizing multiple GPS satellites to determine its location, also utilizes the different channels[5] that each satellite broadcasts from. Doing so can result in a more accurate determination of one's position (it is comparable to having twice as many satellites in view). However, since the different channels from a single GPS satellite are traveling at different

---

[5] Channel L1 at 1.6 GHz, channel L2 at 1.2 GHz.

frequencies, they will propagate through the atmosphere in slightly different ways. If one is able to work out the difference or delay between these two signals (phase shift), one can learn something about the
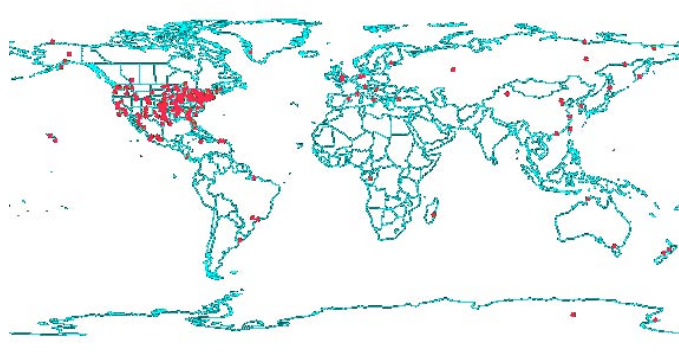

Figure 1-17 - Globally-located SuomiNet sites.

electromagnetic conditions along the paths that those two signals took from the satellite to the receiver. A distributed network of these specialized GPS units already exists across the United States, called SuomiNet [Ware et al. 2000]. The data from the 100 or so SuomiNet receivers (and a dozen or so others scattered across the globe, see Figure 1-17) is available to researchers in real-time over the Internet, and allows them to determine such things as the total electron content (TEC) or the amount of water vapor along the ray path between a satellite and the ground station, from which one can interpolate those conditions across the globe. With a greater network density, one can create more accurate maps of these phenomena and resolve ever smaller features within them.

1.4 Overview of Remaining Chapters

A description of the design, construction, and operation of the solar power system of the robot is the basis for this thesis. Chapter 2 discusses the theoretical aspects of photovoltaic power and describes the unique considerations for solar power in Antarctica. The chapter builds a theoretical model that predicts first the electrical characteristics of one solar cell under a wide range of operating conditions, attempts to expand that model to describe the operation of a full-sized solar panel, then works toward an estimation of the power available from the solar panel box of the Cool Robot in Antarctica. Chapter 2

concludes with a determination of the feasibility of using solar power for the Cool Robot in Antarctica, as well as the special considerations needed for the design and control of the power system.

Chapter 3 deals with the design of the architecture of the robot's power system. A brief comparison of this architecture to comparable systems is provided. A description of the physical construction and implementation of the power system takes up the bulk of the chapter, describing the custom-built solar panels, power electronics needed to condition the panels' output, the power requirements of the robot's electrical components, and the power distribution scheme built to fulfill those needs.

Chapter 4 will begin with an overview of various control methods associated with photovoltaic power. A discussion of the particular control challenges the Cool Robot's power system faces will follow. A description of an adaptive maximum power point tracking algorithm that forms an important part of the overall control algorithm is given, along with simulation results. A detailed explanation of the power system control algorithm, which I term "power matching," is given, along with a description of the other operating modes, and their implementation on the RCM3100 microcontroller using Dynamic C. The results of limited testing of this algorithm are also given.

Chapter 5 provides some conclusions about this work, and its success relative to what I have set out to do. Recommendations for improvement and future work leading up to a field deployment of the Cool Robot are also given.

Additional information related to the power system is given in the Appendices.

2. Solar Power and the Antarctic Environment

2.1 Solar Cell Physics

The photoelectric effect – the ability to generate electricity from light – was first discovered accidentally in the early 19[th] century [Fahrenbruch and Bube 1983]. However, it was not for until many years later, with the origins of quantum physics, that it could accurately be described in terms of solid state physics.[1] The first practical solar cells did not come about until 1954, when researchers at Bell Labs created 6% efficient cells from doped silicon [Tremble I]. Dozens of other acceptable materials and fabrication methods have since been discovered, with the highest performance to date coming from Gallium-Arsenide (GaAs), but the overwhelming majority of commercially-produced solar cells are made from some form of silicon – monocrystalline, polycrystalline, and amorphous. GaAs solar cells, though superior in performance, are considerably more expensive to manufacture and use than silicon, which is far more readily available (silicon makes up a very large percentage of the Earth's crust) and easy to work with (the experience and facilities of the entire semiconductor industry rest behind it). The following discussion describes the operation of monocrystalline silicon solar cells, though it is applicable to many different chemistries. The following discussion will assume that the reader is familiar with some of the solid-state physical properties of semiconductors, primarily the physics of a diode. For a good primer on semiconductor physics, the reader is directed to [Horenstein 1990].

A solar cell consists of little more than a p-n junction, like a diode. P-doped silicon has an excess of mobile positive charge carriers, called holes, which move through a lattice

---

[1] Albert Einstein received his one and only Nobel Prize in Physics in 1921 for his description of the photoelectric effect.

of nuclei that have one extra bound electron.[2]   N-doped silicon is the opposite – mobile negative charge carriers (electrons) in a lattice of fixed positive nuclei. The holes in p-doped silicon and the electrons in n-doped silicon are called the majority carriers, because they are the most prevalent carriers in their respective silicon.   However, there also exists a small number of free electrons in p-doped silicon, as there are free holes in n-doped silicon. These are called minority carriers, because they are usually outnumbered by about 1000 or 100 to 1. The overall electrical charge of both types is neutral, because for each mobile charge there is a fixed charge of the opposite polarity. When the two types of silicon are brought together, the mobile charges from the p- and n-doped sections are free to intermix. When an electron meets a hole, they cancel one another out – the electron fills the void that is the hole. The end result is an area devoid of mobile charges – the junction, or depletion region – where only the fixed, charged nuclei remain. On the n-doped side, these nuclei are positively charged, whereas in the p-doped silicon they are negatively charged. These fixed charges create a static electric field that pushes further charges away. The depletion region is (spatially) quite narrow, and creates a potential of approximately 0.7 V (for doped silicon). In order for an electron to travel through the junction from the n-doped region to the p-doped region, it must somehow traverse this potential. A small number will spontaneously acquire this extra energy (0.7 eV) from thermal excitation. In electronics, the extra energy comes from forward biasing the junction. When a diode is forward biased, electrons move freely from the n-doped region towards the p-doped region, and encounter and recombine with holes that are moving in the opposite direction, which is the flow of current.[3]

---

[2] In actuality, the holes are not physical particles, like an electron, but rather are used to signify the *absence* of an electron, but move through silicon as though they were physical entities.

[3] When the two recombine, a small amount of energy is released, usually in the form of waste heat. In a light emitting diode (LED), this energy takes the form of visible light.

In a solar cell, incoming photons give an electron in the valence band the required 1.12 eV bandgap energy, E_g, needed to jump to the conduction band, leaving a hole in its wake. The electrons and holes will diffuse from high concentrations in the silicon lattice to lower ones, trying to create an even distribution throughout the silicon. Some of these diffusing electrons and holes will recombine after a brief lifetime (on the order of 10^{-4} or 10^{-5} seconds) and release a photon of energy E_g, which is heat. Others, if they diffuse all the way to the junction, will be accelerated by the electric field created there by the fixed
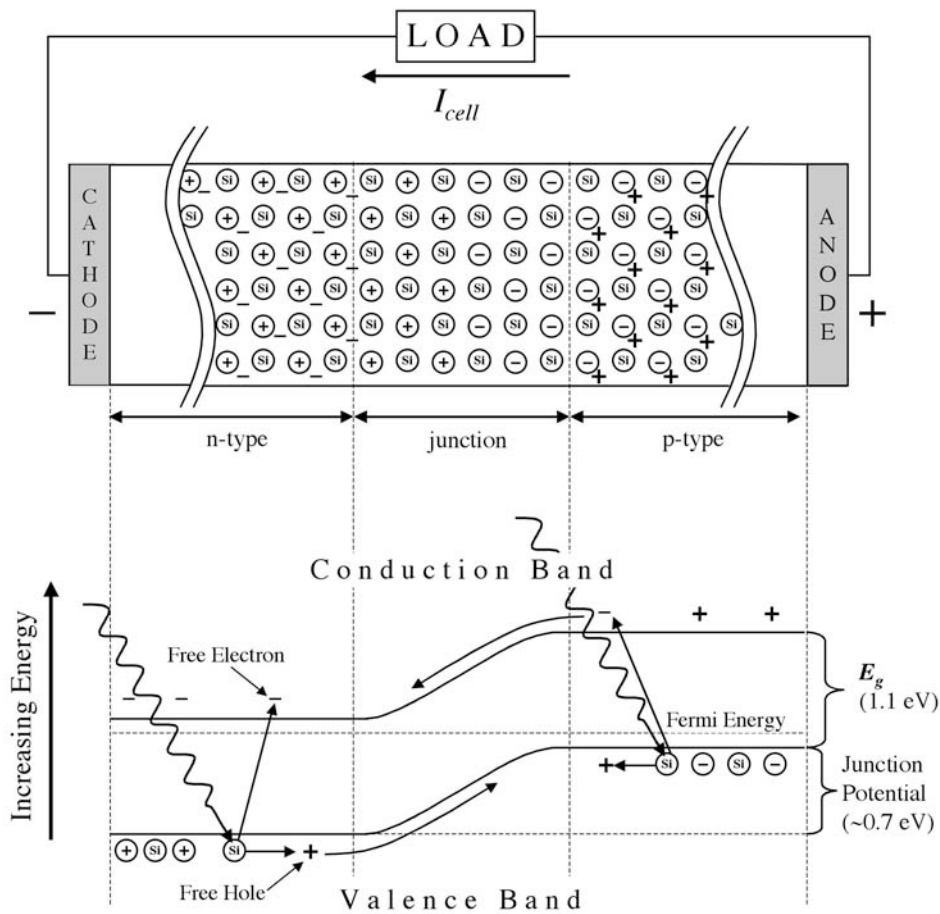

Figure 2-1 - A visualization of how photons interacting with silicon atoms in the n- and p-doped regions of a solar cell are able to create electron-hole pairs, which can drift to the junction and become separated by the static electric field. This image is somewhat oversimplified in the sense that the doping concentrations are greatly exaggerated (they are typically measured in parts per million, not the 1:1 shown above). There are other phenomena not pictured here, either, such as recombination, electron-hole pairs created by thermal excitation or in the junction itself, or what happens to the majority carriers.

charges. The electrons are accelerated towards the n-doped region (the cathode), and the holes toward the p-doped region (anode) [Tremble II]. The junction thereby acts like a collector of the liberated minority carriers. This effect illustrated in Figure 2-1. The motion of electrons in one direction and holes in the other is an electric current, which is referred to as the photocurrent. If the solar cell is connected to a load, the electrons can flow through the load to recombine with the holes on the other side. In this way, sunlight has been converted into electricity to do useful work.

2.2 EQUIVALENT CIRCUIT MODEL

A solar cell can be modeled, to a first order, as a current source in parallel with a diode (Figures 2-2 and 2-3). The current source outputs the photocurrent, and the diode is the intrinsic p-n junction diode of the solar cell. The current source provides whatever voltage is necessary to force the photocurrent, I_{sc}, through the load. If the required voltage is greater than the threshold voltage of the diode (about 0.7 V), the current is shunted through the diode instead of the load. The x-intercept of the accompanying plot is the open circuit voltage V_{OC}, and corresponds to an infinite load. The open circuit voltage, in this idealized
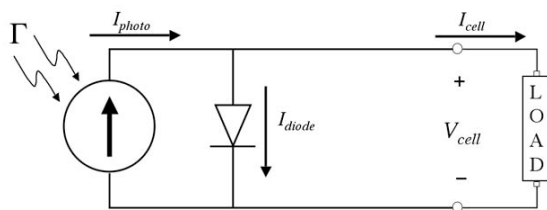

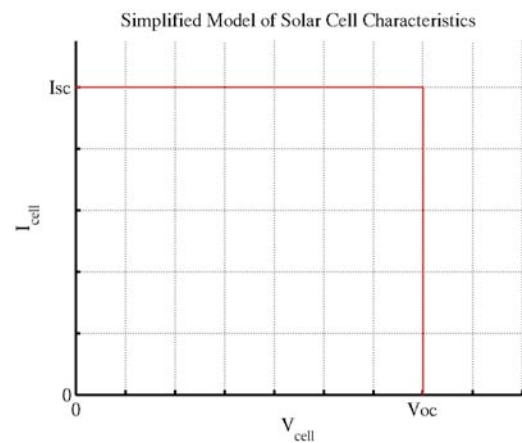*Figure 2-2 - A simple model of a solar cell.*

Figure 2-3 – The current and voltage characteristic of the model 2-2 when a very simple diode model is used.

30

model, is the 0.7 V diode voltage. No current flows through the device terminals and instead is shunted entirely through the diode. If the terminals of the solar cell are shorted together (zero load), the voltage across the diode (and the whole device) is zero, and so all of the photocurrent (I_{sc}) flows through the device terminals.

While this model is appropriate to a first-order approximation, it is far from a complete description. In reality, diodes do not simply turn on or off according to a threshold voltage. Instead, they obey the following (approximate) characteristic equation:

$$I_{diode} = I_0 \left[\exp\left(\frac{qV_{diode}}{k_B T} \right) - 1 \right] \qquad (2.1)$$

where q is one electron charge, V_{diode} is the voltage from the anode to the cathode, k_B is Boltzmann's constant, and T is the absolute temperature. The scaling factor I_0 is determined from the doping concentrations of the silicon, and is on the order of 10^{-11} or 10^{-12} amps. According to equation 2.1 and Figure 2-4, the diode turns on rapidly (but not instantaneously) in the *vicinity* of 0.7 V. As a result, the sharp corner in Figure 2-3 is smoothed over and more gradual, resulting in Figure 2-5.
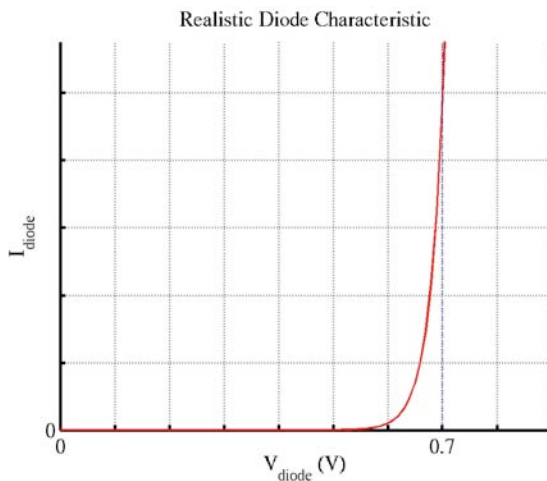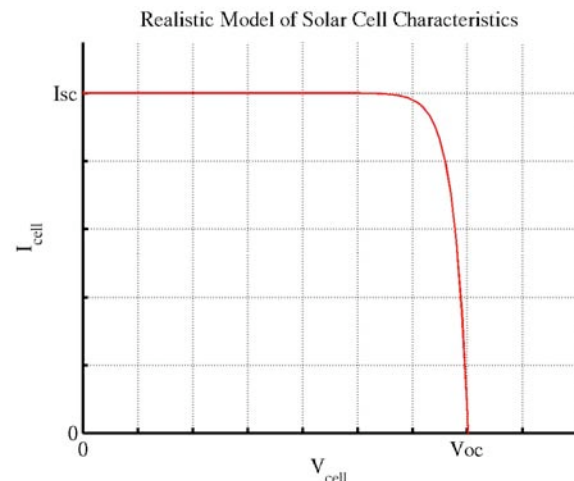

Figure 2-4 - Electrical behavior of a real diode

Figure 2-5 - Electrical behavior of a real solar cell, which I refer to as the I-V characteristic.

A further refinement of the model is the inclusion of a series resistance, which is in series with either the positive or negative terminal of the device. As explained below, the series resistance is a lumped parameter that includes the resistance between the cell and its metal contacts and the distributed resistance of the doped silicon itself. The effect of the series resistance is to further smooth the corner in Figure 2-5 and draw it towards the origin.

While the I-V characteristic is of great interest for modeling the behavior of the solar cell, what is also of interest (and a more useful measure of a cell's performance) is how the cell's *power* output varies with the terminal voltage. That is, the graph of *IV* as a function of
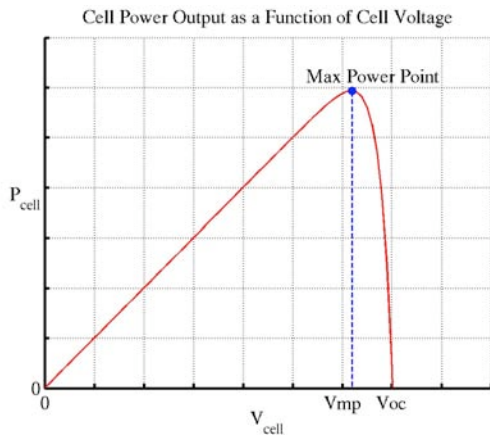
Figure 2-6 - A plot of cell power as a function of cell voltage, demonstrating the existence of a maximum power point.
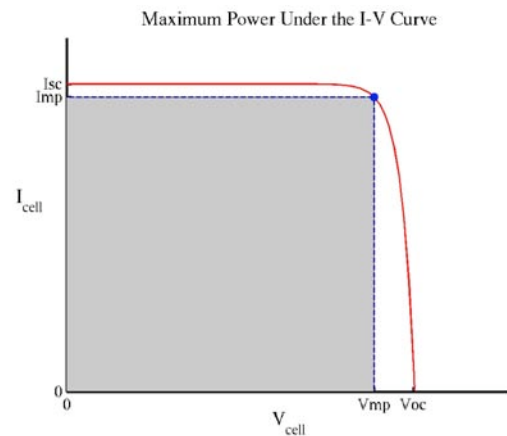
Figure 2-7 - An I-V plot of a solar cell highlighting how the MPP maximizes the rectangular area under the curve.

voltage. Such a graph is given in Figure 2-6. What should be immediately apparent to the reader is that the solar cell has a maximum power point (MPP), which occurs at some optimal (V_{mp}, I_{mp}) location on the graph in Figure 2-6. At this point, the area of a shaded rectangle contained under the I-V curve is as large as possible (see Figure 2-7). A similar (but not as accurate) definition of the maximum power point is where the slope of the P-V graph is zero. That is:

$$\left.\frac{dP}{dV}\right|_{V_{mp}} = 0 \qquad\qquad (2.2)$$

This second definition is true only of a photovoltaic system that exhibits only one maximum. In practice it is possible, due to how the cells are wired together in a panel and how the cells are illuminated, for there to be several local maxima. In that case, equation 2.2 is a description the location of each local maxima, but cannot by itself be used to distinguish from them which is the *global* maximum.

As one might expect, the shape of the I-V curve, and hence the location of the maximum power point, changes greatly with temperature and insolation. These changes will have important consequences for the design and control of the power system, as I will discuss in greater detail later. However, to briefly summarize the results:

1) The current of the solar cell is roughly proportional to the intensity of sunlight that falls upon it. Increasing sunlight will result in a proportional increase in the short circuit current of the solar cell. The open circuit voltage also increases, but only in a small way. The power output is likewise proportional to the insolation. The effect is demonstrated in Figure 2-8.
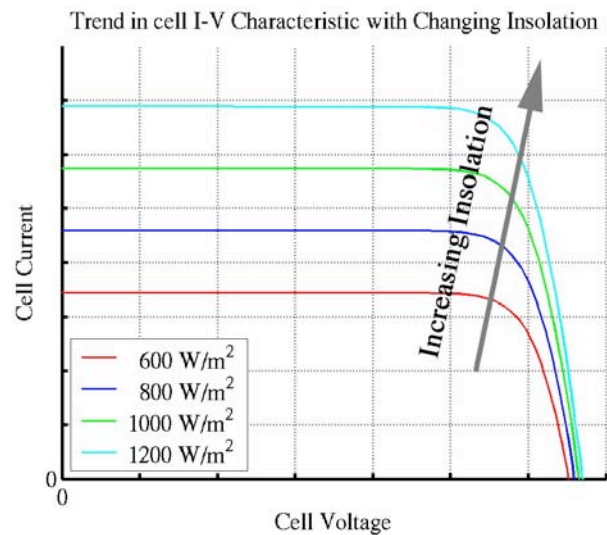

Figure 2-8 - The effect of insolation on the shape of the I-V curve, it also changes the location of the maximum power point.

2) The open circuit voltage is negatively proportional to the cell's temperature. That is, as the temperature increases, the open circuit voltage decreases. The effect is demonstrated in Figure 2-9. The current is directly proportional to the temperature, though this tends to be a much smaller effect. As a result, the

Figure 2-9 - The effect of changes in temperature on the shape of the I-V curve, which also has an effect on the location of the maximum power point.

maximum power output of the cell improves with cold temperatures.

2.3 More Physics and the Limitations Implied

From quantum mechanics, light behaves as both a wave and a particle. Photons are the particle manifestations of light waves, and have an energy directly proportional to their frequency (ν) or inversely proportional to their wavelength (λ):

$$E_{photon} = h\nu = \frac{hc}{\lambda} \tag{2.3}$$

where h is Planck's constant, and c is the speed of light [Eisberg and Resnick 1985]. Only photons with energy above the 1.12 eV bandgap energy are able to liberate an electron from the valence band (where it is bound to the nearest atom's nucleus) to the conduction band (where it is free to move about and conduct current). A photon energy of 1.12 eV corresponds to a wavelength λ_{max} of approximately 1100 nm, which places it in the near-infrared spectrum. Light with wavelengths shorter than this (and hence, of a higher energy) are likewise able to create electron-hole pairs. The light that reaches Earth from the sun has a

34

certain characteristic irradiation spectrum $\Gamma_0(\lambda)$. This spectrum out in space, called the Air Mass 0 spectrum, is filtered as it passes through the atmosphere. In the process, most of the ultraviolet is removed by the ozone layer,
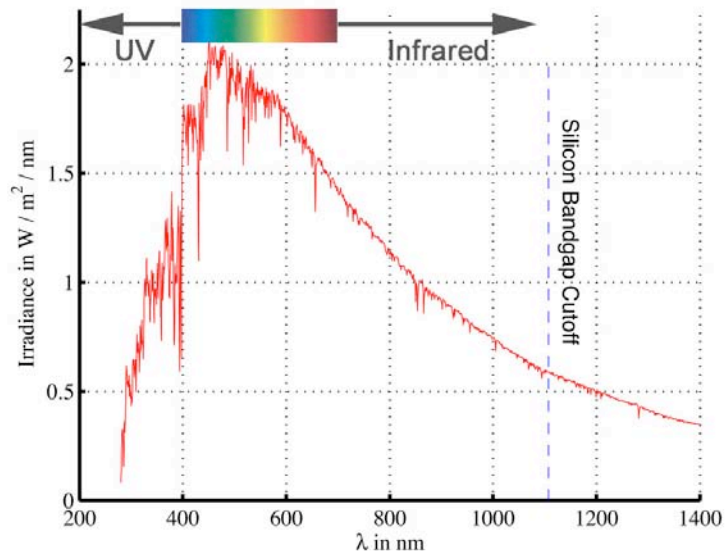
Figure 2-10 - The Air Mass 1.5 Spectrum, with the wavelength corresponding to the 1.12 eV bandgap energy of silicon marked.

and a significant portion of the blue light is scattered and reradiated by the oxygen and nitrogen in the air (which accounts for the blue color of the sky). If it passes straight through the atmosphere, (at it does when the sun is directly overhead) it becomes an Air Mass 1 spectrum. However, the sunlight that reaches the Earth's surface, particularly at high latitudes, takes a longer, slanting path through the atmosphere and becomes the more generally used Air Mass 1.5 or Air Mass 2 standard spectra [Sayigh 1991]. If one examines the spectrum of sunlight that reaches the Earth, as in Figure 2-10, which is the Air Mass 1.5 spectrum $\Gamma_{1.5}(\lambda)$,[4] one can see that the majority of it is concentrated in the visible spectrum (between 400 and 700 nm) and therefore is capable of generating a photocurrent. If one were to integrate the spectrum over all frequencies, one would get the total insolation by the sun on a patch of ground, Λ_s (in W/m^2). A typical value for insolation, under ideal conditions, is around 1000 W/m^2, or 100 mW/cm^2. That is:

$$\int_0^\infty \Gamma_{1.5}(\lambda)\, d\lambda = \int_0^\infty \Phi_{1.5}(\lambda) \cdot \frac{hc}{\lambda}\, d\lambda \equiv \Lambda_s \approx 1000\ \text{W}/\text{m}^2 \qquad (2.4)$$

---

[4] ASTM standard G173-03.

where $\Phi_{1.5}(\lambda)$ is the photon flux density (in photons per unit area per unit time) for the Air Mass 1.5 spectrum.

However, photons with energies in excess of the silicon bandgap energy E_g do not cause the solar cell to produce more energy [Chambouleyron 1991, Rodgers 2005]. Whatever energy is left over from raising an electron into the conduction band, $h\nu - E_g$, is wasted as heat. In this sense, a red photon is just as useful for creating electricity as a blue photon, or an ultraviolet one. One can express the converted power per unit area as:

$$\int_0^{\lambda_{max}} \Phi(\lambda) E_g \, d\lambda = J_{cell} \cdot V_{cell} \tag{2.5}$$

where J_{cell} is the current produced per unit area of solar cell. The unused power from photons more energetic than the bandgap energy is:

$$\int_0^{\lambda_{max}} \Phi(\lambda) \cdot \left(\frac{hc}{\lambda} - E_g\right) d\lambda \tag{2.6}$$

In the design of semiconductor solar cells (no matter what chemistry) there is a tradeoff to be made between the bandgap energy and the effective utilization of the spectrum. A low bandgap energy will allow more photons from the terrestrial spectrum to create electron-hole pairs. On the other hand, a lower bandgap energy would make worse use of the higher-energy photons at the other end of the spectrum. A high bandgap energy would make the most efficient use of the high-energy end of the terrestrial spectrum, but at the cost of ignoring all lower-energy photons. The theoretical optimum bandgap energy to optimally utilize the terrestrial spectrum lies somewhere around 1.4 eV [Tremble II], which is close to silicon's 1.12 eV. GaAs's bandgap energy is even closer, at 1.43 eV [EERE 2005].

There is therefore a theoretical upper limit to the efficiency of the solar cell, which is the ratio of power output from the cell and the power incident to the panel. For a silicon

solar cell, this theoretical limit is around 24% [Rodgers 2005, NREL IV]. Other materials, with different absorption spectra and bandgap energies, will have different theoretical efficiencies. Figure 2-11 shows the this theoretical curve as a function of bandgap energy, along with various common chemistries. For applications where the highest possible efficiency is required, such as in space satellites, triple-junction GaAs solar cells are often used. These solar cells are
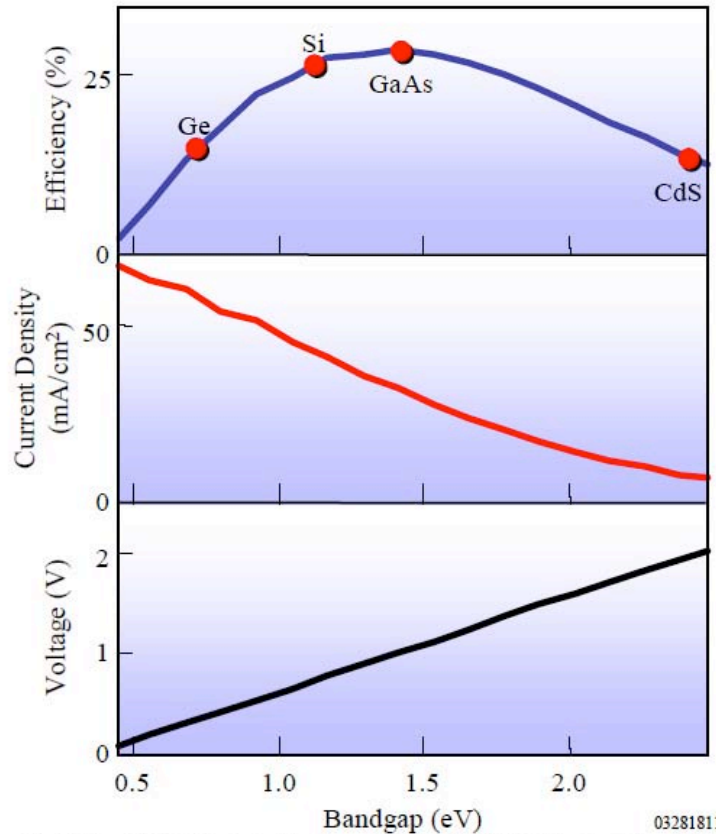
Copied from NREL, "The Basic Physics and Design of III-V Multijunction Solar Cells"

Figure 2-11 - Maximum theoretical efficiency as a function of bandgap energy. Silicon fares well at 1.12 eV, while Gallium-Arsenide performs nearer the theoretical maximum at 1.4 eV. [NREL IV]

called triple junction because, as the name implies, there are actually three p-n junctions in each solar cell, stacked on top of one another, which are tailored to make the most efficient use of different portions of the solar spectrum. One such triple-junction cell, which is used in solar concentrators, uses 1.8 eV GaInP as the top layer, followed by GaAs at 1.43 eV, and lastly Ge at 0.7 eV [NREL IV].

These are ideal, theoretical efficiencies. In practice, the efficiency of silicon solar cells are limited by a host of practical limitations [Tremble II], which I will briefly discuss here:

Shadowing by the Top Contacts – Once an electron-hole pair is formed and separated by the static electric field at the junction, the charge carriers must be collected by metallic contacts (the cathode and anode) if they are subsequently to do any useful work in an electric circuit. However, in most solar cells, the cathode contact, which collects electrons from the n-doped region, consists of a narrow grid of metal on the top surface of the solar cell. This grid overshadows what would otherwise be useful photon-collecting area, and so a sparse grid of thin wires is desired. A dense grid would be desirable, because it lessens the distance that electrons must travel through the doped silicon, which is resistive. The tradeoff between resistive losses (see below) and shadowing losses requires optimization before the solar cells are mass produced. Typically, the front surface contacts take up 7-10% of the usable area.

Reflection – All materials reflect or re-radiate some of the light that shines upon them, and solar cells are no exception. Light that is reflected off the surface of the solar cell (or whatever protective material covers it) is light that cannot be converted to electricity. This is a problem that can be mitigated through surface micro-texturing and anti-reflective coatings, which I will discuss later.

Conversion and Collection – Not every photon incident to the solar cell, even those above the bandgap energy, will create an electron-hole pair. Some, by running the obstacle course of atoms, can pass straight through without interacting. In general, shorter wavelength photons, such as blue and ultraviolet light, are absorbed nearest to the surface, whereas long wavelengths, like red and infrared light, penetrate deeper [Tremble II]. Longer infrared wavelengths, even some above the bandgap energy, have a good chance of sailing right through a thin (100-200 μm) silicon wafer. One can determine, either theoretically or

empirically, the probability that a photon of a certain wavelength will interact with a silicon atom (and create an electron-hole pair) in a wafer of a certain thickness. This is a measure referred to as the quantum efficiency, and is a function of wavelength (or frequency, or photon energy) for a given cell technology. A $QE(\lambda)$ of 1 (or 100%) at a certain wavelength means that every photon of that wavelength that enters the solar cell creates an electron-hole pair, whereas an efficiency of 0 will result in no electron-hole pairs being formed. As the wafer thickness increases, these probabilities approach unity (the wafer becomes opaque). One can also increase these probabilities by making the back surface of the cell reflective, so that photons that have passed through the silicon during their first passage are reflected back in for a second chance at interaction.

Next, not every electron-hole pair created will result in a useful current. Recall that, unless the electron-hole pair reaches the static electric field of the junction, it is doomed to recombine and produce waste heat after a brief lifetime (on the order of 10^{-4} or 10^{-5} seconds) of migrating through the silicon bulk. Therefore, it is important make the solar cell from a thin silicon wafer, so that minority carriers in the bulk are not far from the junction. One would think that it would be advantageous to make the depletion region in the junction very wide, which is accomplished by heavy doping. However, heavy doping is in many ways disadvantageous, as it reduces the average lifetime of the freed minority charge carriers (a heavy doping level means that there are many unpaired majority carriers that will quickly gobble up any liberated minority carriers migrating towards the junction). In practice, solar cells are manufactured to have a doping gradient so that the proper doping level exists at each portion of the solar cell (contact surface, bulk, and junction).

Lastly, not every electron-hole pair that has been separated at the junction is destined to reach the electrical contacts and do useful work as current in the electrical circuit. In order to reduce resistive losses and create a good contact between the silicon and metal contact (see below), the contact areas are heavily doped to make what the semiconductor industry called an ohmic contact. However, heavy doping decreases the average lifetime a minority carrier enjoys before it recombines with a majority carrier., which reduces the useful cell output.

Overall, one can measure the probability that an incoming photon of a certain wavelength will successfully create useful current, and is a combination of the various probabilities just discussed. This is called the *external* quantum efficiency, $QE_e(\lambda)$, and is one useful measure for comparing the performance of one solar cell technology to another [Mulligan 2004]. One would ideally wish for a solar cell whose quantum efficiency was 100% across the entire usable terrestrial spectrum. Current technologies are approaching this limit [Mulligan 2004]. A related measure is the spectral response, $SR(\lambda)$, which gives the current generated by a solar cell per watt of incident light of a given wavelength.

Voltage Factor – The maximum voltage created by the solar cell will always be something less than the bandgap voltage ($V_{cell} \approx 0.7$ V vs. $E_g = 1.12$ V for silicon). The open circuit voltage of the solar cell is determined principally by the difference in Fermi energy levels in the n-doped and p-doped regions of the solar cell. Recall that the bandgap energy is the energy an electron must gain in order to jump from the valence band to the conduction band. The Fermi energy in n-doped silicon is something less than the conduction band energy, and something greater than the valence band energy in p-doped silicon. Quantum mechanics state that these two energies must be equal, which results in the static potential difference at the junction (see Figure 2-1 above). The potential difference between

these two must therefore always be less than the full scale of the bandgap voltage, otherwise it would reside *in* the conduction or valence bands, depending on the temperature. This potential difference, times the photocurrent, is the power created by the solar cell, which will be something less than the ideal quantity described in equation 2.5 above. How close to one band or the other the Fermi energy lies is a function of dopant concentration (which can be tuned to some optimum value) and the temperature (more on this later).

Series Resistance – Silicon is, after all, a *semi*conductor, which means it will always have a higher resistivity, ρ, than a metal. Before the charge carries can reach the metal contacts and be conducted away, it must first migrate through the silicon bulk. The contact between the silicon and metal itself has a resistance (and can produce Schottky barrier effects, which is beyond the scope of this discussion). One can lump these two terms into a series resistance in the equivalent electrical circuit for a solar cell (see Figure A-1). The power lost to this resistance is $I_{photo}^2 R_s$. In order to reduce the series resistance, one must have a dense electrical contact grid (which introduces its own problems, as just noted). One can also reduce the series resistance by heavily doping the silicon (to n^+ or p^+) in the vicinity of the metal, creating an ohmic contact. However, as just described above, high doping concentrations can reduce the migration lifetime of minority carriers, so this presents yet another tradeoff situation.

Curve Factor (a.k.a. Fill Factor or simply FF) – Even without a series resistance, the I-V characteristic curve in Figure 2-5 (above) would still be blunted, and without a sharp corner between the constant current region and the constant voltage region. This is due to the characteristic equation of the diode, which we saw has a sharp, but not instantaneous, turn-on characteristic near 0.7 V. The diode characteristic equation is a function of the doping levels,

temperature, and intrinsic bandgap energy of the material. One can optimize these to good effect, but the I-V characteristic curve will never be as sharp as that given in Figure 2-3 (above). One can quantify how far off from that ideal the actual I-V characteristic is with a quantity called the curve factor [JPL 1976], which is defined as:

$$k_{curve} = \frac{I_{mp}V_{mp}}{I_{SC}V_{OC}} \; < \; 1 \tag{2.7}$$

The I-V characteristic, as already mentioned, is further smoothed by the presence of series resistance. As a result, the series resistance also contributes to reducing the curve factor. Typical values for k_{curve} range between 0.6 and 0.8 [JPL 1976, Mulligan 2004].

Self Heating – As can be seen in Figure 2-9 above, the power output of silicon solar cells have a negative temperature coefficient. That is, they are more efficient and produce more power from the same level of insolation when they are cold. However, as solar cells are far from 100% efficient, they convert the majority of sunlight into heat and correspondingly heat up, like asphalt. What's more, the solar cells are almost always encapsulated with some sort of protective layer, like glass, which traps the heat in. As a result, the temperature of the cell will often be significantly higher than the ambient air temperature, resulting in a lower efficiency than what one would expect for the given temperature. The problem becomes compounded when the panel is not operating at the maximum power point, as the wasted electrical power goes to further heating the cells.

2.4 THE SUNPOWER A-300 CELL

An enabling technology for the Cool Robot's project is the advent of (relatively) inexpensive high efficiency solar cells. Specifically, at the beginning of 2004, SunPower Corporation released a monocrystalline silicon solar cell that, through a variety of technological innovations, achieves an efficiency above 20%. Depending on process

variations, the efficiency can actually be as high as 21.5% [Mulligan 2004]. However, by avoiding many costly manufacturing techniques, such a photolithography, that had been used in a number of
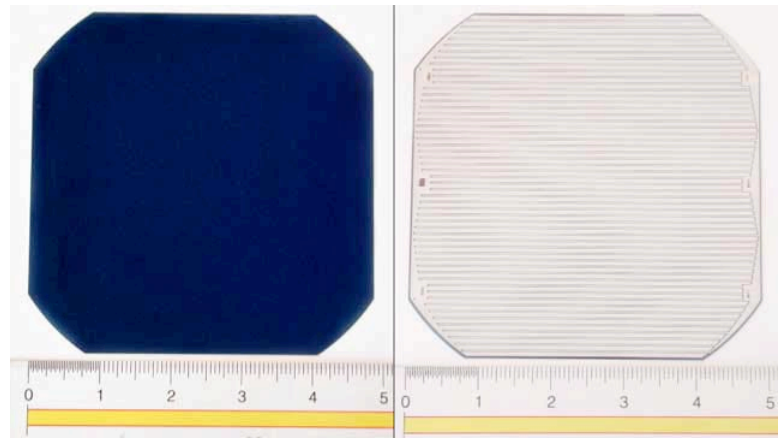

Figure 2-12 - Front and Back faces of the SunPower A-300 cell.

previous solar cell technologies, and by manufacturing in extremely large quantities, the cost per watt of power is relatively low, around $3. After a survey of various competing products [Kahn et al. 2004], the Cool Robots project selected the A-300 solar cell from SunPower as the best balance of efficiency and cost.[5]

The high efficiency of the A-300 comes about from a variety of technological innovations and optimizations. To begin with, the cell has all of its electrical contacts on the back surface (see Figure 2-12). Not only does this improve the aesthetic of the cell, but it means that the top surface is free of obstructions, which results in a 5-10% increase in usable cell area over older designs. Secondly, the top surface of the cell has been prepared in such a way as to capture as much light as possible. Because elemental silicon forms a diamond lattice crystal structure, it can be cleaved or etched along the <1,1,1> planes to create a pyramid-covered surface (see Figure 2-13). This surface texture causes incident light rays to enter the cell at a shallower angle than they otherwise would [McIntosh 2004]. The top surface of the solar cell is also covered with a thin layer of TiO_2 or SiN, both of which are

---

[5] Unfortunately, one factor that hadn't entered our consideration was availability. As with new products and manufacturing, there was a significant time until production was up to full capacity. In the meantime, demand had increased substantially. As a result, although we placed our relatively small order of 500 cells in March of 2004, it was not until the end of the year that we received the bulk of our order.

anti-reflective coatings.[6]
Beneath this layer is a layer of
passivated SiO_2 (glass). The
back surface of the cell, aside
from where back contacts
come through, is likewise

Figure 2-13 - A cross-sectional view of the A-300 solar cell [Mulligan 2004]

covered with SiO_2. The SiO_2 layers trap light inside the solar cell by reflecting internally, just like with a fiberoptic thread. The passivation also helps to reduce losses due to surface recombination. Where the electrical contacts punch through the rear passivation, there is a silicon-metal interface, which further aids in internal reflection. The third interface on the back of the cell is where the electrical contacts are bussed together (in between contacts), which results in a silicon-SiO_2-metal interface, which is also highly reflective.

As a result of these innovations in the solar cell construction, even though the cell is only 250 μm thick, the average light path length is approximately six times that [McIntosh 2004]. The longer a light-ray travels through the solar cell, the greater the chance is that it will create an electron-hole pair. As a result, the relative spectral response (or quantum efficiency) of the A-300 solar cell approaches 100% over a wide range of wavelengths (see Figure 2-14). In other words, the A-300 captures and makes better use of the incident sunlight.
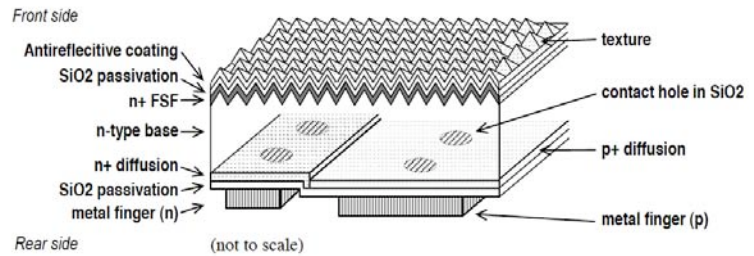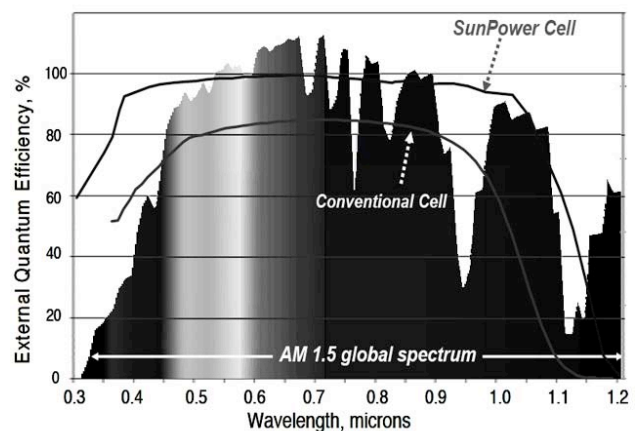

Figure 2-14 - The external quantum efficiency (or relative spectral response) of the A-300 cell [Mulligan 2004].

---

[6] This antireflective coating is the same as on windshield and eyeglasses.

A solar cell with its contacts entirely on the back surface necessarily has its p-n junction diffused through the back surface. A more typical construction has the junction diffused through the top surface, with the cathode contact on top and anode contact on the rear surface. Typically, more energetic photons (blue and ultraviolet) interact with silicon very near the surface, while longer wavelengths penetrate deeper [Tremble II]. Therefore, there are many electron-hole pairs that are created near the surface or in the bulk silicon, a relatively long distance from the junctions. The junctions in the A-300 cell are formed by diffusing alternating rows of n+ and p+ doped silicon in an n-type bulk (see Figure 2-13 above). The diffusions are relatively wide (the n+ diffusion is 600 µm wide), because they are formed with less expensive (and less precise) methods than photolithography. As a result of the relatively large feature size, an electron-hole pair created in the bulk or in the middle of one of those diffusions has a relatively long distance to travel before reaching the junction, where the minority carriers can be collected. Therefore, the n-type silicon starting material for the A-300 cell needs to have a very long minority carrier migration lifetime, so that they have a good chance at migrating all the way to the junction before recombining. The A-300 solar cell uses a relatively inexpensive photovoltaic floating zone (PV-FZ) silicon as its starting material, which has been doped to have a minority carrier lifetime of over 1 ms, which is extremely long compared to most chemistries [McIntosh 2003, Mulligan 2004, Rodgers 2005]. This long carrier lifetime allows most of the created electron-hole pairs to diffuse all the way to the junction. A
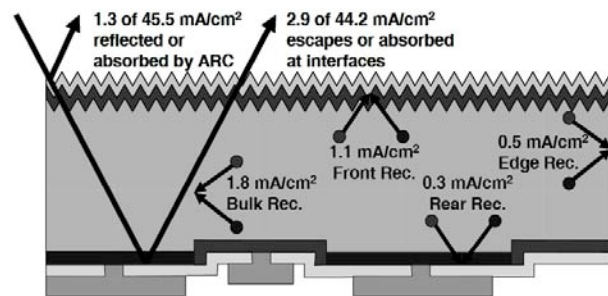

Figure 2-15 - Remaining loss mechanisms in the A-300. The losses are recorded in terms of current density which would ideally be 45.5 mA/cm² for a 1000 W/m² insolation [Mulligan 2004].

breakdown of the various loss mechanisms in the A-300 solar cell, which prevent its efficiency from reaching the theoretical maximum are given in Figure 2-15.

As one can see in Figure 2-12 above, the A-300 solar cell is roughly octagonal in shape. The reason for this is that they are made from 150 mm diameter circular wafers, because semiconductor-grade silicon is formed in cylindrical ingots. However, it is difficult to properly arrange a contact grid, which tends to be linear or rectangular, onto the back of a circular wafer. Circular cells cannot be backed together as closely as rectangular cells. While the difference in packing density may not be critical for some applications, a lower packing density means a lower power per unit area rating of the finished panel, which therefore has more coverglass and backing material, which increases both the weight and cost of the finished panel. On the other hand, if one were to saw the circular wafer into a perfect square, a great deal of otherwise useful silicon would be wasted (not to mention increasing the chances of chipping a corner here and there). The octagonal shape comes from sawing the 150 mm wafers down to 125 x125 mm dimensions, which is a good compromise between packing density and wasted silicon.

In addition to selling their solar cells in bulk to various customers, SunPower also manufactures their own finished panels for residential installations based on the A-300 cell. These panels have an aluminum support chassis and substrate, onto which is adhered a laminated sandwich of the solar cells between two layers of ethylene-vinyl-acetate (EVA), an optically clear material similar to hot glue, which encapsulates the solar cells. This sandwich is in turn laminated to a 4 mm thick sheet of coverglass, which is treated with its own anti-reflective coating. The panels have a 12 x 6 grid of solar cells wired in series, and boast an overall efficiency in the range 17-18% [Mulligan 2004, SunPower II]. These panels are

intended for residential or commercial installations and a 25-year lifetime, during which time they will be subjected to a variety of challenging weather conditions. As a result, they are very ruggedly built, and weigh approximately 225 g per cell [SunPower II].

2.5 THE DEVELOPMENT OF A NUMERICAL MODEL

In order to help predict the amount of power the solar panels could produce when in Antarctica, I developed a numerical model of the A-300 cell. Not only did this prove useful in predicting the power output, but also provided information on the panel I-V electrical characteristics well before the panels were built, which was essential for the design of the rest of the power system and the development of the overall control algorithm. Please refer to derivation in Appendix A for full details of the model of the I-V characteristics of a single solar cell, which shall merely be summarized here for brevity:

$$V_{cell} = k_B T \cdot \ln\left(\frac{I_{sc}(\Lambda,T) - I_{cell}}{I_0} + 1\right) - I_{cell} R_s \tag{2.8}$$

where

$$I_{sc}(\Lambda,T) = I_{sc,ref} \cdot \frac{\Lambda_s}{\Lambda_{ref}} + \alpha\left(T - T_{ref}\right) \tag{2.9}$$

$I_o \approx 32 \cdot 10^{-12} \text{ A}$

$I_{sc,ref} = 5.75 \text{ A}$

$V_{OC,ref} = 0.665 \text{ V}$

$\Lambda_{ref} = 1000 \text{ W}/\text{m}^2$

$T_{ref} = 298 \text{ K}$

$\alpha = 1.75 \text{ mA}/{^\circ}\text{K}$

$R_s \approx 0.1 \text{ }\Omega$

$k_B = 8.6173432 \cdot 10^{-5} \text{ eV}/{^\circ}\text{K}$

T ranges from $-50°C$ to $+50°C$

Λ_s ranges from 200 to 1200 W/m^2

and

$$V_{OC}(\Lambda,T) = k_B T \cdot \ln\left(\frac{I_{sc}(\Lambda,T)}{I_0} + 1\right) + \beta\left(T - T_{ref}\right) \tag{2.10}$$

$$\beta \approx -4.1 \; \text{mV}/_{\circ\text{K}}$$

Once the I-V characteristics of a single solar cell are determined (using the above model), one must then expand that model to determine the behavior of a panel of many such cells wired in series and parallel. For such combinations, Kirchoff's Current Law (KCL) dictates that the current through each element must be the same, and that the currents from parallel circuits are additive. Application of Kirchoff's Voltage Law (KVL) tells us that the overall voltage will be the sum of voltages of the individual cells. That is:

$$I_{panel} = N_{parallel} \cdot I_{cell} \quad \text{and} \quad V_{panel} = N_{series} \cdot V_{cell} \tag{2.11, 12}$$

where $N_{parallel}$ is the number of cells in series, and N_{series} is the number of cells wired in series. In the case of the solar panels for the Cool Robot, 54 cells will be wired in a single series string for the side panels, while the top panel will have 56 cells.

There are two other factors that need to be considered when predicting the behavior of a finished panel of solar cells from the behavior as many ideal individual cells wired together: reflection and self-heating. Reflection occurs due to the fact that the surface of the panel (which can be tempered glass, a silicone compound, or a thin sheet of some polymer like Tefzel) does not transmit all of the light that is incident to it. As a result, the amount of light that actually reaches the solar cells is somewhat less than what is incident to the panel. The relationship between the insolation incident to the panel and what reaches the solar cells can be considered linear:[7]

$$\Lambda_{cell} = \alpha_t \cdot \Lambda_{panel} , \;\; \alpha_t < 1 \tag{2.13}$$

---

[7] It is also influenced by wavelength and incident angle, but we will ignore that for simplicity.

where α_t is the transmission coefficient. The value for α_t must be found empirically for a particular panel construction technique.

As already explained above, the insolation that is not converted into electricity largely is changed into heat within the solar cells. That heat will conduct its way from the cells, through whatever encapsulation material surrounds them, and out to the open air through convection and radiation. For simplicity, I have lumped these three heat transfer mechanisms for the panel – conduction through the silicone, moderate convection at the panel surface, and a small amount of radiation from the panel surface – as a single parameter θ_{HT}, the heat transfer coefficient. Therefore, in steady state:

$$T_{cell} \approx T_{ambient} + \theta_{HT}\left(\Lambda_{cell} - \frac{I_{mp,cell}V_{mp,cell}}{A_{cell}}\right) \tag{2.14}$$

In reality, each of these mechanisms (particularly the convective cooling) are quite complex, and poorly modeled by the relationship in equation 2.14. For instance, the convective cooling is related to not only to the relative temperatures of the panel and the ambient air, but also the panel's orientation with respect to gravity and whether there is any breeze (forced convection). A rigorous treatment of the heat transfer mechanisms is outside the scope of this thesis, so this simpler model will have to suffice.

The self heating can be quite dramatic, causing the solar cells to be 10's of degrees centigrade above the ambient temperature. Usually, one can see the air above the surface of the panel rippling from the heat. At an insolation level of 800 W/m², even with a 1 m/s breeze and an ambient temperature of 20°C, the finished panels from SunPower have a cell operating temperature of 48.5° [SunPower II]. Generally, the self-heating effect is detrimental to a panel's efficiency, as solar cells have a higher power output for a given insolation level when they are cold.

2.6 TUNING AND TESTING THE MODEL

In order to validate the model, tune its parameters for close agreement, and test the fabrication methods that will be used for the Cool Robots project, two small, 12-cell test panels were fabricated and tested both in Hanover, NH and Antarctica. The fabrication method is described in greater detail in Chapter 3 and in Appendix C, so the details will be not be discussed here. One important detail is that the first panel that was fabricated had a smooth, glossy surface, while the second panel was textured like frosted glass. The texturing is helpful for capturing and transmitting a greater fraction of the incident sunlight (i.e., reducing reflection, especially at low angles of incidence) and to a small degree improving the removal of heat through an increased surface area. The second panel also benefited from my having learned from mistakes made during the fabrication of the first panel. As a result, the data presented here comes entirely from the second panel.

Each point on a solar panel's I-V curve corresponds to a certain resistance. If one connects the panel to a known resistance R and measures the resulting terminal voltage V the operating point $\left(V, \frac{V}{R}\right)$ on the I-V curve can be graphed. By varying the resistance from 0 (a short circuit) to ∞ (an open circuit), one can map the entire I-V curve. To facilitate this, I constructed a load box consisting of four power resistors with values 1, 2, 2, and 5 Ω, connected by several DP3T (double pole, triple throw) switches. When the switch positions are changed, the power resistors are connected in various series/parallel combinations, resulting in a total resistance that varies from as little as 0.4 Ω to as much as 10 Ω. To account for wire and contact resistance, the resistance of each possible switch combination was measured to within ± 5 mΩ using a four-point impedance analyzer. A worksheet was crafted, graphically listing the desired switch positions and leaving a space where the

measured panel voltage corresponding to each resistance could be recorded (see Appendix D). The open circuit voltage would be recorded by disconnecting the load box entirely, as would the short circuit current. All measurements were made using a handheld multimeter that was accurate to four significant digits. In all, 16 data points made up each I-V curve, which were post-processed and graphed using MATLAB (Appendix D).

Testing took place during clear days during the summer and fall in Hanover, NH. Testing was done close to solar noon for several reasons: 1) the insolation for that day is at its maximum, so that by tilting the panel at different angles to the sun, a wide range of panel insolations could be achieved, 2) the rate of change on the sun's elevation angle (and hence, the rate of change in insolation) is lowest at solar noon, 3) the sun is highest in the sky at solar noon, meaning the incident sunlight travels through the smallest possible cross section of atmosphere, resulting in the least deviation from the AM1.5 spectrum. Testing was performed on clear, cloudless days for two reasons: 1) to lessen the variation in insolation from one moment to the next and 2) reduce the amount of insolation the panel receives from sources other than the direct rays of the sun such as scattered light from clouds elsewhere in the sky. The diffuse light from the blue sky above is an unavoidable, but beneficial, DC term. Where possible, a site with a clear view of most of the sky was used.

Of greatest interest to us during the testing of the smaller panels was the variation of the I-V characteristic as a result of changes in the total insolation. Because the insolation from the sun is assumed to be a constant during the time span of these tests, the way to vary the insolation to the panel is to vary the angle of incidence between the panel and the sun. The angle of incidence, ϕ, is here defined as the angle at which the incoming sunlight strikes the panel. When the panel is directly facing the sun and the sunlight is normal to the panel's

surface, $\phi=90°$. When the panel is edge-on to the sun, $\phi=0°$. This is naturally different than the sun's elevation angle from the horizon, which I define as θ. These two angles are illustrated in Figure 2-16. As one would expect, the
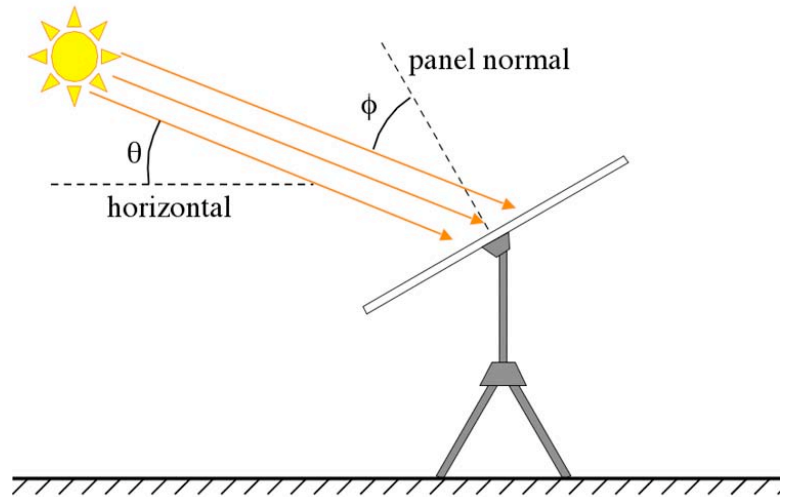
Figure 2-16 - The definition of the angle conventions used. The sun's elevation angle above the horizon is θ, while the angle at which the sunlight hits the panel is ϕ.

relationship between the insolation normal to the panel and the incident insolation is related to the sine of the angle between:

$$\Lambda_{panel} = \Lambda_s \cdot \sin\phi \qquad (2.15)$$

By tilting the panel at various angles to the sun, we are in this way able to adjust the amount of sunlight the panel receives. Generally, a full test would plot an I-V curve at roughly a half-dozen angles: $\phi = 90°$, $\approx 60°$, $\approx 45°$, $< 30°$, $< 10°$, and $0°$. These angles can be realized in several different ways that do not require direct measurement. For the moment, let us consider the case where the $\theta = 35°$:

$$\text{Panel Horizontal}: \quad \phi = \theta = 35°$$
$$\text{Panel Vertical}: \quad \phi = 90°\text{-} \theta = 55°$$

Other angles can be realized by observing the shadow the panel casts behind it:

$$\text{Maximized Shadow}: \quad \phi = 90°$$
$$\text{Zero Shadow}: \quad \phi = 0°$$

Other angles require measurement. From the geometry and angle conventions of the setup in Figure 2-16, it can be shown that:

$$L_{shadow} \cdot \sin\theta = L_{panel} \cdot \sin\phi \qquad (2.16)$$

so that if one knows the sun's elevation angle and the length of the panel, one can adjust the panel's shadow length on the ground (L_{shadow}) to produce a desired angle. Alternately, one can simply eyeball the desired angle ϕ, measure the resulting shadow's length, and later back-calculate what the actual angle of incidence was. A second way to measure the angle of incidence is with an instrument I crafted that consists of a ruler of sorts with a spire perpendicular to it. The height of the spire is known to within several thousands of an inch. The length of a shadow cast on the ruler is related to the spire height and the tangent of the angle of incidence. The "ruler" is placed flat on plane of the panel, and the sun casts the spire's shadow onto it. The markings on the ruler measure the angle of incidence that results in the shadow being a certain length.

The test data consist of 14 voltages resulting from a different loads, V_{OC}, I_{sc}, and the location and timestamp for the data. The sun's elevation angle, θ, can be determined to great accuracy from the day, time, and location of the test using a star-mapping program, or from various online almanacs [Gronbeck]. One very comprehensive algorithm can be found from the National Renewable Energy Laboratory [NREL II]. A reasonably accurate result can be obtained from the following equations [Lever]:[8]

$$D(N) = 23.45° \; \sin\left(\frac{2\pi}{365}(284 + N)\right) \qquad (2.17)$$

$$\theta(N,L,t) = \sin^{-1}\left[-\cos(L)\cos(D)\cos\left(15\left[\tfrac{°}{hr}\right]\cdot(12 - t)\right) - \sin(L)\sin(D)\right] \qquad (2.18)$$

---

[8] This set of equations is for the southern hemisphere. For the northern hemisphere, a minus sign must be changed here and there.

where D is the declination of the Earth to the sun, N is the Julian Day,[9] L is the degrees of south latitude, and t is the local (military) time in hours.

The angle of incidence, ϕ, for a particular set of test data can then be determined using the methods described above. This allows one to plot the I-V curve as a function of incidence angle. An example of the processed test data is given in Figure 2-17.
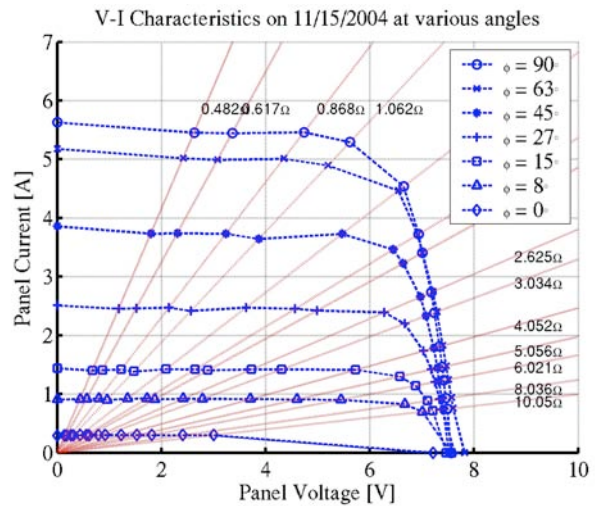
Figure 2-17 - An example of the I-V as a function of ϕ data taken with a 12-cell test panel.

The insolation at a given time of day was obtained from a pyrometer on the meteorological station at CRREL. Their instrument measures the insolation (in W/m^2) normal to a horizontal patch of ground. Using these insolation data, ϕ, and θ, one can determine the amount of sunlight incident to the panel, Λ_{panel} at the time of the test.

One can also plot the P-V curve, as in Figure 2-18. With only 16 discrete points, the actual location of the maximum power point P_{max} must be extrapolated by sketching a continuous curve through the discrete
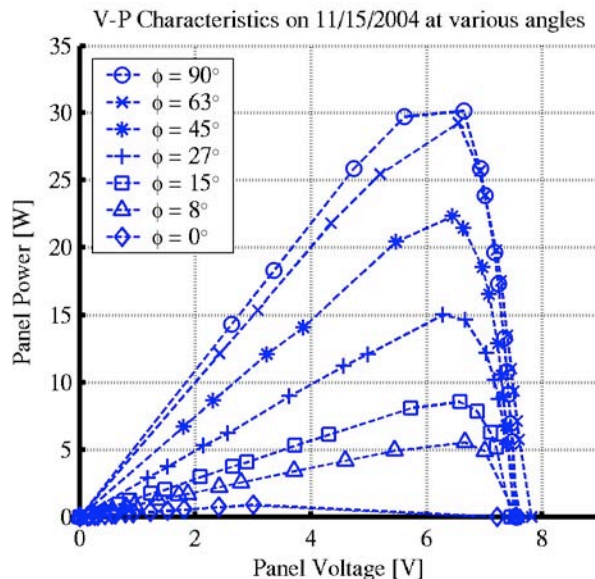
Figure 2-18 - The graph of power versus panel voltage from the same data presented in Figure 2-17.

---

[9] The Julian Day is defined by January 1 = 1, December 31 = 365, etc.

points. However, once the maximum power point is measured from the graph, one can determine the panel's efficiency:

$$\eta_{panel} = \frac{\Lambda_{panel} \cdot A_{panel}}{P_{max}} \qquad\qquad (2.19)$$

where A_{panel} is the area of the panel (0.180 m^2 for the 12-cell panels). Depending on the operating conditions, the efficiencies ranged from 14% to 18%.

Finally, one can plot the maximum power versus the angle of incidence, as in Figure 2-19. One would expect that the maximum power would be related to $\sin\phi$, and in general this is the case. Near to $\phi = 0°$, the experimental results are higher than the sine law due to the fact that, even though no direct insolation is reaching the solar cells,
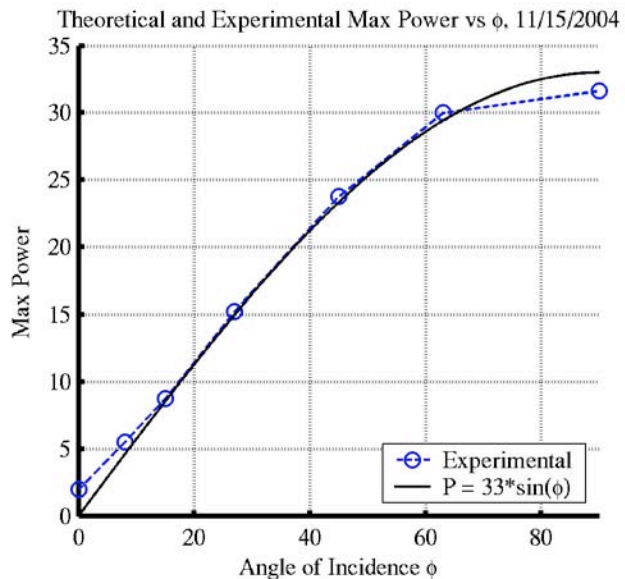
Figure 2-19 - The panel's maximum power output as a function of the angle of incidence, ϕ. From the same data set as Figures 2-17 and 2-18.

there is still a significant amount of diffuse radiation from the sky and surroundings. Near to $\phi = 90°$, the experimental results are lower than the sine law due to self-heating, which reduces the panel's performance.

As discussed previously, the numerical model for the A-300 solar cell was tuned to predict the behavior of the solar cells themselves. In order to predict the behavior of a panel, one must account for reflection from the panel's surface and self-heating. These two phenomena are accounted for in yet another layer to the model on top of the cell model

through the parameters α_t and θ_{HT}, respectively. The transmission coefficient, α_t, reduces the amount of insolation incident to the panel that actually reaches the solar cells to be converted to electricity by equation 2.13 from above:

$$\Lambda_{cell} = \alpha_t \cdot \Lambda_{panel} \tag{2.13}$$

And the resistivity of the heat transfer path, θ_{HT}, causes the cell temperature to be elevated using equation 2.14 from above:

$$T_{cell} \approx T_{ambient} + \theta_{HT}\left(\Lambda_{cell} - \frac{I_{mp,cell}V_{mp,cell}}{A_{cell}}\right) \tag{2.14}$$

These two equations allow the operating conditions for the panel as a whole to be converted into the operating conditions for the individual cells within the panel. Those operating conditions in turn are used by the model given by equation 2.8 to create I-V curves for a single cell. The results from the single cell model are scaled up by the number of cells in the panel. In the process, the parameters α_t and θ_{HT} are tuned to create good agreement between these generated I-V curves and the experimental ones. The results of this tuning and
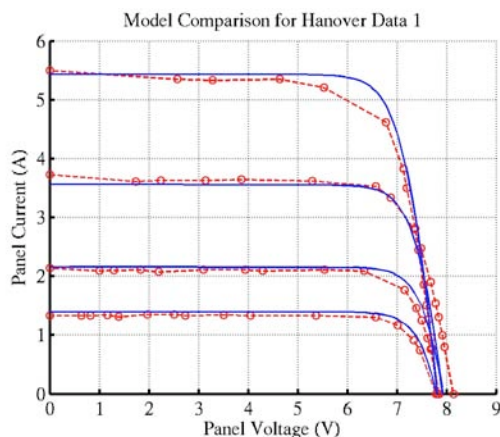

Figure 2-20 - Comparison of the tuned theoretical model of a 12-cell panel and test data of the same from Hanover. The red dotted lines show discrete test data, while the solid blue lines are the model predictions.
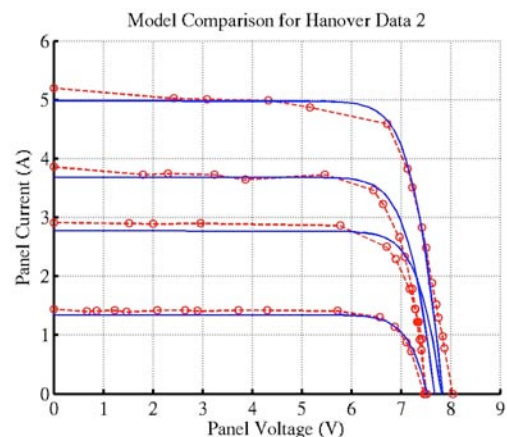
Figure 2-21 - A second comparison graph between the model and test data. Again, the dashed lines are test data, the solid lines are model data.

comparison are presented in Figures 2-20 and 2-21 (divided into two graphs for clarity). The solid blue lines represent actual test data, while the dashed red lines are the I-V curves generated by the cell/panel model for $\alpha_t = 0.78$ and $\theta_{HT} = 0.0375$ ˚C-m^2/W. As one can see, despite the long route taken to get to a result, the model has successfully been tuned to be able to reproduce the behavior of finished panels over a broad range of insolations and a relatively narrow range of temperature conditions.

Once the model was tuned for the A-300 cell and our fabrication method using test data taken in Hanover, N.H., it was then applied to the test data from some weeks later taken by Dr. Lever in Antarctica. The results of one such test are given in Figure 2-22, and another in Figure 2-23. As one can see, there is little agreement between the two, or at least, not nearly as much as there exists with the Hanover data. The discrepancy is particularly large in the case of the panel directly facing the sun. In general, the predicted open circuit voltages are higher, and the short circuit currents much lower. While there is sizeable disagreement between the model and test data, the model data consistently under-predicts what the maximum power should be, given the testing conditions. So, predictions made about power

Figure 2-22 - Comparison of the tuned 12-cell model to test data taken in Antarctica. The red dashed lines show the test data, while the solid blue lines are model predictions.
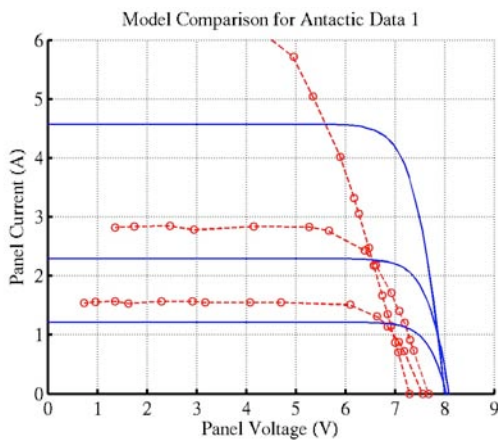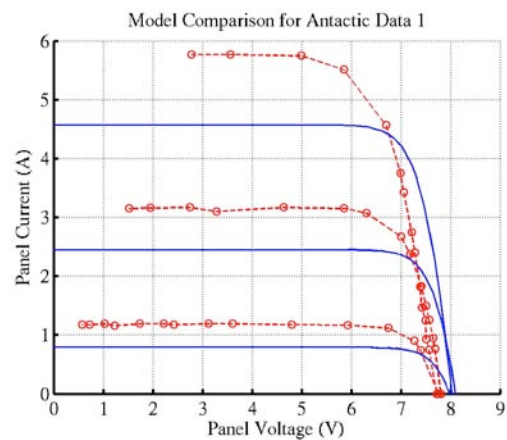

Figure 2-23 - A second comparison graph of the 12-cell panel model and data taken in Antarctica.

availability using the numerical model will be, if anything, overly conservative.

There are at least two possibilities that can help explain the discrepancy. The first is that, although measurements of the insolation were made with a recently-calibrated radiometer, there could have been significantly more sunlight available to the solar panels in Antarctica. As I will explain shortly, the amount of sunlight available from reflections off the snow is tremendous. And while no explicit modeling was done to evaluate its contribution, the diffuse light that comes from the sky above could be adding a significant amount of power that was unavailable in Hanover. Because of the thinner atmosphere, reduced ozone layer, and particularly low humidity, the spectral content of the insolation in Antarctica could be significantly different than in Hanover. If the pyrometer was calibrated to give insolation readings based on an Air Mass 1.5 spectrum, it could have given erroneous readings in Antarctica. Secondly, the thermal resistance coefficient, θ_{HT}, was tuned to match the conditions in Hanover. There are so many different factors that determine this coefficient (such as wind and other convective effects), that it would not be surprising if that θ_{HT} did not carry over to a different environment, or even a different day.[10] Better agreement for the Antarctic data could be obtained by retuning the α_t and θ_{HT} parameters, but then the model is no longer predictive, merely explanatory. Testing of the solar panels in Greenland this upcoming summer might help in the development of a more accurate model.

---

[10] During testing of the actual panels, in order to be assured of a nearly constant sun angle, the tests were conducted quite rapidly from direct insolation to edge-on, not slowly enough to allow the panel to come to a steady state temperature. Equation 2-14 makes the assumption that the panel's temperature is at steady state, which is another factor that can help explain the discrepancy.

2.7 Controlling the Operating Point

In all the discussion of solar cell modeling, I-V characteristics, and maximum power point tracking, one thing has not been discussed: how *does* one dictate where a panel will operate on its I-V curve? There are, in fact, at least three ways that one may accomplish this.

The most common way, which in some ways is how most maximum power point trackers work, is to present the panel with a certain resistance across its terminals. Each point on an I-V curve corresponds to a particular load or resistance. That is, for each location (V_{cell}, I_{cell}) on the curve there is a corresponding load

Figure 2-24 - A typical I-V curve of a solar panel with lines of equal resistance overlaid on it.

$R = V_{cell}/I_{cell}$. If one were to connect the solar panel to this resistance, the solar panel would operate or reside at that point on the curve. One could overlay an I-V curve with lines of equal resistance, which start at the origin and ray outward (Figure 2-24). The slope of the lines are $\Delta I/\Delta V = 1/R$, so a high slope corresponds to a low resistance, and vice versa. As the I-V curve changes due to changes in temperature and insolation, the operating point will translate along this "line of equal resistance." By varying the resistance (and hence, the slope of the line) one can traverse the entire span of the I-V curve, from zero load (a

Figure 2-25 - Dictating a panel's operating point on its I-V curve using a fixed resistive load.

59

short circuit, the y-intercept) to infinite load (and open circuit, the x-intercept). This arrangement is illustrated in Figure 2-25. It was in this way that the experimental I-V curves of the 12-cell test panels (and of the full sized panels) were obtained.

Because the solar cell has a maximum power point (V_{mp}, I_{mp}), it also has an optimal load $R_{opt} = V_{mp}/I_{mp}$. However, a fixed resistive load will not be able to track the movement of the maximum power point as it shifts due to changes in temperature and insolation. What's more, a fixed resistor does not make a very useful load – all it can produce is heat. The load does not actually need to be resistive; in practice, it almost never is. Instead, it is often the input impedance of a DC/DC converter, which can be adjusted by varying the duty cycle of its switching frequency.

There is one other commonly used way to dictate where a solar cell (or panel) will operate on its I-V curve. One can place a voltage source, such as a battery, across the terminals of the solar panel, which is
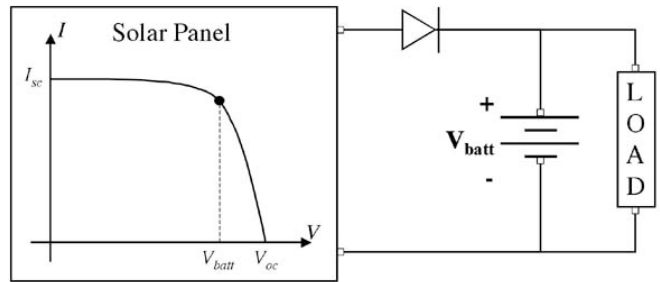
Figure 2-26 - Dictating the operating point of a solar panel using a voltage source.

shown in Figure 2-26. The battery, as a semi-infinite current sink or source, clamps the panel's voltage at V_{batt}. The solar panel's current is then whatever corresponds to the battery's voltage on the I-V curve. In order to prevent the battery from accidentally forcing current back through the solar cell, a diode is usually placed in series, so that current can only flow out of the panel. If there is a good match between the battery voltage and V_{mp}, then this combination can operate at close to peak efficiency. However, this is rarely the case, as the battery voltage changes with its state of charge, and the panel's V_{mp} can change

independently of the battery. The load, in whatever form it may be, can then be powered

from the V_{batt} voltage. This arrangement has the benefit that the battery can automatically

supplement the solar panel's power.

A third way to dictate where a
panel will operate on its I-V curve is
to place a current source in series
with the panel and the load. This is
illustrated in Figure 2-27.
Analogously to the case of a battery,
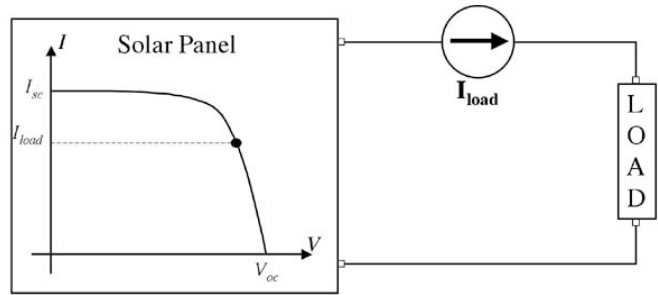
Figure 2-27 - Dictating the operating point of a solar panel using a current source.

the current source will clamp the panel's output current at a certain level. The intersection of

that current and the I-V curve is the operating point. An example of a constant-current load

would be a motor configured in torque control. However, as this is a very specialized case,

current control of solar panels is rarely encountered in practice.

2.8 Modeling the Availability of solar Power

Despite the fact that it is a
harsh and forbidding environment to
humans, Antarctica is actually an
ideal location for using photovoltaic
power. The cold temperatures and
steady light breeze enable the solar
cells to remain quite cool, which
increases their efficiency by as much
as several percentage points. The
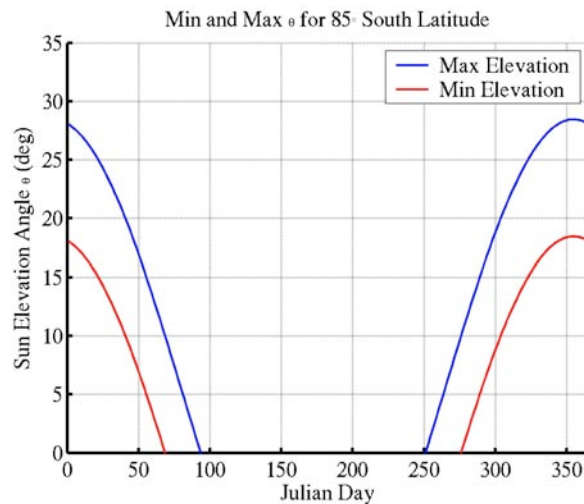
Figure 2-28 - Maximum and Minimum sun elevation angles at 85˚ S latitude throughout the year.

Antarctic plateau is quite high – the average elevation of the plateau is 3000 m, and the average elevation of the whole continent is 2000 m, which is considerably higher than any other continent [Linacre and Geerts 1999], which means the atmosphere above is thinner than at sea level. The cold, dry, stable air of the interior is crystal clear and nearly pollution-free and mostly isolated from the weather patterns of the rest of the southern hemisphere, which further improves transmission of the sun's rays to the surface. The thinning of the ozone layer, which is overall a bad situation, allows more of the ultraviolet end of the solar spectrum to penetrate to the surface. During the austral summer, the Earth is actually slightly closer to the sun than during summer in the northern hemisphere [USNO 2003]. The net result is that the sunlight that reaches the plateau is of a significantly higher intensity and quality than can be found anywhere else on Earth. What's more, during the austral summer, the sun never sets, but merely wanders about the horizon with a varying elevation angle (see Figure 2-28), providing an ever-present and relatively time-invariant power source.

Figure 2-29 shows data collected at the South Pole Station during the 2002 calendar year. The x-axis is the Julian Day. The y-axis records the vertical irradiance received; that is, the insolation that is normal to a horizontal surface (as the CRREL meteorological station does). Because the sensor is facing the sky, reflected light from the snow does not influence these measurements. If one accounts for what the sun's elevation angle was at the time when each data
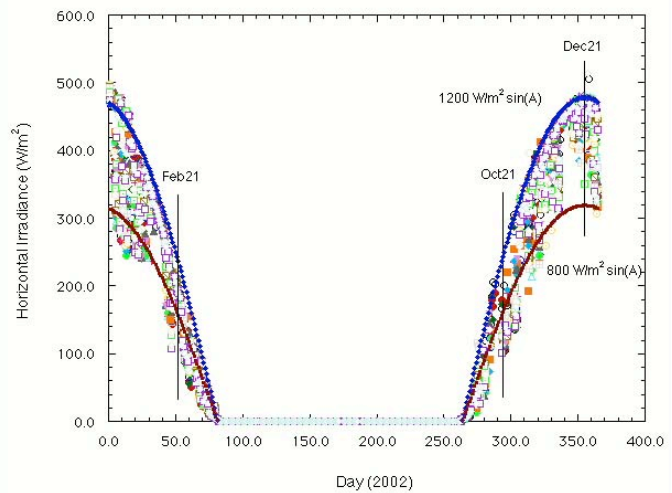
Figure 2-29 - Vertical irradiance data from South Pole Station. The solid lines represent the upper and lower bounds of the direct insolation from the sun, which varies between 800 and 1200 W/m².

point was recorded, one can back-calculate what the total insolation from the sun was. As one can see from the upper and lower bounds that have been added to the scatter plot, the insolation from the sun varies between 800 and 1200 W/m^2.

The abundant reflection of sunlight off the vast snowfields of the plateau is a further, and as we will see, quite substantial benefit. The albedo of Antarctic snow is extremely high, above 95% [Grenfell 1994, Warren 1998]. Because it consists of many crystals packed in a random fashion, snow scatters light of all wavelengths equally in all directions. Investigations of the snow reflection indicate that a vertical panel receiving no direct sunlight (i.e., edge-on to the sun) could still receive up to 30% as much power as if it were directly facing the sun [Kahn et al. 2004]. Even the panel facing away from the sun and shadowing part of the snowfield that would reflect light back to it can receive an insolation on the order of 10% of what the sun is producing.

To back up this assertion, Dr. Jim Lever made a theoretical analysis of sunlight reflecting off an infinite snowfield onto a vertical solar panel. A smaller problem is to consider the
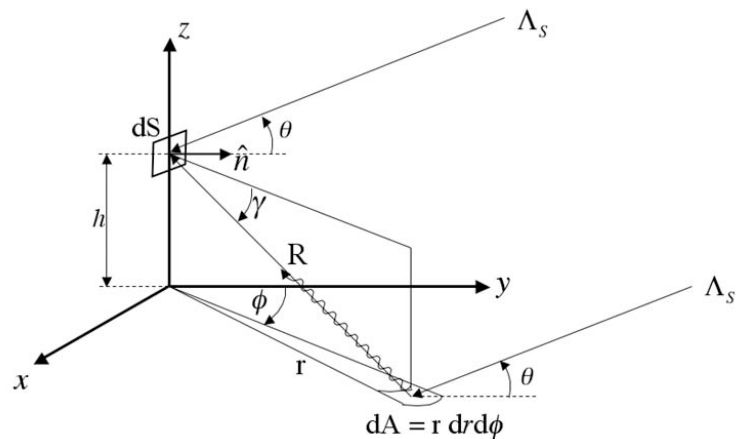
Figure 2-30 - Setup for the development of the theoretical analysis covered in Appendix B.

light reflecting from a small patch of snow, d*A*, onto a small piece of that vertical panel, d*S*, as shown in diagram given in Figure 2-30. The full details of the calculation are given in Appendix B. While a thorough determination is impossible without knowing more of the

physical properties of snow, an *approximation* of the insolation reflected to the panel (and normal to its surface) is given by $\Lambda_s \sin\theta$, which is in contrast to the direct insolation, which is $\Lambda_s \cos\theta$. Even more surprising than the simplicity of this result is the fact that, if one then considers the ratio of direct insolation to reflected insolation on a vertical panel, one will find that it is approximately equal to $\tan\theta$. For a $\theta = 20°$, that ratio equates to 36%, which means that the amount of insolation reflected to a panel is 36% as large as the direct insolation, which is quite significant.

A more difficult calculation is to determine the amount of power a panel will receive if it is facing away from the sun and thus shadowing some of the snow that otherwise would be reflecting energy onto it. Rather than attempt to perform such a calculation (which ultimately would need to be performed numerically), we have instead chosen to rely on empirical data taken during the winter of 2004 [Kahn et al. 2004] in Hanover, N.H., and also by Dr. Lever in Antarctica during the 2004-2005 austral summer. This empirical data, along with the theoretical analysis, has allowed the Cool Robots project to make estimates of the amount of electrical power that we can hope to extract from the solar panels while in Antarctica.

Consider the situation when the robot is driving straight towards the sun, which is at some low angle ($\theta < 25°$) above the horizon and produces some insolation Λ_s. Let us define the horizontal insolation that the front panel receives directly from the sun as $\Lambda_0 = \Lambda_s \cos\theta$. There is, however, also the reflected energy, which from empirical data and the theoretical analysis is approximately $0.35\Lambda_0$, which the front and both side panels receive. However, due to self-heating phenomena, the front panel's efficiency goes down, so let us say instead say that it simply receives less reflected insolation ($0.25\Lambda_0$, based on empirical data) and

retains its original efficiency. The rear panel, which is in shadow, receives approximately $0.1\Lambda_0$. The top panel receives direct insolation only, $\Lambda_s \sin\theta$. These relationships are crudely shown in Figure 2-31.
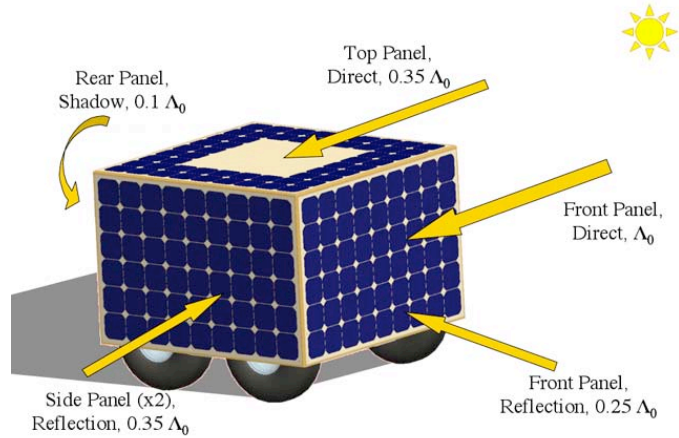
Figure 2-31 - Predicted insolation inputs to the five sides of the Cool Robot solar panel box.

The total power input to the robot is given by:

$$P_{T,insolation} = \Lambda_s\big[(1+0.25+2\cdot0.35+0.1)\cdot A\cos\theta + A\sin\theta\big] \qquad (2.20)$$

where A is the panel area.[11] The resulting power input from the solar panels is:

$$P_{T,panels} = \eta\cdot\Lambda_s\big[(1+0.25+2\cdot0.35+0.1)\cdot A\cos\theta + A\sin\theta\big] \qquad (2.21)$$

Using reasonable values for the efficiency of the manufactured panels (18%), solar cell area of a panel (0.81 m^2), sun elevation angle (20˚), and insolation (1000 W/m^2), one can evaluate the above equation to find get an expected power input of 330 W. This result in excess of the estimated power requirement for the robot, somewhere between 200 and 270 W, which means that the robot may not even make use of all the power available to it. In other words, there could very likely be a power *surplus*.

The situation just described, where the robot is driving directly towards the sun, is not even the optimal case. As a matter of fact, it is a minimal case, because only the front and top panels receive direct insolation. The robot can expect to receive its maximum solar power input when the sun is facing a corner of the robot, and illuminates not just the front and top panels, but also one side panel, as illustrated in Figure 2-32. When the sun is 45˚ to

---

[11] The top panel actually has two solar cells more than any of the side panels. However, since this is an estimation anyway, we will ignore this small difference.

the left or right of the robot, the front, top, and one side panel all receive direct insolation. If before, when the robot drove directly towards the sun, the front panel received horizontal insolation Λ_0, it would now receive $\Lambda_0' = \Lambda_s \sin 45° \approx 0.7\Lambda_0$, and the same amount of reflected insolation as before. However, the side panel,
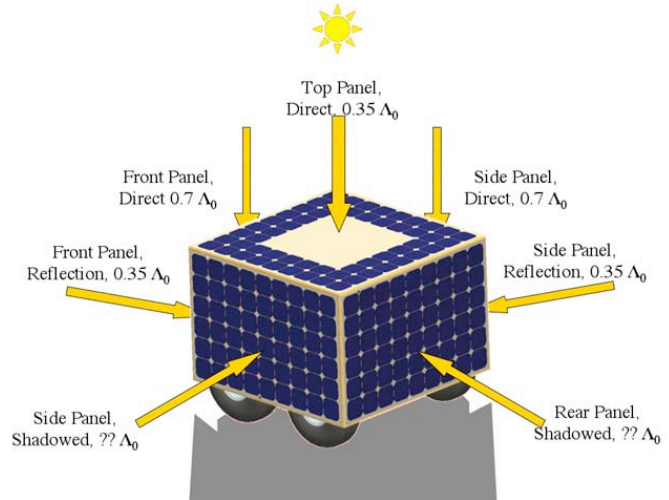
Figure 2-32 - Estimated insolation input to the Cool Robot's panel box when the robot is not directly facing the sun, but rather at 45° to it.

which before had only received reflected insolation $0.35\Lambda_0$, now also receives the same direct insolation as the front panel, Λ_0'. The result is an increase in insolation input to the solar panels of $0.4\Lambda_0$ (just Λ_0 before, now $2 \times 0.7\Lambda_0$). The top panel receives the same amount of insolation as before, $\Lambda_s \sin\theta$. The rear panel, which before had received only a small amount of reflected energy due to its shadow, now receives somewhat more insolation, because its shadow has been skewed to one side. The other side panel, which had received significant reflected insolation $0.35\Lambda_0$, now receives somewhat less because it is casting a shadow identical to the rear panel's. On the whole, however, the decreased output from the shadowed side panel is outweighed by the increased output of the other side and rear panels, resulting in a net increase in power. The exact amount of this increase is difficult to determine, since it is difficult to evaluate the effect of the skewed shadows on the insolation to the side and rear panels. When the robot is first deployed, determining the power availability in this case will be a matter of great interest. In any event, if we encountered a significant power surplus before, when the sun directly faced one of the panels, we can

expect an even greater surplus when the sun is at some angle to the robot, which will be the case the majority of the time.

2.9 Conclusions on Solar Power in Antarctica

It should be apparent from the above discussion about the availability of solar power in the Antarctic that we can be confident in having enough power to drive even at full speed continuously. The analysis of the power input when the robot is driving directly to the sun, which we have shown to be a minimal case, demonstrates that the solar panels should be able to provide power in excess of what is required by the robot to drive at full speed. In other cases, where a side panel is illuminated in addition to the front and top panels, the expected solar input should increase.

It is difficult to overstate the extraordinary usefulness of the reflected power illuminating the four side panels, which contributes some 45% of the total power input. The reflected power, in fact, makes the mission possible. Were the reflected power unavailable (in a terrain with lower albedo) or inaccessible (because there were no solar panels facing away from the sun) the power input to the robot would be significantly less, and the robot would be unable to drive continuously for long periods without stopping to recharge, which would make its ability to traverse 500 km in less than two weeks less certain.

It is important to note that, no matter where the sun is in relation to the robot, the conditions that each panel are exposed to (temperature and insolation) will be different from the others. As the robot drives over *sastrugi*-strewn terrain, its orientation, and the orientation of each panel, relative to the sun will rapidly change. I approximate the period of this motion to be 2 seconds. What's more, as the sun circles around the horizon and changes its elevation angle with the time of day, the sun's output and location in the sky will also

change. As result, the I-V characteristic of each panel will be different from the rest, and will change rapidly in time. Because we should have a surplus of available power from the panels, no matter their respective operating conditions, not all the panels will have to operate at their respective maximum power points. Even if we had perfect information about the operating conditions of each panel, the numerical model of a panel may not be accurate enough to be able to predict, in real-time, what the I-V characteristic of each panel is.

These facts have important ramifications for the design of the rest of the power system. Each panel must operate independently from the others, implying that the power electronics that condition the solar power input from each panel must operate independently from the rest. The power electronics must be able to adapt to the wide range of time-varying operating conditions of the solar panels will be exposed to. Finally, the control mechanism that oversees the power system and dictates where the solar panels will operate on their respective I-V curves must be able to recognize changes in the panels' characteristics and adapt quickly in a closed-loop fashion. The design of these power electronics and the overall architecture of the power system is the focus of the next chapter, while chapter 4 is devoted to the development of the control system.

3. Design and Construction of the Cool Robot Power System

3.1 The Architecture of a Photovoltaic Power system

A photovoltaic-based power system generally has five parts: the photovoltaic source, input power conditioning, energy storage, output power conditioning, and control. The characteristics of the photovoltaic power source have already been described in exhaustive detail in Chapter 2. Input conditioning, output conditioning, and storage will comprise the majority of this chapter, while the control system is the focus of Chapter 4. The role of the parts other than the solar panels I shall briefly describe here.

Input Power Conditioning – As has already been established, the operating characteristics of a solar panel can change greatly with changes in insolation and temperature. As a result, the location of the maximum power point on the I-V curve, (V_{mp}, I_{mp}), also changes over time. In order to achieve the greatest use and efficiency from the photovoltaic power system, the solar panels are generally connected to a Maximum Power Point Tracker (MPPT). The power point tracker ensures that the panels are operating at their optimal location, and usually consists of a switched DC/DC converter. The large signal input resistance of the DC/DC converter can be tuned in real-time (by modulating the duty cycle) to be the optimal value V_{mp}/I_{mp}. However, it is possible to have a PV system with no explicit MPP tracking, and yet still require some sort of input power conditioning to change the panel's terminal voltage to a DC or AC voltage of a desired (and regulated) value.

Energy Storage – In most places in the world, the sun only shines during certain hours of the day, and not necessarily when the demand for power is at its greatest. What's more, with a fixed installation, such as panels mounted to a home's roof, the intensity of the

insolation changes with the time of day. Therefore, it is advantageous to have some form of energy storage that can accept excess power from the solar panels when it exists and then later discharge it when the power from the solar panels is insufficient. Usually, a bank of batteries serves as the storage medium, although one can find other means (flywheels, reversible fuel cells, ultracapacitors, etc).

Output Power Conditioning – The power available from the bus, which comes from the input conditioning and/or the energy storage, usually needs to be further transformed to be useful. For instance, the bulk power on the bus could pass through an inverter to create a typical 120 V, 60 Hz electrical supply typical of residential PV systems. In other cases, the power needs to be changed through myriad voltage supplies to power the internal electronics, as in a satellite power system.

Control – Each of these blocks needs some sort of control in order to operate properly, either as a stand alone unit or as part of the overall PV system. For instance, the input conditioning requires some sort of control algorithm in order to seek out and track the time-changing location of the maximum power point. Sometimes, the topology of the system uses the input power conditioning as a charging regulator for the batteries, and the output conditioning as a discharge regulator. Other times, there is a special bit of hardware that further isolates the batteries from the power bus. In either case, the state-of-charge of the battery needs to be monitored and regulated so that it is neither over-charged or –discharged.

3.2 LITERATURE REVIEW

In [Krumman 2001], Krummann describes the topology and operation of a solar power system for a next-generation GOES satellite. The satellite uses a single solar panel which is divided into a number of independent strings, which experience the same uniform

insolation. Each string is connected to its own DC/DC converter that performs MPPT, which provides a certain amount of fault protection and redundancy. Because the solar cells are identical and receive the same uniform insolation, the MPPTs receive a common duty cycle command, so that they all operate in the same fashion. The output of the MPPTs are connected in parallel to a power bus that provides power to the internal electronics and payload. A pair of large batteries provide energy storage, and receive power from the bus via a dedicated charge controller. The batteries are also discharged through the charge controller. Any excess power from the solar panels is shunted through resistors, and in this way the bus voltage is regulated to $53.1 \text{ V} \pm 0.25 \text{ V}$.

In [Shirbacheh 1997], Shirbacheh gives an overview of the photovoltaic power system of the Pathfinder mission to Mars from 1996.[1] The overall design emphasized simplicity and low-weight over performance. For instance, there was no explicit maximum power point tracking (apparently the benefits were not apparent for a Mars mission). Instead, the solar panels of the lander were directly connected to a common bus. The bus could be regulated in a crude way using a number of fixed shunt resistors that could be attached to the load in 8 discrete steps. For storage the lander used a pair of Silver-Zinc batteries, which were likewise directly connected to the bus. Current would flow in or out of the batteries depending on where the bus voltage was (due to the presence or absence of the shunt resistors) in relation to the battery voltage. To control the flow of power to the various subsystems, the team used a number of latching relays (mechanical switches that are flipped on and off by pulses of current).

---

[1] The power system on mission reused several designs from the Cassini mission, which is intriguing since Cassini utilizes a radioactive power source, not solar power, and is of a considerably larger size.

Reference [Tan and Tseng 2003] deals with the design of a power system for a group of small satellites (< 500 kg, which is small by satellite standards). Each satellite has two solar panels that are continuously oriented towards the sun, and a smaller array of solar cells affixed to the body. Each panel has an independent MPPT, and their outputs are combined together and feed into a battery charge unit, which regulates the main power bus voltage and so controls how much power flows in or out of the batteries. The paper, after giving an overview of the designs, then goes on to demonstrate the effectiveness of several different MPP tracking algorithms.

Moving away from space-based applications, [Bhide and Bhat 1992] describe the design and control of a modular photovoltaic installation in India designed principally for charging lead-acid batteries. The design is modular in the sense that each panel has its own power conditioning electronics, which all contribute (in parallel) to the common power bus. The maximum power point tracking is done using a microcontroller (microcontroller) that merely needs to know how many modules are attached to the system to be able to operate all of them. The design of the system was made to be extremely simple and allow for easy expansion or reconfiguration. Likewise, the control algorithm is computationally very simple, allowing many arrays to be controlled simultaneously using inexpensive digital hardware.

Reference [Sullivan and Powers 1993] describes the design of the maximum power point tracking system of a solar car built by the University of California at Berkeley. The solar car had a single solar panel, which would often be exposed to rapidly changing insolation conditions as it drove in and out of shadows on the road. As a result, the control algorithm had to be both fast and robust. The power output from the solar panel fed to a

72

power bus, from which the main drive motor for the car would extract the majority of power. A set of batteries also was directly connected to the bus, and their presence on the bus acted as a buffer or clamp – dictating what its voltage would be. This fact was utilized in the design of the MPP tracking system. An extensive analysis was made to determine the optimum design and selection of power components to make the tradeoff between weight and converter efficiency.

Pacheco et al., in [Pacheco 2003], designed a rather ingenious maximum power point converter that incorporates the charge/discharge function of an attached battery to either accept excess or provide supplemental power to a load. The design of the DC/DC converter itself was unorthodox, a combination of boost and buck topologies, which utilized two switches. One of the switches, which was responsible for extracting power from the solar panel, was controlled by a relatively simple analog circuit that was responsible for the MPP tracking function. The second switch was controlled by another, independent, analog circuit that regulated the output voltage to the load using a basic PI controller. So, in a single device, the functions of input and output conditioning combined , along with regulating the storage.

This is just a sampling of the large amount of work that has been published on the design and control of photovoltaic power systems. The majority of systems that one finds in the literature, as well as in practice, utilize maximum power point tracking all the time to make the best use of the solar panels. Subsequently, one needs to make provisions to dump surplus power once the storage devices (batteries) have been fully charged and there is no place left for the surplus power to go.

The design of the Cool Robot power system is in some ways unique, and differentiates itself from the literature and current state-of-the-art in several ways that will be described in turn. As an example, one does not often find in the literature the situation where multiple panels, with vastly different insolation and temperature operating conditions, feed power to a common bus. The closest analog to our situation would come from solar cars which, do to their body shape, usually have multiple independent panels and, when driven around, can encounter rapid changes in insolation. Generally, as in [Krummann 2001], [Tan and Tseng 2003], and [Bhide and Bhat 1992], when multiple independent panels exist, they are all oriented the same way and receive the same uniform insolation. While in theory each panel could be drastically different from one another, that usually is not considered in the design. The fact that the solar panels of the Cool Robot necessarily have to operate under different conditions requires us to plan this into the design.

Another common feature in the design of PV systems is a separation between maximum power point tracking, bus voltage regulation, and battery management, usually having explicit power electronic devices managing each function. As in [Krummann 2001]. [Tan and Tseng 2003], and [Pacheco 2003], the MPP trackers are separate units that place power onto a central bus. Regulating the bus voltage is a combination of sending power to the load, diverting power to/from the batteries (which are isolated from the bus by a charge controller), and shunting any extra power through resistors as heat (i.e., increasing or decreasing the load). The control of each of these units – the MPPTs, the battery charge controller, and the shunt load – can be centralized through a microcontroller, or can be completely decentralized, in which case each unit operates independently from the others. In the Cool Robot power system, the control is completely centralized by the microcontroller.
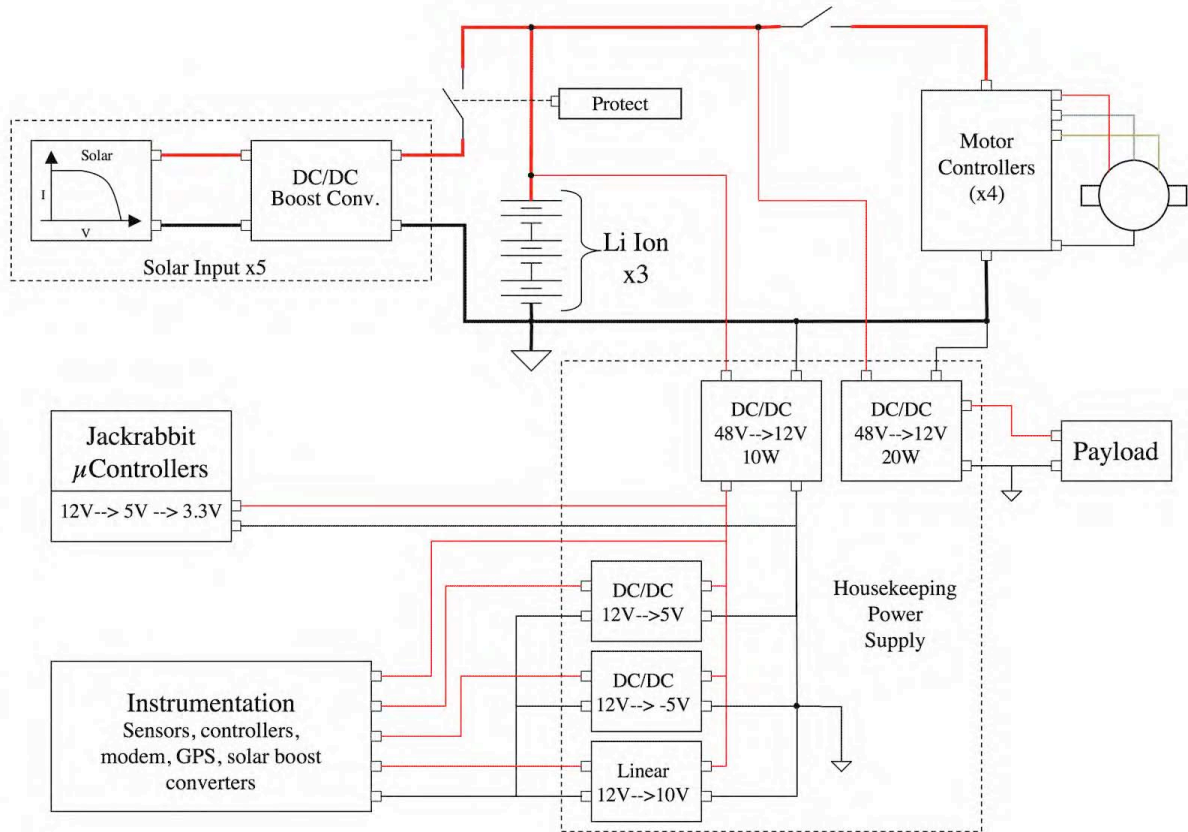
Figure 3-1 – Block diagram and overview of the Cool Robot power system.

However, the robot does not use maximum power point tracking all the time, nor does it have

a shunt load, nor does it have an explicit battery charge controller. By avoiding using these

separate devices, we hope to make the design of the power system less complex.

3.3 Overview of the Power System Architecture

A block diagram of the power system of the Cool Robot is presented in Figure 3-1. A

stack of Lithium Ion (Li-ion) batteries establishes the main power bus voltage at (nominally)

48 V. The four motor controllers draw power from this bus and use it to power the brushless

3Ø DC motors. The path to each motor controller has a 3 A slow-blow fuse (not shown) to

prevent a fault (such as an excessive current draw) from one of the motors (or controllers)

from propagating through the rest of the power system. In addition, there is a MOSFET switch

that can completely cut off power to the motor controllers, which would be advantageous when the robot is stationary, as the motor controllers dissipate power even when the motors are not turning. Each of the five solar panels delivers power to the bus via its own DC/DC boost converter, which raises the panels' (nominally) 32 V up to the level of the bus. There is a second MOSFET switch that can disconnect the output of the boost converters from the rest of the system.[2]

Power is drawn from the main bus to a variety of off-the-shelf DC/DC converters to provide supply voltages to the rest of the electronics in the system. One converter steps down the bus voltage to 12 V and can supply up to 20 W, which will be dedicated to the science payload and several specific circuits that will be discussed later. A second 12 V converter can supply up to 10 W, from which other converters establish +10, +5, and –5 V for the rest of the other instruments and electronics. The Jackrabbit microcontrollers are mounted on an evaluation board. The evaluation board accepts power from the smaller 12 V converter and steps it down to its own +5 V rail, from which the evaluation board derives a 3.3 V rail to supply the microcontrollers.

An overview of the various signals associated with the power system is given in Figure 3-2. The current and voltage of each panel will be periodically sampled using a 16-channel analog-to-digital converter (ADC). The ADC will also sample the bus voltage and battery current, leaving four unused analog channels for future use. The microcontroller will also be able to talk to an 8-channel digital-to-analog converter (DAC), which will provide

---

[2] This is perhaps an unnecessary complication to the design, since the boost converters can individually or collectively be disabled by the microcontroller. Whenever this MOSFET is disabled, the DC/DC boost converters must also be disabled, or else they would continue to extract power from the solar panels and, having no place to send it, fail in a spectacular fashion.
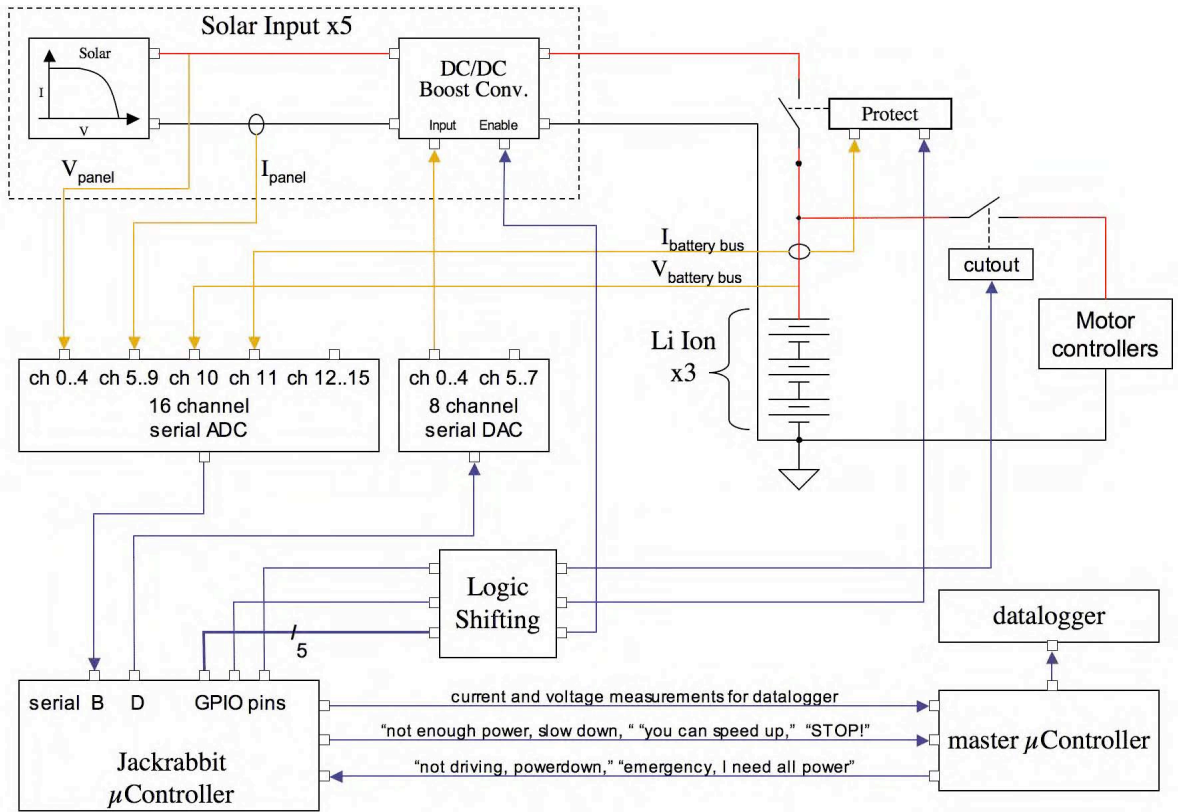
Figure 3-2 – Overview of the various measurements and control signals of the power system.

analog commands to the DC/DC converters, which will leave three unused output channels for future use.

A number of general purpose input/output (GPIO) pins on the Jackrabbit will be used to send digital commands to various parts of the power system. For instance, each solar DC/DC converter has an Enable function, allowing them to be individually turned off. The MOSFET that can disconnect power to the motor controllers receives its command from a GPIO pin, as does the MOSFET that can disconnect the solar power input. Through the master/slave 8-bit data bus, the Jackrabbit that controls the power system can pass messages back and forth to the master Jackrabbit, which controls the navigation and communication. The messages can be used to send current and voltage measurements to the on-board datalogger as well.

3.4 Description of Major Elements

In order to give a better understanding of the operation of the Cool Robot power system, a description of most of its major components will now be given. In some cases, the details that went behind the design (such as in the case of the boost converters) will be left to the appendices.

Solar Panels

Based on the modeling of solar power in Antarctica and examining the space constraints of the Twin Otter aircraft, the robot design uses the following arrangement: each side panel consists of 54 cells in a 9x6 grid, and the top panel would have a double ring of 56 cells.[3] For each panel, all the cells would be wired in series, which would create an input voltage range of between 30 and 40 V and an input current from as little as one ampere to as much as 6 or 7 A.

However, even once we had settled upon and ordered 500 of the A-300 solar cells from SunPower, and determined their configuration on the Robot, we still had to figure out and teach ourselves how to change those cells into the finished panels for the robot. There are several manufacturing techniques that are used in the industry, and have various tradeoffs for weight, durability, and transmission efficiency. A typical panel made for a residential installation or satellite system follows the construction given in Figure 3-3.
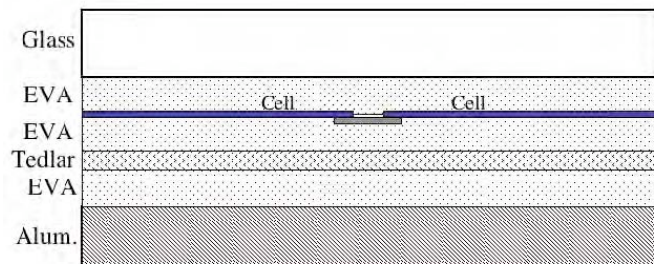
Figure 3-3 – A cross sectional view of a solar panel constructed for residential installations. (not to scale)

---

[3] A ring of 9 cells to a side, with a ring of 7 cells to a side within it (see Figure 1-5).

The cells are encapsulated between two layers of ethylene vinyl acetate (EVA) or a material similar to it. EVA is often used, because it is optically clear, lightweight, flexible, resistant to long exposure to ultraviolet light, available in sheets of standard thicknesses (0.5 mil through 100 mil), and fills in any voids between cells by reflowing at high temperatures during construction, like hot glue. Oftentimes, the backside EVA is pre-laminated to a layer of Tedlar (polyvinylideneflouride) or Tefzel (ethylene tetrafluoroethylene or ETFE), which is a thin, clear, polyvinyl plastic sheeting similar to Teflon that offers durability and substantial electrical resistance. On the front side, the EVA is also laminated to a thin layer of tempered glass, which has been treated with an antireflective coating to improve light transmission, especially at shallow incidence angles. For residential installations, which have to guard against the weather, this glass layer is especially tough, and able to withstand rain, snow, and hail [Rodgers 2005]. For satellite applications, the glass is often treated to absorb extremely short UV rays that could degrade the layers underneath over time. The glass is also very thin to save on weight, only several thousandths of an inch thick, which substantially increases the cost. This whole sandwich is laminated to a substrate of plate or honeycomb aluminum. In residential installations, the lamination is given additional mechanical support and provided with mounting points by an aluminum frame called the junction box.

While this construction method is quite common and durable, it tends to either be quite heavy or expensive, or both. For the SunPower SPR-210 panel, which uses the same A-300 cells as the Cool Robot, the finished weight is roughly 225 g per finished cell. For a satellite application, the finished weight is much less, but the manufacturing cost is considerably higher. What's more, this kind of construction technique needs to be performed in a vacuum oven, lest air pockets develop between the layers. The Thayer School has no

such facility, certainly not one large enough to fit an entire panel. Investigations of companies within several hours' drive that have or sell such equipment yielded no one who was willing to do such a small lamination job under subcontract. In any event, the large number of layers, particularly the glass, trap heat inside, resulting in a higher cell operating temperature, which will reduce the overall efficiency. For instance, while the A-300 solar cell is 21% efficient, the SPR-210 solar panel made from those cells only has a minimum guaranteed efficiency of 17% [SunPower II].

In applications where weight is a prime factor, such as solar racing vehicles, the construction given in Figure 3-4 is preferred. Instead of heavy glass as the top surface, a layer of Tefzel is used instead. While this is not nearly as durable as solid glass, it is still remarkably sturdy, despite being only a few thousandths of an inch thick. Being of a similar
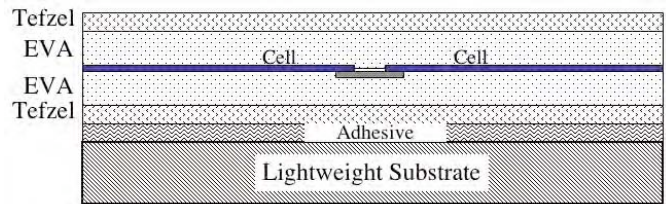
Figure 3-4 – A cross sectional view of a solar panel construction for solar racing vehicles. (not to scale)

chemical composition to Teflon, it is very easy to clean as well. The laminations of Tefzel and EVA still need to be performed in a vacuum oven, however. When complete, the flexible lamination needs to be adhered to a sturdy substrate, which is often the body of the racer itself. The body panels of the racer are often made from a composite construction, much like the honeycomb panels we are using, but are often curved to create a streamlined body. Adhering the lamination to the body can be difficult due to the fact that Tefzel (or Tedlar) is hard to bond to. One can skip the backside layer of Tefzel, but then one needs to have a facility large enough to laminate directly to the body panels. Nevertheless, this

construction method is very light (< 40 g per cell, depending on the substrate), and has excellent heat removal due to the thinness of the lamination.

A variation of this construction technique, which did not require the use of a vacuum lamination facility, was used by the Pathfinder solar-powered airplane, as described in [Carey 1994]. This unmanned craft, developed in the 1990s, was essentially a very large flying wing, whose entire top surface was covered with a dense grid of solar cells. Because weight was at an ultimate premium in this application, the laminations were made using adhesive-backed sheets of Tedlar

Figure 3-5 – Cross section of the construction method used for the Pathfinder solar airplane. [Carey 1994]
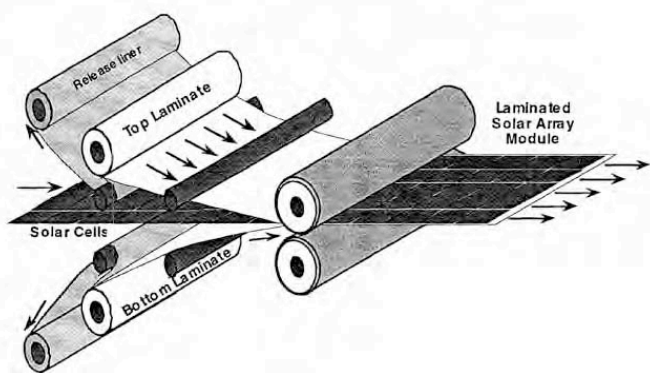
Figure 3-6 – Rolled lamination construction technique used for the Pathfinder. [Carey 1994]

film, as thin as scotch tape (Figure 3-5). The laminations were performed using very large rolling machines, like a printing press (Figure 3-6). Were we to use a similar process with similar materials and our own honeycomb panel substrate, the finished panels would have weighed only 20 g per finished cell. While this yielded excellent results, conversations with the subcontractor who did the rolling laminations for NASA, Advantage Converting, determined that the material would probably be too expensive for our application, not available in the correct dimensions [Alchimisti 2004], or simply proprietary information that

could not be divulged [Sinton 2004, Aceves 2004]. Even if we did have the material on hand, it seemed unlikely that we would be able to create our own laminations without: 1) allowing air pockets to be created between the film and the solar cells or 2) breaking the solar cells from pressing down too hard while trying to avoid introducing air bubbles [Gay 2004]. What is really needed for this method is the specialized rolling equipment of the manufacturer, which was unavailable to the Cool Robots project.

The construction method we developed for the Cool Robots project is based on silicones. This is a technique sometimes used by solar racing teams that do not have the resources to do thin laminations in a vacuum. The silicones are liquid before curing, which means that air bubbles are largely avoided on the surface of the solar cells. Their rubbery texture means that they are very durable and protect the cells against impact and vibration. While this technique is lighter than that used for commercial panels, it is heavier than thin laminations, and the insulative properties of silicone trap heat more effectively than laminations. However, the equipment and facilities to do this sort of construction were readily available at the Thayer School, and the material cost was relatively low.

To develop and test the construction technique, I built two small-sized test panels of 12 cells each, arranged in a 4x3 grid. A full
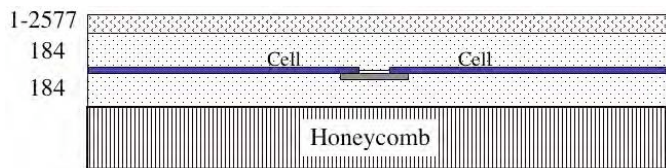
Figure 3-7 – Cross section of the silicone-based construction technique used for the Cool Robot's solar panels. (not to scale)

description of the construction technique is given in Appendix C. A cross section of the construction is given in Figure 3-7. In brief, the solar cells are encapsulated between two layers of a two-part silicone. The first layer is poured onto the 1/4" honeycomb panel substrate. Strings of several cells that have been soldered together in series are placed on top

of this layer while it is still wet. For the test panels, three strings of four cells were used; for the full-sized panels, nine strings of six cells were used. The top panel required a trickier arrangement. Once the first layer has cured, the final electrical connections between strings are made and a second layer is poured to cover the front surface of the cells. Unfortunately, the rubbery texture of the two-part silicone means that dust (and anything else) sticks to it quite tenaciously. A different top-coat would be needed.

Initially, it was thought that the top surface could consist of a thin film of Tedlar or Tefzel. The film would cling, rather than adhere, to the smooth surface of the two-part silicone. However, experiments with samples of the films revealed that it was impossible to lay the film down without trapping some air bubbles and dust beneath. The trapped air would have significantly decreased light transmission. It would perhaps have been possible to lay the film down over the second layer of silicone while still wet, as is done in [Komp 1987]. However, to do this over the whole surface of the panel would have been frightfully messy, so it was soon abandoned.

Another alternative would be to use a one-part silicone for the top surface. Several such silicones exist, so-called conformal coatings, which are used primarily by the electronics industry to weather-proof circuits. When they have dried, their surfaces are not rubbery but hard and clean easily. One could, in theory, use these conformal coatings for the entire panel construction, bypassing the two-part silicones, but for three key issues. The first issue is that the coatings are quite hazardous. Rather than curing through a chemical reaction that has no byproducts, as the two-part silicones do, the conformal coatings cure by evaporating a nasty and likely carcinogenic solvent (see Appendix C-3 for the Materials Safety Datasheet). The second issue is cost – the conformal coatings cost at least twice what

the two-part silicone does, which does not include the increased hazardous material shipping and disposal cost. The third issue is that the conformal coatings are meant to be used thinly (< 1 mm thick), and so would encapsulate the solar cells far less effectively than the bulkier two-part silicones. So, this silicone was used only as a very last step, to produce a durable and easy-to-clean top surface.

Another problematic issue with this construction procedure is low-angle transmittance. The glass used in commercially-produced panels has a micro-textured surface with an antireflective coating, much like the solar cells themselves do. Both the two-part and one-part silicones, when they cure under the influence of gravity, produce a very smooth surface that, while aesthetically attractive, reflects a significant amount of light. The transmittance generally is worse at shallow angles of incidence (< 15°). The only way to decrease the reflectivity of the top surface was to give it a texture, much like frosted glass, what is referred to as a stippled surface. The effect is described visually in Figure 3-8. One way we discovered to do add this texturing was to lay a woven fabric of Teflon-coated fiberglass over the one-part silicone as it was drying. Once dry, the silicone would retain the imprint of the fabric. We tested this on the small scale with one of the two 12-cell test panels – texturing its surface with fiberglass while leaving the other one glossy. In testing conducted both in Hanover and in Antarctica, we found that the second panel, with the textured surface, did indeed have higher performance, particularly at low angles.

This technique of texturing the surface did not, however, scale up well to the full-sized solar panels. We attempted to achieve the same textured surface using the same procedure. However, when the Teflon-coated fiberglass fabric was rolled back, a deal of the textured silicone came with it, leaving the second layer of two-part silicone behind.
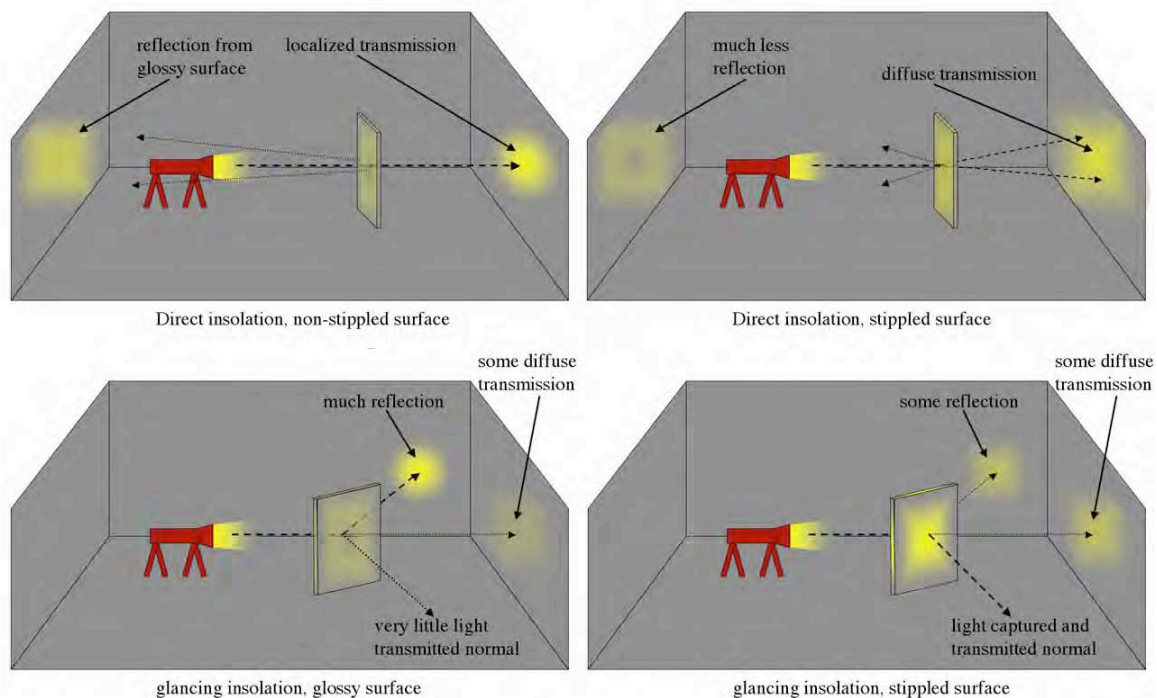
Figure 3-8 – The effect of stippling on direct and low-angle transmittance in comparison to a smooth, glossy surface.

Obviously this was not an acceptable finish, so we were forced to abandon this procedure. Thankfully, tests with the non-stippled panel indicate that, for most angles of incidence, the need for a stippled surface is not paramount. The only panel which would truly need a stippled surface would be the top panel, which will always be subjected to low-angle insolation in Antarctica. For that panel, the stippling can be done in smaller segments, rather than the whole panel at once.

Rather than build all of the full-sized panels with a glossy surface, however, an alternative was serendipitously discovered. Just like the two-part silicone, the one-part silicone needs to be spread out to ensure a uniform coating. If one continues to work the silicone, however, after several minutes the tacky surface of the silicone will cause the spreader to chatter across the surface. As it chatters, it creates indentations in the surface. Doing so for several minutes and in random directions will create a textured surface that,

while not nearly as neat and uniform as that obtained from the fabric imprint, is at least considerably less smooth than the glossy surface (see Figure 3-9).

In this way we were able to construct the full-sized panels for the robot. In general, these panels had an efficiency of greater than 18%, and yet with a final weight

Figure 3-9 – The surface finish achieved through active texturing.

of about 70 g per cell (3.5 kg/m^2). At low angles, the efficiency would decrease due to the reflective losses. At high angles of incidence (i.e., close to direct insolation) the efficiency of the panels would also decrease, as the test panels had done, due to self-heating from all the additional unconverted sunlight.

In these regards – weight, low-angle reflection, and self-heating – the construction procedure is far from ideal. In addition, this construction technique requires a large amount of delicate manual labor, and takes nearly a week from start to finish to complete one panel. While this method was acceptable for this first prototype, and was capable using the facilities available at the Thayer School, there is considerable room for improvement. The weight and self-heating problems can be alleviated by using substantially less silicone per cell. Indeed, with each successive panel that was made, the amount of silicone used per cell was steadily decreased. Were more panels to be made using this method, it would not be difficult to produce panels weighing less than 50 g/cell.

Ultimately, however, best way to reduce the weight of the panels, improve the removal of heat, and perhaps to increase the low-angle transmittance, is to use a lamination

procedure like that used by solar race cars. Late in this research project, contact was made with a Alain Chuzel of SunCat Solar, who specializes in custom solar laminations. Using the bare cells and honeycomb panel we sent to him, a spare, 5th side panel was constructed to evaluate its performance relative to our own. However, due to the fact that the Cool Robot's project did not discover this outlet until after our own construction of the solar panels was underway, and the lengthy turnaround time from SunCat, we did not have the time or financial resources to subcontract all of the panel fabrication. However, based on the evaluation of this panel upon its return, future generations of Cool Robots may use this laminate construction. The cost per panel in parts and labor (minus the cost of the solar cells and honeycomb substrate) would be roughly $500. This is in comparison to the roughly $125/panel silicone part cost, and approximately 12-15 hours of labor.[4]

DC/DC BOOST CONVERTERS

As concluded in Chapter 2, each of the five panels of the Cool Robot will operate under different conditions of insolation and temperature. That is, the I-V curve of each panel will be different from the others, and will change in time. What's more, it will often be the case that the robot has a surplus of power available from the solar panels. As a result, it is necessary to have some block of power electronics to fuse the power from each solar panel onto the common bus. Based on the modeling and testing performed in Chapter 2, it is known that the solar panels' V_{mp} will vary in the range of 32-36 V, and V_{OC} in the range of 35-42 V. The bus voltage, which is set by the state of charge of the Li-ion batteries, varies between 42 and 49.6 V, with a nominal value being 48 V. A DC/DC converter is therefore

---

[4] As a grad student, the cost of my labor is vaguely unquantified, but it is reasonable to expect a labor cost of $20-30/hr for this level of skilled work.
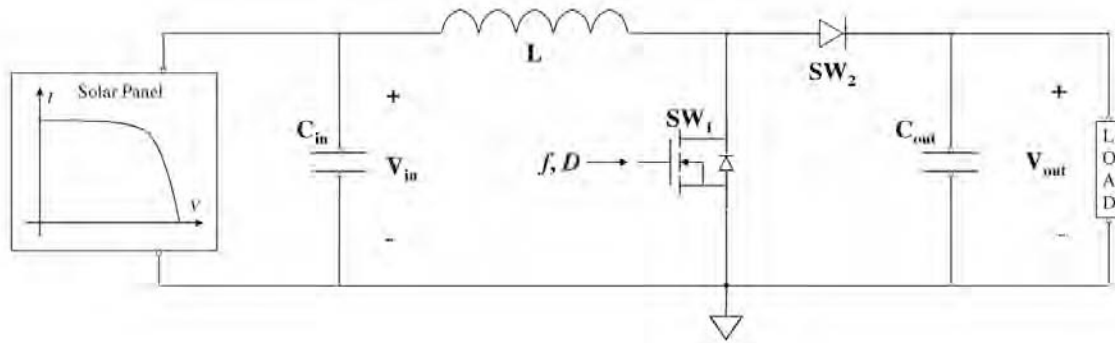
Figure 3-10 - The boost converter topology used to convert the output of the solar panels onto a common bus.

necessary, one for each panel. As the panels' voltage will always be less than the bus voltage, but not a great deal less, a simple boost converter topology will be utilized.

The boost converter topology, shown in Figure 3-10, consists of an inductor (L), two switches (SW$_1$ and SW$_2$), and one input and one output capacitor (C$_{in}$ and C$_{out}$). The main switch, SW$_1$, is an actively controlled switch and is typically implemented as a MOSFET or IGBT. The second switch can also be an active switch, but is here implemented as a diode.

The operation of a boost converter can be summarized as follows. The active switch receives a pulse-width modulated (PWM) command of frequency *f* and duty cycle *D*. When SW$_1$ is closed (during the *D* or On time of the frequency period), the inductor has V_{in} placed across it. By the component rule of an inductor:

$$\frac{dI_L}{dt} = \frac{V_L}{L} \tag{3.1}$$

the change in current through the inductor is proportional to the voltage placed across it. The larger the inductor value *L*, the slower the change. So, the current through the inductor ramps up, drawing from both the panel and the input capacitor C$_{in}$.[5] The ramp-up continues for a time $D/f = DT$, then the input to the gate of the MOSFET switches off, and SW$_1$

---

[5] Usually, one assumes that the input capacitor is sufficiently large such that the actual power source only supplies a DC power, while the capacitor handles the AC fluctuations.

becomes an open circuit. The current in the inductor wants to continue flowing, and does so by flowing through the diode, which becomes forward biased in the process. The voltage across the inductor is now $V_{in} - V_{out}$. In order to prevent the current through the inductor from continuing to ramp up to infinity, we require that $V_{out} > V_{in}$, so that the voltage across the inductor is now less than zero, and the inductor
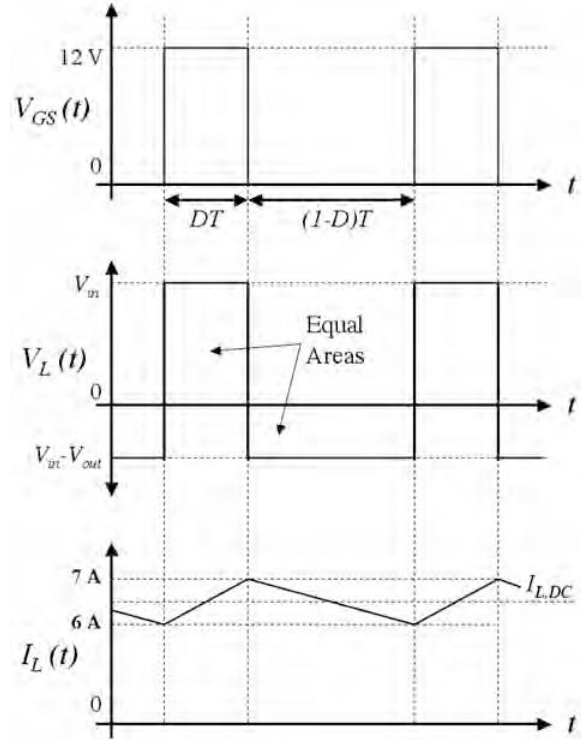
Figure 3-11 – Current and voltage waveforms for the inductor in a boost converter. The V_{gs} of the MOSFET is given as a timing reference.

current ramps down. This current ramp-down continues for a time $(1-D)T$. The inductor waveforms are given in Figure 3-11. When the boost converter has reached a condition called periodic steady state, the input and output voltages have reached a semi-constant level. In order for periodic steady state to exist, there can be no net change in inductor current over the course of one period. That is, the amount of ramp-up must equal the amount of ramp-down. In other words:

$$V_{in} \cdot DT = \left(V_{in} - V_{out}\right) \cdot (1-D)T \tag{3.2}$$

Therefore:

$$V_{out} = \frac{V_{in}}{1-D} \tag{3.3}$$

This relationship between the input and output voltages, which is characteristic of a boost converter, is generally applied to the case where the input voltage is fixed (like from a

battery or power supply) and one wishes to find the duty cycle that will raise that input voltage to a desired higher output voltage. For instance, an off-the-shelf DC/DC boost converter is used in the robot to take power from the 12 V payload power supply and boost it up to 15 V to power the magnetometer. However, in the case of the solar power system, we are presented with the opposite case. The input voltage is not fixed; indeed, it is something that we want to adjust in order to control where the panel sits on its I-V curve. The *output* voltage, however, is fixed by the presence of the batteries, which clamp the bus voltage to a certain level. In this case, it is more appropriate to rearrange equation 3.3 to be:

$$V_{in} = (1 - D)V_{out} \qquad (3.4)$$

By fixing the output voltage with the batteries and using the microcontroller to adjust the duty cycle of the boost converter, one can force what the input voltage of the boost converter is and, by extension, what the terminal voltage of the solar panels will be. More on this will be discussed in Chapter 4.

Ideally, the boost converter has 100% efficiency. The efficiency of a DC/DC converter, η, is the ratio of the output power over the input power:

$$\eta = \frac{P_{out}}{P_{in}} = \frac{V_{out}I_{out}}{V_{in}I_{in}} < 100\% \qquad (3.5)$$

In reality, one must factor in such things as the voltage drop across the diode, resistance through all of the components, the power required to operate the device, and power dissipated through the MOSFET when it is switching between on and off. These various losses detract from the output power, and thus decrease the overall efficiency of the converter, although efficiencies above 90% are not difficult with careful design. Even though the Cool Robots team believes that there will often be a surplus of power available to the robot, it is still advantageous to build a high-efficiency converter. There are several

practical reasons for wanting to design for high efficiency. The first reason is that whatever power is lost due to inefficient design is lost as heat, which must be removed from the components to prevent their burning out. Removing heat from components requires ever larger and heavier heatsinks, which is detrimental to a lightweight robot design. Secondly, while the robot can expect a power surplus, it is likely that future generations of Cool Robots would utilize smaller and lighter panels that provide no surplus, requiring the power system to make the most efficient use of the available solar power.

With efficiency and light weight in mind, the following specifications for the DC/DC boost converters were developed:

Input Range – As stated earlier, the voltage range of the solar panels between the maximum power point and the open circuit voltage is approximately 30–42 V. This will be the required input range for the converters, although a wider range would be easy to accommodate. The input current can be as high as 6.5 A in bright sunlight.

Output Range – The output voltage, which is the common bus voltage, is clamped by the presence of the batteries. Therefore, the output voltage will range between about 42 and 50 volts, depending on the batteries' state of charge, and have a nominal value of 48 V.

Control – The boost converter will accept an analog input in the range 0–5V, which will modulate the duty cycle.

Operation – The switching frequency will be 100 kHz.[6] In order to minimize oscillations about the operating point on the I-V curve, the input voltage ripple will be less than 0.5 V_{p-p}. The inductor's input current ripple will be less than 2 A peak-to-peak. The

---

[6] The choice of switching frequency tends to be a compromise. A high switching frequency allows one to use substantially smaller components. However, the power losses in the inductor and capacitors, operating power to drive the MOSFET, and switching losses tend to scale up with frequency (sometimes more than linearly), resulting in lower efficiency.

output voltage ripple will be less than 0.5 V_{p-p}. The boost converter must be able to operate in an ambient temperature of –40˚C. In order to fulfill the projected power estimations, the efficiency of the boost converter must be at least 90%.

Weight – Initially 500 g was allocated for each converter, which proved to be overly conservative. A more reasonable upper limit for the weight of each converter is 300 g.

Size – There is ample space inside the robot, as well as between the chassis and the panel box. Therefore, there is no hard criteria for the size of the boost converter.

The analysis that led to the component design and selection is provided in Appendix F, and a schematic of the final circuit is given in Appendix I-1. The input capacitor C_{in}, which keeps the solar panel operating at nearly a constant location on its I-V curve, consists of a 1000 μF electrolytic capacitor in parallel with a 10 μF polyester film cap. The electrolytic ensures a large capacitance, while the poly-film ensures a low impedance at the 100 kHz operating frequency. The 100 μH inductor has a custom design optimized for low losses and an acceptable weight. The MOSFET switch is an International Rectifier 1010EZ, which is rated for 60 V and 75 A and has a low on-resistance of 8.5 mΩ. The diode is another IRF product, the 30CTQ60, which is rated for 60 V and 30 A and has a low forward voltage drop of 0.425 V at the predicted operating conditions. The datasheets for these parts are provided in Appendices H-2 and H-3, respectively. The output capacitor uses another 1000 μF electrolytic, but also four 10 μF polyester film capacitors to ensure a low output voltage ripple.

The design of the FET drive circuitry, which is responsible for accepting an analog command from the power microcontroller and producing the PWM signal that turns the MOSFET on and off, is also given in Appendix F. A pair of 555-timer circuits establish the

100 kHz clock frequency and use the analog input command to modulate its duty cycle between 35% and 99%. In order to isolate the low power signal circuitry from the MOSFET driver and the power components, the PWM signal then passes through an opto-isolator, the H11L1M, which also inverts the signal, so that the duty cycle of its output varies between 2% and 60%. All of the components are commonly grounded. The PWM command is received by the MOSFET driver, an IXDD404 (see Appendix H-4 for datasheet), which provides the bursts of current needed to the MOSFET gate in order to rapidly turn it on and off.

The design outlined in Appendix F represents the second generation. The first generation was designed with a lower switching frequency of 50 kHz and much tighter specifications in terms of the allowable input and output voltage ripple. As a result, that design used considerably larger and heavier components. The first generation design used a 250 µH inductor that weighed over 200 g just by itself. In order to have very low output voltage ripple, four 1000 µF capacitors were in parallel with a single 10 µF polyester film capacitor, which likewise weighed a lot on their own. The first generation prototype overall weighed some 350 g. The power handling capabilities of this boost converter were well within specification, and boasted a 97% conversion efficiency. Although this design satisfied the specifications given above, several possible improvements were obvious.

The execution of these improvements led to the construction of the second generation prototype. This design utilized the 100 kHz switching frequency, which allowed the inductor and capacitors to be substantially reduced. In particular, the inductor was reduced from 250 µH to 100 µH. The large output capacitor stage was slimmed down to utilize four of the polyester film capacitors and only one of the electrolytics. Not only did this reduce the weight of the converter, but it also ensured that the operation of the device will remain clean

at low temperatures.[7]   The second generation prototype had a weight of 250 g, and its efficiency remained at the same high level as the first generation prototype.

While the second generation prototype, as a circuit design, was an improvement over the first generation, it would have been quite difficult and labor intensive to produce another five or six of them for the Cool Robot. The reason is that the wiring of all the circuits, both the PWM signal circuitry and the power components, was done by hand. This technique of cutting each wire to length and soldering the many connections one wire at a time is good for producing one-off prototypes, but ill suited for even medium-scale production. There was also an opportunity to make the circuit more compact and lightweight.

Therefore, a printed circuit board (PCB) implementation of the second generation circuit was designed. There are too many considerations to be made in designing the component layout and manufacturing process for printed circuit boards to be covered in even modest detail in this document. PCB design is not an area that is integral to the research of this thesis in any case. The finished board measures 2.5 x 4 inches. A total of six such converters were made – one for each panel plus one spare. A set of 200% scale images of the copper bottom, copper top, and silkscreen top layers are provided in Appendix G-2. The
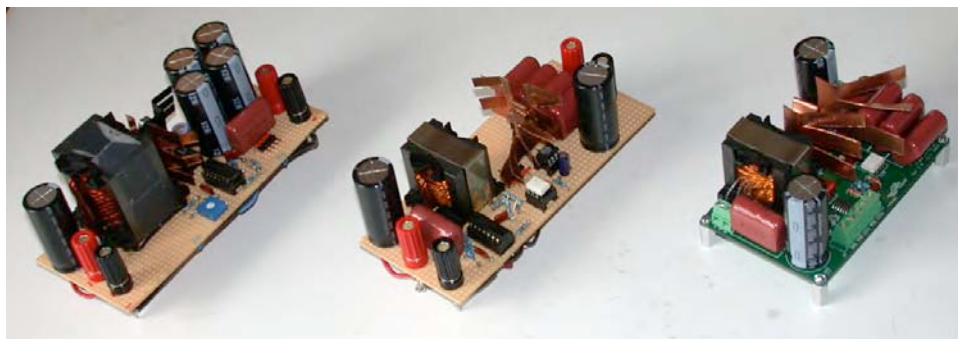
Figure 3-12 – The progression of the DC/DC boost converter design from first prototype to finished third generation. Along the way, the weight dropped from 350 g to 200 g.

---

[7] The reason for this is that the electrolytic gel that gives these capacitors their name becomes far less effective when it is cold, and the operation of the device becomes more resistive and less capacitive.

finished board is not only more compact physically than the hand-wired board, but weighs a mere 200 g. Compared to the original specification and weight budget that allowed 500 g per converter, this represents and impressive weight savings of 1.5 kg. The progression from first generation to PCB third-generation is shown in Figure 3-12.

Testing of each of the finished boards was accomplished using several high power power supplies, a dual-output bench power supply, a number of Fluke multimeters, an oscilloscope, and a bank of high power resistors. The boards (and the prototypes that preceded them) were designed and built well before the solar panels were constructed.[8] As a result, the boost converters could not, for a long while, be tested in the way that they would ultimately be used. Instead of a solar panel, whose output voltage and current vary greatly, the boost converters were connected to a pair of heavy-duty power supplies each capable of producing 10 A at 20 V. Instead of an output voltage clamped by the presence of a large battery, the converters were connected to a fixed resistor load. While this testing arrangement was not representative of the boost converters' ultimate application, it was sufficient to test their functionality as boost converters. So long as the boost converters function as designed in this testing arrangement, they would also operate as planned when finally installed in the robot's power system.

The testing covered a range of input voltages, and duty cycles. Input voltages varied from 5 to 35 V, and duty cycles of 10%, 25% and 50% were used. The load was a fixed 6 Ω, 250 W resistor. For each set of operating conditions, simultaneous measurements of V_{in}, I_{in}, V_{out}, and I_{out} were taken using the Fluke multimeters. Because operational power required to

---

[8] In this, it was advantageous to have the numeric model of the A-300 solar cell and the completed panels, so that we could have a sense of the range of currents and voltages that we could expect the panels to eventually provide.

create the PWM signal and drive the MOSFET should not change with the amount of power being shuffled from the input to the output, a single measurement of the 12 V supply input power was made. The measurements of V_{in}, I_{in}, V_{out}, and I_{out}, along with the fixed operating power of the boost converter allowed measurements of the net efficiency at various operating conditions to be made using equation 3.5.[9] On the whole, the PCB-based boost converters all functioned admirably. The lowest efficiency measured was 94%, with a short duty cycle and low input voltage. However, as the amount of power being converted increased, so did the efficiency of the device, as the various fixed losses of the device became increasingly insignificant. On the whole, the efficiency of the devices was above 97%. What's more, the power handling capabilities of the device were also quite high. In various burn-in tests, the boost converters could handle as much as 300 W of power. In actuality, the only limiting factor during the tests were the capabilities of the power supplies and the resistive load. As expected, the diode dissipated the greatest amount of power, and required a heatsink.[10] The MOSFET, under high current, would at times be warm, but not hot. The inductor remained at room temperature, as did the capacitors. We therefore have confidence in the boost converter's ability to satisfy the needs of the power system.

There are, however, further opportunities for improvement. For instance, if one examines an oscilloscope trace of the gate-to-source voltage, V_{gs}, of the MOSFET during a switching transition, as in Figure 3-13, one can see that it does not have the idealized form shown in Figure F-2. Instead, there is a great deal of ringing. Ringing was also found on V_{ds}, which in turn appeared on V_{out}. While the ringing does not appear to have adversely

---

[9] The operating power was subtracted from the output power when applying equation 3.5.

[10] Due to the density of components on the board, an off-the-shelf heatsink would not fit. As a result, a custom-shaped one needed to be created from copper sheeting.

affected the efficiency of the circuit, it did place a considerable amount of high frequency switching noise on both the input and the output – as much as 500 mV_{p-p} and on the order of 30-50 MHz. While it is unknown exactly what effect this noise may have on the overall operation of the
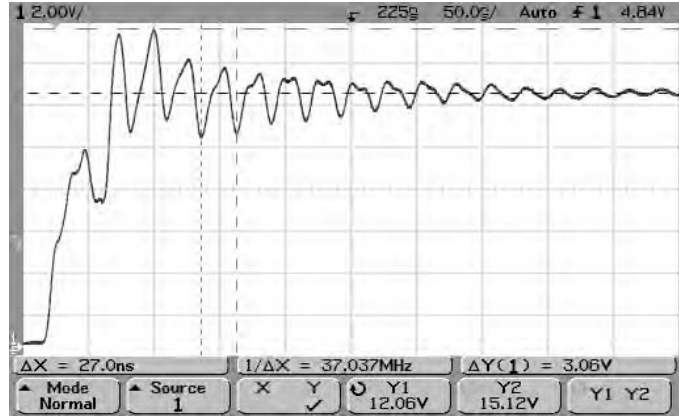
Figure 3-13 – Oscilloscope trace of the V_{gs} of the Mosfet during a low-to-high transition. The ringing is at a frequency of approximately 37 MHz.

power system, it almost certainly is detrimental to the accuracy and cleanliness of whatever measurements the instrument payload is making.[11] Therefore, both upstream and downstream of the boost converters, chokes will be used to try and reduce the amount of high frequency noise that propagates to the rest of the power system. This is discussed further later on. While it is rarely possible to remove this ringing on V_{gs} from a DC/DC converter, there are ways to mitigate it. Simply adding a low-value resistor in the path between the FET driver and MOSFET gate would add damping to the parasitic inductance and capacitance along that path. A simple circuit called a snubber, consisting of a resistor and capacitor, placed across the drain and source, would help to dampen the ringing on V_{ds}. However, a both solutions would also slow the switching transitions of the MOSFET, thus increasing its power dissipation. A more careful PCB layout could reduce the parasitic inductance and capacitance that is the likely cause of the ringing.

---

[11] For this reason, there is an operational mode, described in Chapter 4, where the robot disables the DC/DC converters and the motor controllers and runs entirely on battery power to reduce the amount of noise in the measurement circuits and instrument payload.

Considering the success of the boost converter design thus far, there is also a possibility to further increase the switching frequency of the circuit, which would allow still smaller components to be used. With a switching frequency several hundred kHz, it is conceivable that the weight of the converter could be further reduced below 200 g. A further refinement would be to replace the Schottky diode, which currently is the greatest source of loss, with another MOSFET that is switched in a complementary fashion to the primary MOSFET. Such a circuit is called a synchronous rectifier. The advantage is that the conduction losses at that point in the circuit would be reduced from $I_{SW2}V_{diode}$ to $I_{SW2}^2 R_{DS,on}$, which would be a decrease of at least an order of magnitude. Another improvement would be to completely remove the large electrolytic capacitors from the design, as they have significantly larger losses compared to polyester film or ceramic capacitors due to their higher equivalent series resistance. Late in the project, it was discovered that, although the polyester film capacitors have a lower ESR, the electrolytics have a significantly lower impedance due to their overwhelmingly large capacitance. As a result, most of the current will preferentially flow in and out of the electrolytics, which will reduce the overall converter efficiency.

HOUSEKEEPING POWER SUPPLY

The combination of solar panels, DC/DC boost converters, and batteries creates a source of bulk power. However, the various electronics of the robot – microcontrollers, analog and digital circuits, GPS, modem, and science payload all need very specific and regulated voltage supplies. Table 3-1 gives the voltage requirements for each of the main components. The supply requirements for the payload are somewhat unknown at this time, since the whole premise of the Cool Robot is to be able to carry any number of different

payloads. However, for the first payloads of a magnetometer, dual frequency GPS, and datalogger, a single 12 V supply will be sufficient. While it would have been possible to custom design a number of DC/DC converters, or perhaps a flyback converter with multiple windings, to supply these various voltages from the bus

Subcircuit	Voltage Specification	Voltage Provided
Jackrabbit Development Board	8-24 V	12 V$_s$
ADC evaluation module	5, ±12 V	±5, 10 V
DAC circuits	±5 V	±5 V
GPS evaluation module	10 V	10 V
Radio Modem	6-25 V	10 V
Handheld Radio	4.5-6 V	5 V
Iridium Modem	5 V	5 V
Motor Controllers (x4)	20-80 V, PGND	V$_{batt}$, PGND
DC/DC boost converters (x5)	12 V, PGND	12 V$_{p, s}$, PGND
Solar Input Board	±5 V, PGND	±5 V PGND
Power Distribution Board	5, 12 V, PGND	5, 12 V$_s$, PGND
Payload	Unknown	12 V$_p$
CR1000 Datalogger	12 V	12 V$_p$

Table 3-1 – Voltage supply requirements of the various subcircuits. It is assumed that each circuit will need at least a signal ground connection.

voltage, it was far easier to simply buy off-the-shelf voltage converters. While these off-the-shelf components are not as efficient as custom-built ones would be, they are considerably more compact and easy to interface to the rest of the power system. In addition, the time savings compared to building custom circuits is considerable.

In order to accommodate the voltage requirements of the various circuits, supplies for +12 V, +10 V, and ±5 V were needed. In the interests of ensuring the health of the robot against any fault that might occur in the payload (and vice versa, since it is conceivable that the eventual payloads could be considerably more valuable and sensitive than the robot), it is desirable to have the payload power supply and robot power supply separated as much as

possible. It is generally good practice
to provide a low power and high
power ground, as the higher current
levels in certain circuits (such as the
DC/DC boost converters) can generate
significant noise for lower power

Part Number	Input	Output
PT4523	V_{batt}	12 V_p, 20 W
PT4224	V_{batt}	12 V_s, 10 W
LM2937	12 V_s	10 V, 5 W
PT6101	12 V_s	5 V, 5 W
PT78NR105	12 V_s	-5 V, 5 W
PT5042	12 V_p	15 V[1]

Table 3-2 – Part numbers used to establish the various supply voltages in the Cool Robot.

(usually analog sensing and conditioning) circuits. These two grounds should be separated

from each other throughout the robot, and originate from one location ver close to the battery

negative terminal.

As the voltage supplies needed are relatively standard voltages, a large number of off-

the-shelf DC/DC converters are available. Many require a certain number of external

components, such as MOSFET switches or inductors. What's more, most of these devices

have a specific range of input voltages, most of which were well below the voltage of the bus

(48 V). However, a number of suitable converters were found from Power Trends[12] that

were, essentially, plug-in solutions that contained all the necessary circuits and components

(aside from some low-ESR capacitors) within a single module. Aside from their compact

form factor, these converters have a relatively high conversion efficiency between 70% and

80%, depending on the load. All of these models are temperature rated for -40˚C to +85˚C.

Within this product line, it was not possible to find converters that would step down the bus

voltage to such a relatively small voltage as ± 5 V. However, models existed that would

convert from the 12 V supply down to these levels. Table 3-2 provides information on the

devices used, and Appendix G-5 provides a circuit schematic. As one might notice, there are

two 12 V DC/DC converters. The first, rated for 20 W, is meant to power the payload, and a

---

[12] A subsidiary of Texas Instruments.

small number of circuits that should be isolated against failure from the rest of the power system. This is the "payload" supply, $+12\ V_p$. The second, from which the 10 V and ± 5 V supplies are derived, is the $+12\ V_s$ or "signal" supply. If it is necessary, future generations of Cool Robot could be outfitted with additional DC/DC converters to power more extensive scientific payloads.

In order to have clean supply rails, the housekeeping power supply circuits were constructed on a 1 oz copper-clad board. The large areas of copper would ensure a low resistance, low inductance path for current return back to the batteries. A printed circuit board would have been an equivalent option, but the simplicity of the connections and the higher cost of having such a board fabricated made this a less attractive option. The board was laid out on paper, and then islands were cut using a 1/32" ball end mill (what I refer to as "anti-traces"). The filter and bypass capacitors, sized according to the recommendations on the product datasheets, were soldered directly to the copper across the anti-traces. The filter capacitors (each for the +5 V and -5V converter outputs) consisted of three 33μF tantalum capacitors rated for cold operation. Ceramic 1 μF bypass capacitors were used throughout. In one or two places, high-value resistors were added to the converters' adjustment pins to tune their output level. In order to aid testing and allow robust connections between the input and output of the power board, banana plug binding posts were added to the board. Figures 3-14 and 3-15 show top and bottom views of the finished housekeeping power supply board, respectively.

Although the binding posts create mechanically robust connections to the supply voltages, a look at Table 3-1 above shows that a great many connections to these supplies are needed. In fact, there were too many connections needed to try and fit them all on the
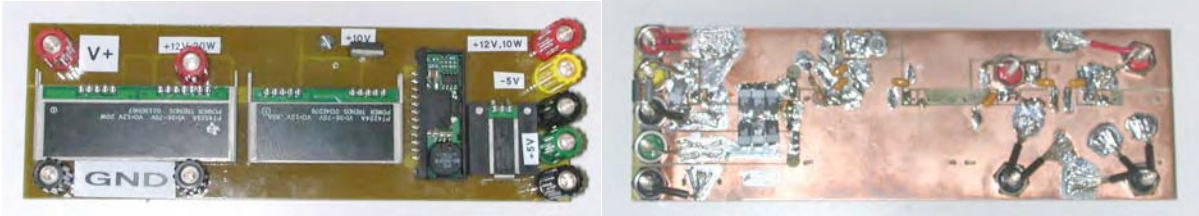
Figure 3-14 – Topside view of the housekeeping Figure 3-15 – Bottom view of the housekeeping power supply. The board measures 3 x 8 inches. power supply board.

housekeeping power supply in Figure 3-14. Therefore, a breakout board was fabricated. The supply voltages were laid out in parallel vertical busses, and each horizontal column has the following pinout:

Pin#	1	2	3	4	5	6	7	8
Function	Orient	GND	+12 V_p	+12 V_s	+10 V	+5 V	PGND	-5V

Each subcircuit, then, accesses those pins of the 8-pin connector that corresponded to the voltages needed. The first pin, labeled as "Orient," is extra long, and is there to ensure that the connectors for each subcircuit are oriented correctly into the connector. The original breakout board simple used a grid on exposed pins, to which each subcircuit could plug into using female MT connectors. Before deployment to Greenland, however, this setup will be replaced with more robust connectors that have a stronger, vibration resistant connection. Additionally, each supply line will be fused to isolate a downstream failure or short circuit. This redesign will be part of a larger overhaul of the wiring of the robot to make it more secure and rugged; less susceptible to vibration or jolts.

Future generations of Cool Robot should consider a redesign to the housekeeping power supply. The addition of a second 12 V, 20 W DC/DC converter to power additional payloads would be a worthwhile addition. Each of the 12 V converters is equipped with an "Inhibit" pin, which allows the user to disable to DC/DC converter's operation to conserve power. While this functionality was not utilized in the Cool Robot power system (each of the converters is always on), it may be beneficial to power up and power down some or all of the

102

payload power supply. Finally, any redesign of the power supply should consider adding more fault protection, such as separate fusing for the internal and payload supplies, so that a failure on one side cannot spread to the other. Adding fuses to each supply line on the breakout board will afford a good level of protection, short of remaking the housekeeping power supply board in Figure 3-14.

Solar Input Board

The operation and control of the DC/DC boost converters requires some knowledge of the operating point of each solar panel along its I-V curve. From a more practical standpoint, some central place to securely connect the solar panels' terminals is needed. Therefore, a board was created that, in addition to providing screw-down connection points for the solar panels and DC/DC boost converters, would also measure the panels' current and voltage. I refer to this as the solar input board.

The voltage of the panels was predicted by the model to be in the range of 30-40 V. Therefore, the voltage measurement of each panel begins with a resistor divider that reduces the panel voltage by a factor of 11.5. The signal is then passed through a second-order Butterworth low-pass filter with a bandwidth of approximately 50 Hz to remove any high-frequency noise that may exist in the signal. The sensitivity of this circuit is scaled to 87 mV/V. Figure 3-16 and Appendix G-3 provide a schematic of this circuit, which was replicated five times over, once for each panel.

Instead of a simple resistor-divider circuit, an improved design for this
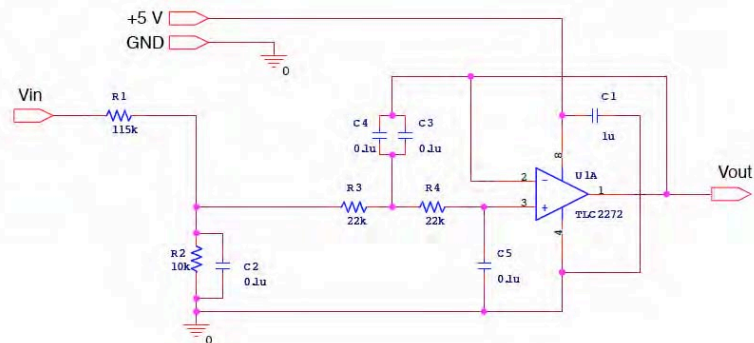
Figure 3-16 – Schematic view of the panel voltage measurement conditioning circuit.

circuit would disregard the large bias in the panel voltage and amplify only that 10 V range of interest between 30 and 40 V. Rather than the 87 mV/V sensitivity of the current implementation, such a circuit would have a sensitivity of 0.5 V/V. However, there was not time during the construction and testing of the robot to implement this more complicated circuit. However, being able to detect a panel voltage of less than 30 V may be an advantage, as it would help identify that something has malfunctioned within the panel or the boost converter.

The current coming from each panel could be up to 6 A. Measurements of currents this large could be made using a hall-effect sensor, which has two important advantages over the more typical method of measuring the voltage drop this current makes across a low-value precision resistor. The first advantage is that hall-effect sensors present almost no resistance to the current, and so have almost no influence on the circuit they are measuring. The second advantage is that hall-effect sensors electrically isolate the current path from the measurement circuitry. As a result, they can be directly connected to high voltage circuits, yet produce isolated low-voltage outputs. However, hall-effect sensors are relatively expensive (several dollars apiece), are better suited for measuring tens of amperes, and have poor accuracy and precision when dealing with small current measurements of less than a few amperes.[13] The precision-resistor method was utilized instead for the solar input board. For each panel, a 40 mΩ, 1% resistor was placed in line with the negative terminal. These resistors are a good compromise between low resistance and good temperature stability.

The current sensing resistor needs to be placed in series with either the positive or negative panel lead before they are connected to the DC/DC boost converter. However, the

---

[13] This is gradually changing, however, and hall-effect products may be viable options in the future.

positive panel lead sits at 30-40 V. While the voltage *across* the resistor would be small (on the order of 100 mV), the large DC bias on which is sits would be well outside the allowable input range of most op-amps. One could solve this problem by using a resistor-divider to scale down these two voltages, as was done to measure the panels' voltage. However, using a resistor divider to reduce the DC component by, say, 10x would likewise reduce the voltage drop across the sensing resistor by 10x, greatly reducing the overall measurement sensitivity. Therefore, the lowside measurement is utilized, where one end of the sensing resistor is connected to the system ground. In order for the DC/DC boost converter to work properly, it is necessary for its input and output to have the same ground as the rest of the system. Therefore the current sensing resistor is oriented such its signal is actually negative with respect to ground. To help illustrate this topology, the reader is directed to
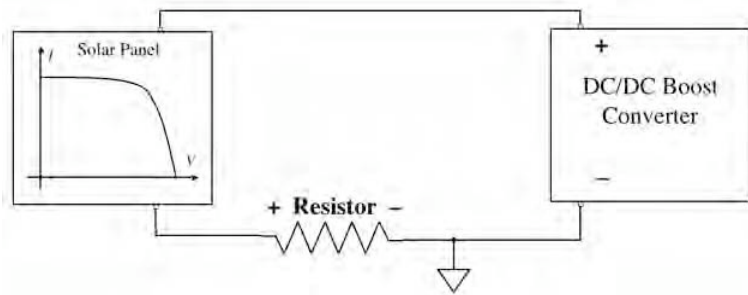
Figure 3-17 – Location and orientation of the current sensing resistor in relation to the solar panels and DC/DC converters.

Figure 3-17. The signal is passed through a Butterworth lowpass filter with a bandwidth of approximately 50 Hz to remove high-frequency noise. Because of the low value of the sensing resistor, the resulting signal is quite small, on the order of -100 mV. Therefore, the output of the Butterworth filter is passed through an inverting op-amp circuit with a gain of 12x, to produce a signal that is scaled to 0.48 V/A. The final circuit is provided in Figure 3-18.

An improvement to this circuit, which could be implemented in future generations of Cool Robot, is to use an instrumentation amplifier (or an difference amplifier op-amp circuit)
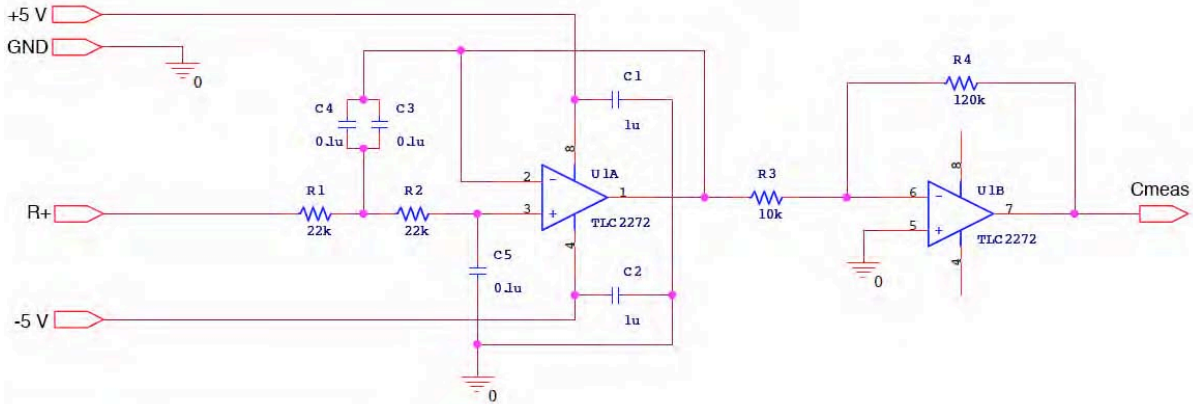
Figure 3-18 – The circuit used for measuring the panel current.

to amplify the voltage difference between power ground and the positive end of the resistor, both with respect to signal ground. In hindsight, I should have realized that there may be a significant difference between the analog ground of the sensing circuits and the power ground that the resistor is attached to. An instrumentation amplifier would have handily taken any difference between analog ground and power ground into account. The revised circuit could also implement a higher gain, which would make better use of the 5 V dynamic range of the ADC. Low pass filtration would have to be performed by an op-amp circuit downstream from the instrumentation amplifier.

The solar input board, aside from the connectors, required these ten dual op-amp circuits (current and voltage measurement for five panels). In order to reduce the footprint of so many circuits, surface-mount components were used in their implementation. Rather than have these circuits fabricated on printed circuit boards, which would have been difficult to modify, and would have been a considerable expense for just one or two prototypes, the circuits were implemented on surfboards. A surfboard, such as the one shown in Figure 3-19, is an generic
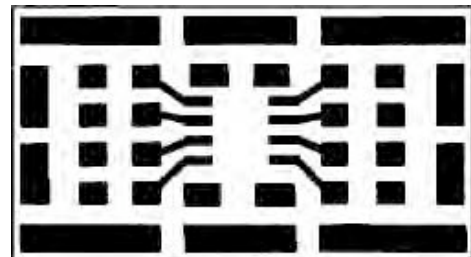
Figure 3-19 – Pad and trace layout of an 8-SOIC-based surfboard. Shown at twice actual size.

printed circuit board that contains the pads needed to mount an SOIC-sized chip, with a periphery of pads onto which one can attach 0805- or 1206-sized components[14] in various configurations to create circuits around the IC. A ground plane of copper foil is added to the rear side with adhesive tape. Images of the finished voltage and current sensing circuits are provided in Figures 3-20 and 3-21, respectively.
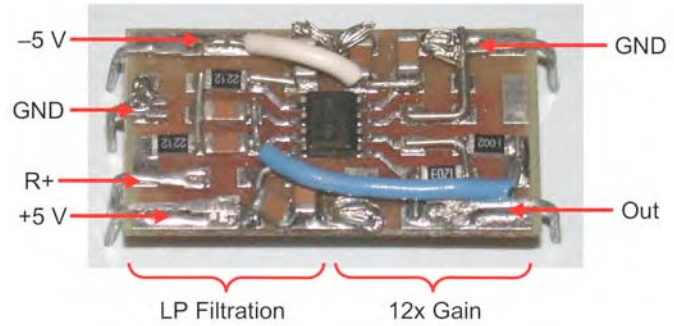
Figure 3-20 – Annotated view of the surfboard implementation of one current measurement conditioning circuit.
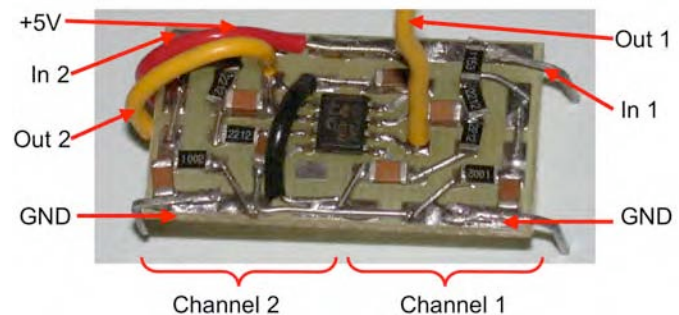
Figure 3-21 – Annotated view of the surfboard implementation of two panel voltage measurement conditioning circuits.

In future generations of Cool Robot, it is recommended that the current and voltage measurement circuits be implemented on one or two printed circuit boards. Using a custom-designed PCB would allow these circuits to be produced with much less labor and in a smaller footprint. Using PCBs would also provide an opportunity for redesign, so that the improved measurement circuits could be implemented. Lastly, it would be beneficial to have the analog-to-digital converter for these 10 signals located on this same board, so that the conditioned analog signals can be immediately digitized. In the current

---

[14] This notation is used to describe the physical size of the surface mount components such as resistors and capacitors, which appear as tiny monolithic slabs in finished circuits. The first two digits are the width in hundredths of an inch, the second two are the length. So, a 1206 component measures 0.12 x 0.06 inches. The height varies depending on the component.

implementation, the analog signals are brought to the ADC board over wires, which can pick up electronic noise in the process.

One last design detail for the solar input board is the use of chokes. These devices are made from materials similar to those used in inductor cores. They are used to increase the high-frequency impedance of the wires that pass through them, so that high frequency noise on that line (in the MHz range) is dampened. In the case of the solar input board, these were used to prevent any high-frequency noise from the DC/DC boost converters from propagating back towards the solar panels. While this step was probably unnecessary, it was done as a precaution, since it was unknown at the time what effect the high-frequency noise might have on the operation and long-term life of the solar panels. Future generations of Cool Robot could do without them. The finished solar input board is shown in Figure 3-22, and schematics are provided in Appendix G-3.

POWER DISTRIBUTION BOARD
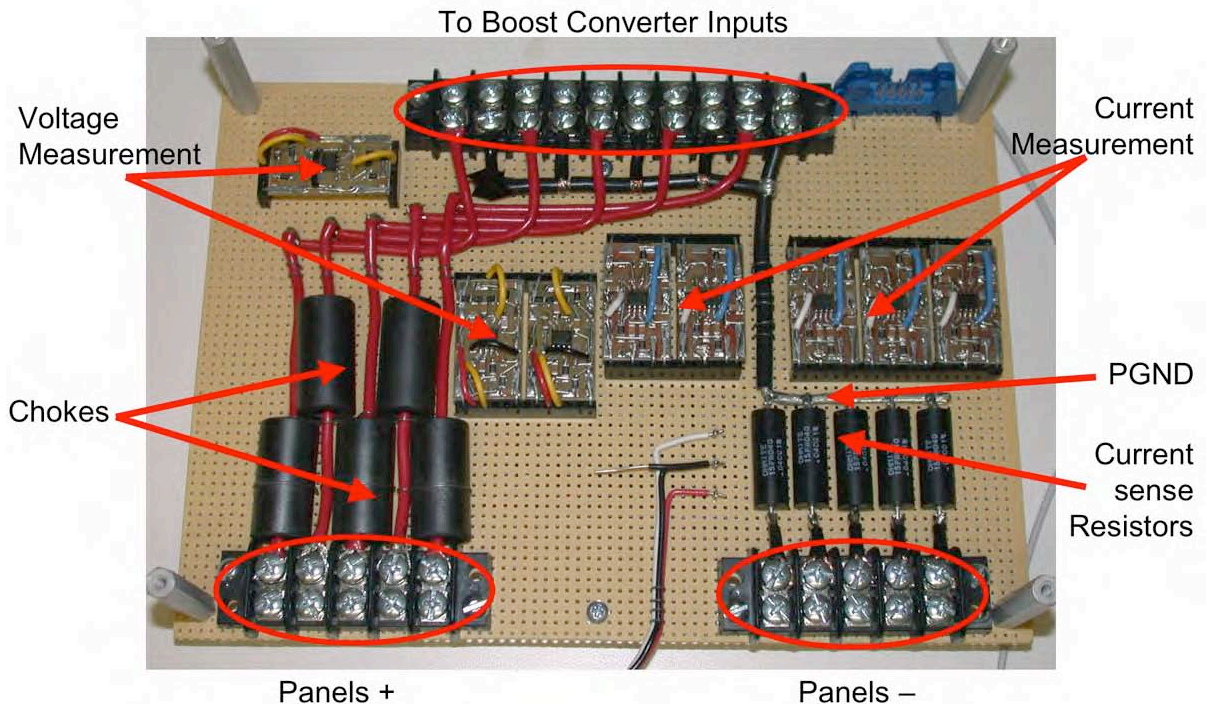
Figure 3-22 – Annotated view of the completed solar input board.

Once the power from the solar panels has passed through the solar input board and been measured, it moves on to the DC/DC boost converters. The outputs of these five converters need to be connected to the battery bus, then passed on to the housekeeping power supply and motor controllers. For control purposes, it is necessary to provide a means for isolating the flow of power among these components. Also, the current in and out of the batteries, and their combined voltage, needs to be monitored. These functions are implemented on the power distribution board.

There are two breaks in the power bus which are used to control the flow of power (see Figure 3-23). The first lies between the DC/DC boost converters and the batteries.
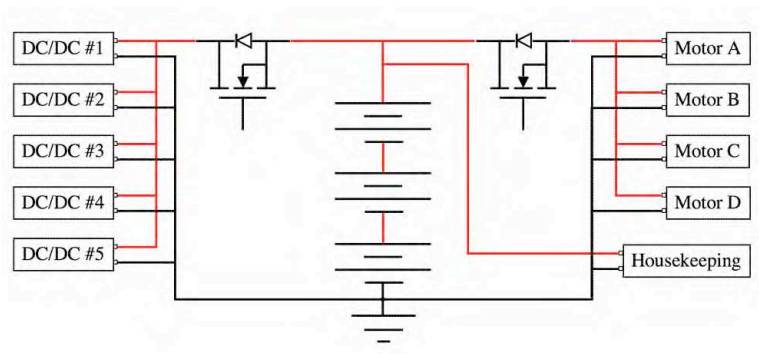
Figure 3-23 – Basic architecture of the common power bus.

While this switch is normally closed, there are several conditions under which is might be opened. One occasion for disconnecting the solar input would be to run the robot entirely off battery power, which might be desirable during times when the payload is making measurements that require the robot to run "electromagnetically quiet." While there was not time to implement it during this thesis, an automatic circuit, faster than the microcontroller, that could open this MOSFET switch if it detected an excessive charge current, would be a prudent addition to the power system. Whenever this switch is opened, it is necessary to disable the DC/DC converters. If the DC/DC boost converters continue running with their collective output disconnected, they would continue to charge up their output capacitors until either the capacitors or Schottky diode broke. It is significant to note that, due to the

placement of this switch and the rest of the architecture of the power system, it is impossible for the DC/DC boost converters to deliver power to the robot without being first connected to the battery bus. That is, whenever the DC/DC converters are operating, their output voltage will be clamped to whatever the battery voltage is. In hindsight, the addition of this Mosfet switch was probably redundant, as disabling the DC/DC boost converters accomplishes nearly the same thing.

The second breakpoint in the bus lies between the batteries and the motor controllers. Such a break could be necessary to limit power consumption. Even when the motors are not running, the motor controllers consume significant power.[15] Secondly, while it should not be necessary to do so, it may be necessary to disconnect the batteries from the motor load to prevent the batteries from being overdischarged. Lastly, like with the DC/DC boost converters, disabling the motors so that the robot will result in a substantial decrease in the amount of electromagnetic noise in and around the robot.

Both of these switches are implemented using the same 1010EZ MOSFETs that are utilized in the solar boost converters. However, unlike the way they are used in the boost converters, these MOSFETs operate in a high-side topology. In the case of the boost converters, the source pins of the MOSFETs are tied to ground; a so-called "low-side" topology. Therefore, in order to raise V_{gs} up to the desired 12 V to turn the MOSFET on, one need only supply 12 V to the gate. In the case of the MOSFETs on the bus, however, the source is not tied to ground. This situation, when the source is not grounded-referenced, is referred to as a "high-side" topology. In the case of the MOSFET that isolates the solar power input, the source is always at V_{batt}. In the case of the second MOSFET, the source is either at

---

[15] Without even measuring the quiescent power draw directly, one can tell this from the fact that the motor controllers are warm to the touch, even when the robot is not moving.

V_{batt} when power is being supplied to the motors,[16] grounded through the motor controllers when the switch is open, and somewhere in between during a switching transition. In order to make a MOSFET conduct, it is necessary to provide the gate with a voltage that is at least a threshold voltage above the source, and track the changes in the source voltage. The threshold voltage of the 1010EZ is approximately 3 V, although a V_{gs} of 10 V or more is desirable for good conduction with low power loss (see Appendix H-2). In the DC/DC boost converters, the 1010EZ is driven with a V_{gs} of 12 V. To do the same with the MOSFETs on the bus, a gate voltage of up to 60 V is required.

In certain DC/DC switching converters, a high-side connected MOSFET is unavoidable. In the case of this main power bus, a low-side topology for one or both switches would have been possible. However, I made the choice to use a high-side topology for these two switches so that the DC/DC converters and the motor controllers would always have a solid connection to the system ground. There are several ways to construct robust and fast high-side driving circuits for MOSFETs, and the reader is directed to [IRF 2003] for a good design reference. However, there is at least one easy way to provide the necessary higher voltage for such circuits. Consider the circuit provided in Figure 3-24. The IC in the center is a high-side FET driver which, depending on the logical input, alternately connects the MOSFET gate to the source (whatever V_s may be) or to V_B, which lies at one end of the capacitor C. The other end of
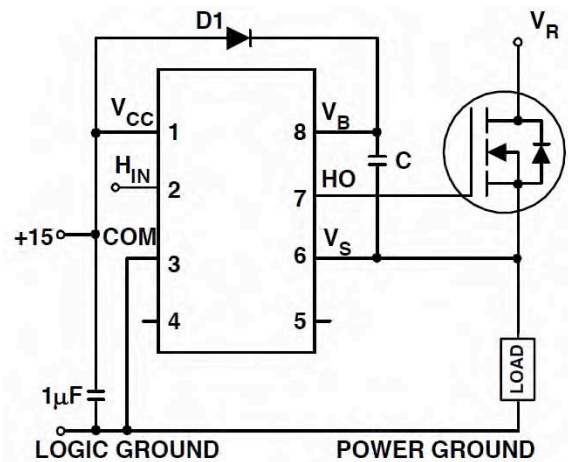
Figure 3-24 – High-side drive circuit for a switched circuit centered on a bootstrapped capacitor [IRF 2003].

---

[16] Minus a small drop due to $I_{motors}R_{DS,on}$

the capacitor is connected to the source. When the logical input (H_{in}) is low, the gate and source are connected by the FET driver, so $V_{gs} = 0$ The MOSFET is therefore off, the source is grounded through the load ($V_s = 0$), and the diode *D1* located above the FET driver becomes forward biased and charges the capacitor to $V_{cc} - V_{diode}$, or about 14.3 V in this example. Therefore, $V_B = V_s + V_{cc} - V_{diode} = 14.3$ V.

When the logical input goes high, the FET driver connects the MOSFET gate to V_B, which causes $V_{gs} = 14.3$ V, which causes the MOSFET to turn on. The conduction from drain to source causes V_s to rapidly rise to $V_s = V_R$. Because the negative end of the capacitor (called a "bootstrap capacitor") is tied to the source, its voltage also rises to V_R. The charge on a capacitor cannot change instantaneously, so the positive end rises in turn, always remaining $V_{cc} - V_{diode}$ above the voltage of the source. Therefore, when the device is fully on, $V_B = V_s + V_{cc} - V_{diode} = V_R + 14.3$ V. Therefore, when the low-to-high transition is finished:

$$V_{gs} \equiv V_g - V_s = V_B - V_R = V_{cc} - V_{diode} = 14.3 \text{ V} \tag{3.6}$$

In this way, one can keep V_{gs} at a constant level throughout the low-to-high transition.

Ideally, this sort of circuit would work for the Cool Robot, where the switches are nearly always conducting, because the bootstrap capacitor will never discharge, and so V_B will remain 14.3 V above V_s. However, all semiconductor devices have a small amount of leakage, and eventually the bootstrap capacitor would discharge through the MOSFET gate and bootstrap diode, causing V_B to decay. The drive circuit works for a switching regulator, however, because thousands or millions of times per second, the MOSFET is turned off, the source becomes grounded, and the capacitor is recharged from V_{cc} through the bootstrap diode. It is conceivable that one could still use this drive circuit for the Cool Robot, if one

were to turn each of the two switches off and back on again at regular intervals (as infrequently as once every 10 seconds or longer, depending on the size of the bootstrap capacitor). However, this would cause disruptions to the driving and solar panel control, and should be avoided if a different solution can be found.

The alternate solution uses the same circuit as Figure 3-24, but continuously keeps the bootstrap capacitor (and, by extension, V_B) charged up using a circuit called a charge pump. This circuit is given in Figure 3-25. The circuit centers around a 555-timer chip, which is configured to produce a regular square wave. When the clock signal is low, the 10 nF capacitor becomes charged
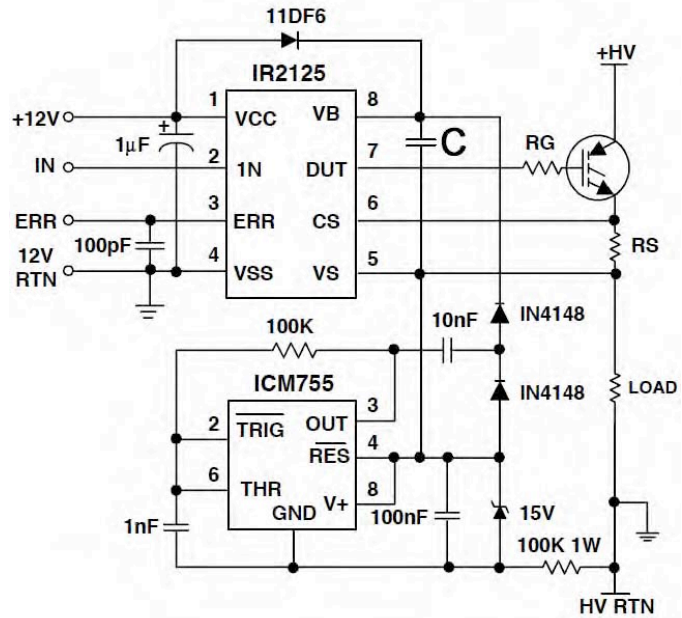
Figure 3-25 – A continuous high-side drive provided by a charge pump circuit [IRF 2003].

from V^+ to $V^+ - V_{diode}$. When the clock output goes back high, the lower side of the capacitor is raised from GND to V^+, so that its higher end now resides at $2V^+ - V_{diode}$. This is another bootstrap operation. The second diode allows this bootstrapped capacitor to provide a trickle of charge at a voltage $2V^+ - 2V_{diode}$ to circuits downstream.

It is important to note that this charge pump circuit's supply voltage is the source voltage V_s, and its "ground" is formed with a Zener diode. The Zener diode forces a certain potential difference between the V^+ and GND pins of the 555. Whatever V^+ (or more accurately, V_s) is, the GND of this circuit is always $V^+ - V_{zener}$. In this way, the Zener diode

has created a virtual ground, which is able to track changes in V_s. For instance, if $V^+ = V_s = V_{batt} = 48$ V and the diode is a 12 V Zener, then the GND of the charge pump circuit will reside at 36 V. In those cases when $V_s < V_{zener}$, such as when the source is grounded, the Zener diode is not sufficiently reverse biased to conduct, and the 555 is not supplied with any power. In that case, the charge pump circuit is disabled, and the main bootstrap capacitor C is charged from +12 V. The operational voltages of the charge pump, explained in the previous paragraph, are made with respect to the virtual ground $V^+ - V_{zener}$. So, continuing on with our example of $V^+ = V_{batt} = 48$ V and $V_{zener} = 12$ V, the output of the charge pump is:

$$V_{pump} = \left(\left(2V_{cc} - 2V_{diode} \right) - \left(V_{cc} - V_{zener} \right) \right) = V_{cc} + V_{zener} - 2V_{diode}$$
$$V_{pump} \approx 48 \text{ V} + 12 \text{ V} - 1.4 \text{ V} = 58.6 \text{ V}$$

(3.7)

The charge pump ensures that the bootstrap capacitor of the FET driver is continually charged up so that $V_B = V_s + 10.6$ V. Using the circuit in Figure 3-25 as a guide, the drive circuitry for the two bus switches was constructed. The final circuit, which was replicated for both MOSFETs, is given in Appendix G-4. The FET driver used is an International Rectifier 2117, whose datasheet is given in Appendix H-5.

The logical inputs that control these two switches is the logical sum of a number of different logical conditions, and will therefore be the output of some logic gate. The logic level will be either 3.3 V (directly from the microcontroller pin) or could be 5 V from something like a 74LS discrete logic chip. However, the 2117 requires that its logical input be of the same level as its supply voltage: 12 V. Therefore, it is necessary to shift the logic level from 5 V up to 12 V. This could be accomplished with some sort of open-collector device, like a comparator. There also exist various ICs that can perform logic-shifting

functions from dual voltage supplies, most translate 3.3 V to 5 V logic, and vice versa, and may provide the best solution in future designs. The power distribution board uses the circuit provided in Figure 3-26 (and
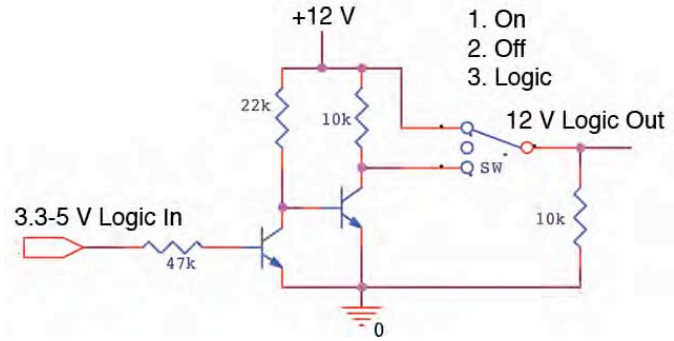
Figure 3-26 – The logic-shifting circuit and manual override switch.

Appendix G-4). This is a resistor-transistor logic buffer, which can perform the 5→12 V logic shifting operation with a delay of only 320 ns on both the rising and falling transition. A similar circuit for making a 12→5 V logic shift would use a similar circuit, though with different resistor values. It is desirable to have some way to manually override the logical signal during testing (for instance, a kill switch). A toggle switch was added to the power distribution board that allows the user to switch between Always On, Always Off, or Logic Controlled. This three-way function is implemented using a SPDT switch and is given as part of the circuit in Figure 3-26.

Another function of the power distribution board is measuring and scaling the battery stack voltage. The battery stack, depending on the state of charge (SOC) of the three individual lithium ion batteries, can vary from as little as 36 V to as much as 49.6 V. The batteries will hopefully never be allowed to discharge so fully, so we set a lower operational and measurement bound of 42 V (14 V per battery, a less than 10% state of charge). The upper end of the measurement range will be 50 V. In order to fit within the 5 V dynamic range of the analog-to-digital converter, one could simply scale this voltage by a factor of $1/10^{th}$. However, to do so would waste more than 80% of the measurement range, as the measurement signal would vary only between 4.2 and 5 V. It is better to scale and measure

only the 42-50 V range of interest; that is, create a conditioning circuit that would scale a battery voltage of 46 V to a measurement voltage of 2.5 V, 42 V to 0 V, and 50 V to 5 V.
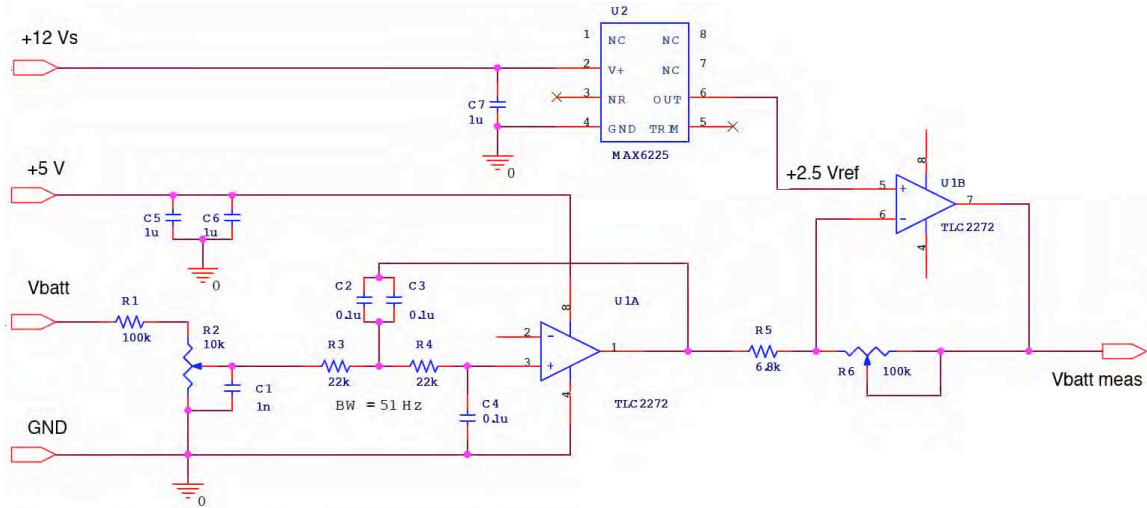
Figure 3-27 – The circuit used to measure and amplify the battery stack voltage.

The circuit given in Figure 3-27 (and Appendix G-4 in larger format) implements this function. The battery voltage is first scaled by a factor of $2.5/46 = 0.05434$ using a resistive divider. Then, the difference between this scaled voltage and a precision 2.5 V reference is scaled by an inverting amplifier. The DC transfer function of the circuit is given by:

$$V_{out} = 2.5 \text{ V} - \frac{2.5}{4}\left(V_{batt} - 46 \text{ V}\right) \tag{3.8}$$

Which not only utilizes the full 5 V measurement range of the ADC, but has a measurement sensitivity of 0.625 V/V, whereas simply scaling the battery voltage would have resulted in a sensitivity of just 1/10. When sampled using the ADC, I found the measurement accuracy of the circuit to be better than ±50 mV across the whole 42-50 V range, with a precision of better than ±20 mV.

In addition to the battery voltage, the power distribution board also needs to measure the battery current. If the robot is being powered entirely from the batteries (i.e., there is no

116

solar power provided), then the battery current will be, at most, 3 A for each motor and approximately 3/4 A for the housekeeping supply, for a total of 12.75 A. Usually, the actual current requirement would be about half that. Although it is difficult to envision a situation where it would be necessary, the batteries can provide a 5 second pulse of up to 36 A. Meanwhile, the maximum charging current for the batteries will be 3 A. We will define the convention that current leaving the battery is positive, and charging currents are negative. So, we need to be able to reliably measure a bipolar current over a range of 15-16 A, and scale the measurement so that it fits within the 5 V measurement range of the ADC. A hall-effect sensor is used for this current measurement. The current sensor used for this purpose is the ACS704ELC-015 from Allegro Microsystems (see Appendix H-10 for datasheet). This sensor has a measurement range of ±15 A, with a sensitivity of approximately 100 mV/A, and an analog output centered on 2.5 V. At the same time, it presents a resistance to the circuit of approximately 1 mΩ. It is also electrically isolated from the current path that it is measuring, so that it can be directly connected to the power bus, yet be powered by and

connected to a more sedate 5 V circuit. As Hall Effect sensors go, it also has good precision, accuracy, linearity, and bandwidth. As with the battery voltage sensing circuit, the output from the measurement needs to be lowpass filtered, though no
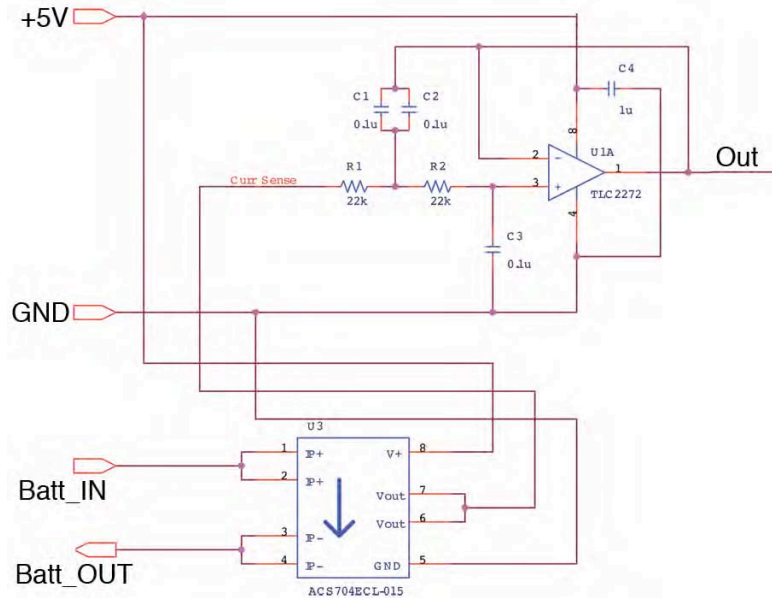
Figure 3-28 – Circuit used for measuring and amplifying the battery current.

117

amplification is used. The finished measurement and conditioning circuit is shown in Figure 3-28, and in Appendix G-4.

Connections are also provided on the power distribution board for the motor controllers and housekeeping power supply. Each positive motor connection has a 3 A slow-blow fuse in series, which at the colder temperatures the robot will operate at in Antarctica should only blow when a motor controller substantially exceeds its preprogrammed current limit for several seconds. Screw-down terminal connectors are used for both the power and ground of each motor, while housekeeping supply receives its power through banana plug wires. The maximum rated power of the housekeeping power supply is 30 W, so a 3/4 A
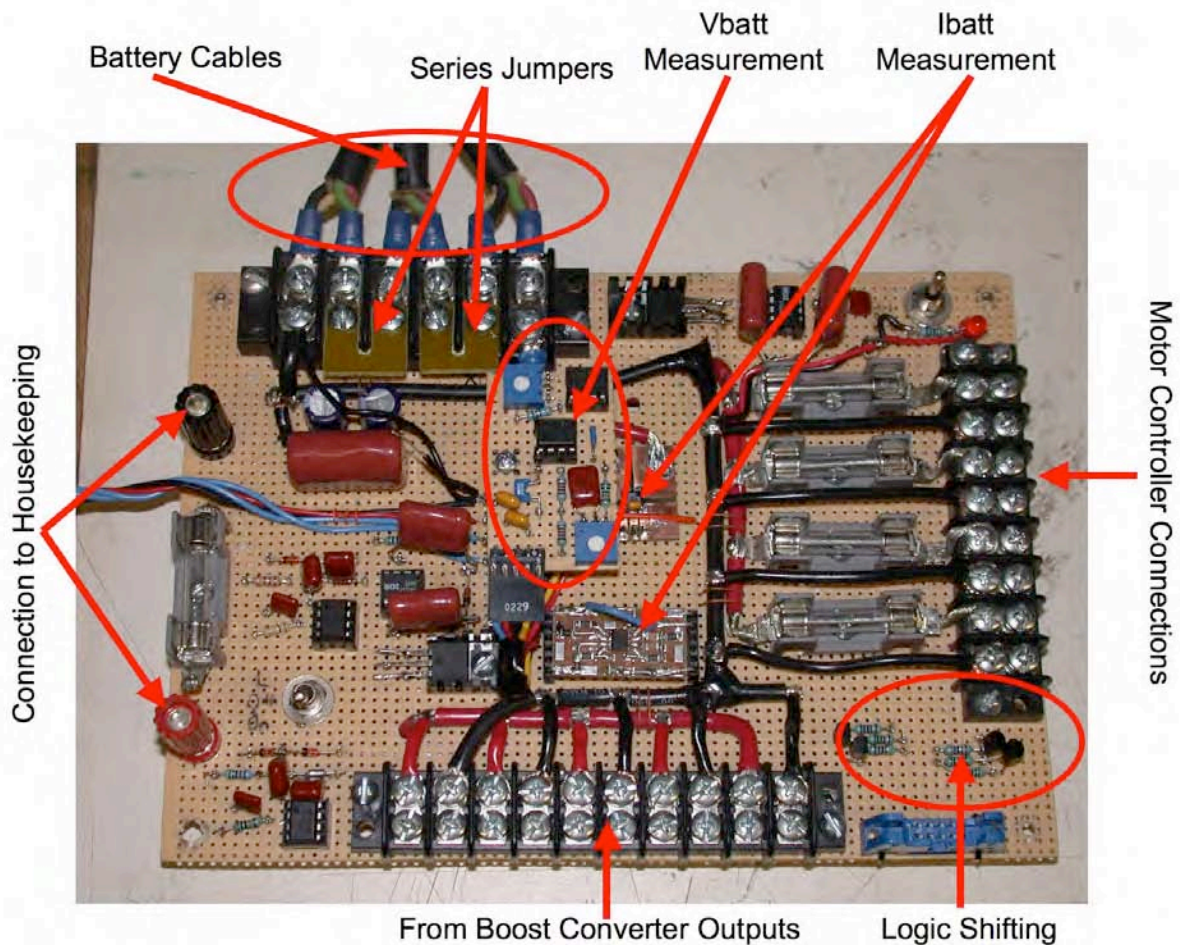
Figure 3-29 – Annotated view of the completed power distribution board.

fuse is inline with this connection. The complete power distribution board is shown in Figure 3-29.

BATTERIES

The energy storage and backup power device on the Cool Robot is a stack of Lithium-Ion (Li-ion) rechargeable batteries. Li-ion batteries have the highest power density – both energy per volume and energy per mass – of any widely used rechargeable battery chemistry. It is the same chemistry that is used for laptop and cell-phone batteries. A complete discussion of the Li-ion chemistry is beyond the scope of this thesis. However, because understanding of some of the characteristics of the Li-ion chemistry is crucial to the development of the control system, some of the important points will be covered here. A comparison table of Li-ion to several other mainstream chemistries is given in Appendix E.

Li-ion batteries have several advantages over other rechargeable chemistries, aside from their greater energy density. Unlike some chemistries, Li-ion batteries exhibit little self-discharge, which gives them a long shelf-life. Secondly, Li-ion batteries do not exhibit a memory effect [Buchmann I-5]. The memory effect is a problem generally associated with nickel-based chemistries such as nickel-metal-hydride (NiMH) and nickel-cadmium (NiCd) [Buchmann I]. The memory effect comes about from repeatedly discharging a battery only partway before recharging. Over time, the unused capacity becomes permanently unavailable. To avoid this situation, one needs to periodically discharge nickel-batteries fully to exercise their full capacity [Buchmann II].

However, Li-ion batteries are perhaps the most fragile and finicky chemistry available. Deep discharging or overcharging can be damaging to any battery chemistry, but the risks with Li-ion batteries are substantially greater. Deep discharging a Li-ion battery can

result in it being impossible to fully recharge. Overcharging will result in plating the electrodes with metallic lithium, which seriously degrades the battery's capacity, and can lead to the battery exploding [Buchmann II, UBI 2004, Appendix H-16]. Rapidly discharging a Li-ion battery can also result in lithium plating [Buchmann I]. The best way to prolong a Li-ion battery's life is to partially discharge it, less than halfway, then top it off again. For these reasons, one will almost always find Li-ion batteries contained within a pack that also has a number of protective circuits that can disconnect the battery from the pack's terminals.

Lithium Ion batteries have stringent charging requirements. A NiCd battery can be attached to a voltage source for charging, and sink a high amount of current as it does so until the battery's voltage reaches that of the charge source. A Li-ion battery, by contrast, can only safely accept a limited amount of current as it charges. The limitation stems from the speed at which the chemical reaction can be run in reverse to store energy. Too high a charge current can result in a permanent loss of capacity. For this reason, the charge cycle of a Li-ion battery is general divided into two controlled segments. The first segment is the constant-current (CC) phase, where a charge current up to a certain limit can be applied to the battery. As this happens, the voltage of the battery will rise in an almost linear fashion, much like a capacitor. At some point, the battery voltage will approach its final value, and the battery enters the constant voltage (CV) stage of the charge cycle [Buchmann I, UBI 314], where the charger asserts the final desired battery voltage, and the current draw begins to decrease (as it would with a charging capacitor). Eventually, when the charge current has decreased to some arbitrarily small amount (oftentimes 1/10[th] the constant-current limit), the charge cycle is declared to be finished [Appendix H-16]. The battery charger, then, acts

much like a bench power supply that tries to create the desired voltage across its terminals, but can only provide a certain limited current. In fact, a very crude way to charge a Li-ion battery is to simply attach it to a bench supply that has been current-limited. In laptop and cellular phone batteries, the charge circuitry is often contained within the pack, along with all the protection circuitry, so all one need do is provide it with bulk power [Maxim 189, 344]. In other, less sophisticated batteries, the charge circuitry is external to the battery [UBI 5133, 5134].

The typical electrochemical potential of a single Li-ion cell is approximately 3.6 V. For some applications, such as cellular phones, this is adequate. In other applications, such as laptop batteries, multiple cells are placed in series within the battery pack, increasing the voltage at the terminals in multiples of 3.6. Cells are wired in parallel to provide a higher current-producing capacity. One problem that multi-cell Li-ion batteries exhibit is drift. When the battery pack is brand new and the cells are charged, each cell in series will exhibit the same voltage. However, as the battery is partially discharged and recharged over and over, the cells in series may begin to charge or discharge at different rates. The end result is a situation where one or more cells has a significantly different voltage than the others. The greater the number of cells in series, the greater potential for drift. Anything greater than four cells in series is generally avoided; when one can only measure the terminal voltage of the entire series string, the differences from one cell to the next are unobservable. The only way to correct cell drift is to slightly overcharge or overdischarge the battery to force all the cells to the same level. This solution, however, degrades the battery's capacity, as it inevitably causes damage to some or all of the cells. One way to mitigate the problem is to occasionally fully cycle the batteries, rather than merely partially cycling. Another way is to

be able to monitor and individually charge each of the individual cells, which is the solution sometimes seen in computer batteries.

Li-ion batteries are relatively tolerant to a wide range of operating temperatures, but only to a certain extent. The dangers of over charging or discharging are greatly exaggerated by temperature. In cold temperatures, the electrolyte, which is a thick liquid supposed to help move charges between the anode and cathode, begins to gel to the point where it can no longer move charges effectively and the battery stops functioning. At elevated temperatures, it is far easier to overcharge the battery, causing a reduction in the battery's capacity. What's more, at elevated temperatures, Li-ion batteries (like many chemistries) start to produce a small amount of gas. Because the batteries are sealed, however, this gas cannot escape unless it builds up to such a degree that the cell explosively releases it. In order to prevent this from happening, the protective circuits associated with Li-ion batteries will often also guard against excessive temperatures or pressures.

The batteries used in the Cool Robot are three UBI-2590 from Ultralife Batteries (Figure 3-30, Appendix H-16). These are four-cell batteries, whose terminal voltage is nominally 14.4 V (min 12 V, max 16.4 V). Each battery actually contains two packs (A and B), which can be connected in parallel for a 14.4 V, 12 Ahr battery, or in series for a 28.8 V, 6 Ahr battery. In the Cool Robot,

Figure 3-30 – The Ultralife 2590 Lithium-Ion battery used in the Cool Robot.

122

they are wired in parellel. Together, the three batteries can provide enough power for several hours of routine driving. These batteries are commercial versions of an identical battery used by the military. Each battery weighs 1.44 kg, and measures 127 x 112 x 62 mm. A special cable is needed to interface with a special 5-pin connector on the battery. The other end of the cable yields the positive and negative terminals of the A and B packs. These terminals are connected to the power distribution board with heavy-duty screw down connectors. Two pieces of copper-clad board make up jumpers that connect the three batteries in series. Please refer to Figure 3-29.

The manufacture's recommended charging cycle is to assert 16.4 V per pack, current limited to 3 A, until the current dies off to 300 mA. It is possible to charge the A and B packs independently, in which case the charge current is 3 A per pack. The manufacturer's battery charger implements this independent charging function. In the case of the Cool Robot, the two packs are not independent, but rather wired together in parallel. In that case, we do not have independent control of the voltage of the A and B packs of each battery. Indeed, we don't even have independent control of the voltage of each of the three batteries. In this situation, the manufacturer recommends a 3 A charge current. In that case, the charge profile for the whole battery stack, given in Figure 3-31, is to assert a voltage of 49.2 V, while limiting current to 3 A. In this case, the CC phase of the charge cycle will bring the battery stack up to over 80% of its full capacity. As I will describe in greater detail in chapter 4, it is necessary for the battery to only be partially charged, so that it can accept as well as deliver power. Therefore, the batteries will almost always be operated with the 3 A charge current limit.

One final function that a battery charger must take into account is temperature. As has already been explained, it can be damaging to the battery, and dangerous to the user, to charge or operate a battery at high temperatures. In the case of the UBI2590, the operational temperature range is -20°C to +60°C.

Figure 3-31 – The desired charging profile for the Cool Robot battery stack, which is identical to the manufacturer's recommended profile, but expanded by a factor of three in both time and voltage.

Batteries are heated by their environment but also by their operation. As current passes in and out of the battery, it creates heat through a lumped parameter called the series resistance. All batteries have a certain amount of series resistance, and is related to the chemistry and construction, rather than a physical resistor within the battery. Particularly during charging, Li-ion batteries must have their temperature monitored and, if it rises too high, be disconnected from the charge source. In the Cool Robot, the batteries will be contained in their own insulated enclosure, to ensure that they are warm enough to produce power-on-demand. A temperature sensor will be placed in this enclosure, which will allow the power system to monitor the battery temperature. If the battery temperature becomes too high (which is unlikely, considering the Antarctic environment, the power system bring the system to a halt and disconnect both the motors and the solar power input, to allow the battery to passively cool. If the battery temperature becomes too cold, then some sort of heating

element within the enclosure will be turned on to warm them. However, there was not time during this thesis to implement this functionality, and must be left to future work.

ADC, DAC

In order for the control algorithm developed in Chapter 4 to function, the power microcontroller needs to be able to interact with all of the circuitry described to this point. The interface between the digital processor and the analog world are the analog-to-digital and digital-to-analog converters (ADC and DAC, respectively). As the names suggest, these devices translate analog voltages into digital numbers, and vice-versa.

All of the analog signals of interest to the power microcontroller – the panel currents and voltages of the five solar panels, and the battery current and voltage – are digitized by a Analog Devices AD7490 (see Appendix H-7 for datasheet). This is a 12-bit, 16 channel device.[17] The user supplies an external voltage reference, such as 2.5 V, 4.096 V,[18] or 5 V, and the AD7490 is capable of converting the data in the range of 0 to V_{ref}. In the case of a 2.5 V reference, the AD7490 can produce a two's-compliment output in the range 0 to 2 V_{ref}. The analog input circuitry and the digital output circuitry can be powered by two different voltage supplies. In the Cool Robot, 5 V is used for the analog side, and 3.3 V for the digital side, as it needs to communicate with the 3.3 V microcontroller. The communications between the two devices is handled by the microcontroller's serial port B, and uses Serial Peripheral Interface (SPI). SPI is a widely-used protocol among integrated circuits, so a library for using SPI on the RCM3100 microcontroller was readily available from the Z-World website. That library was augmented and improved for use by the Cool Robots

---

[17] While we could have gotten away with using just a 12 channel ADC for the power system, the master microcontroller uses the 7490 to monitor such things as the motor currents and wheel speeds, and does require all 16 channels. Using the same device for both allows us to reuse code that had already been developed for the navigation system.

[18] While this may seem like an odd reference, It makes sense when one views it as 2^{12} mV.

project by Goetz Dietrich during his development of the navigation system [Dietrich and Zettl 2005].[19] One alteration that I made to that original code was to extract it from the navigation and driving libraries and place it in its own library, `AD7490EV.lib`, which is presented in Appendix I-1. By doing so, the driving and navigation code is device-independent, and the code could be used by the power microcontroller without redundancy.

The actual format of the data that passes between the microcontroller and the AD7490, such as the command words and digitized data, is tedious and largely unimportant, and the interested reader is referred to the datasheet in Appendix H-7, and to [Dietrich and Zettl 2005]. A breakdown of the channel assignments for

ADC Channel #	Corresponding Signal
0-4	Panel Voltages 1-5
5-9	Panel Currents 1-5
10	Battery Stack Voltage
11	Battery Current
12-15	(currently unused)

Table 3-3 – Analog channel assignments for the AD7490 evaluation board

the power system's AD7490 is provided in Table 3-3.

The AD7490 comes in a 28-pin TSSOP package, which is quite small to use in prototype circuits. The 0.025" pitch of its pins is generally beyond the soldering skill of human hands. Therefore, in order to aid the integration of the device into the system, an evaluation board was purchased which already had the AD7490 mounted to is, along with a number of peripheral circuits and connectors. The ancillary circuits include a 96-pin connector, ICs for creating or accepting the reference voltage, and op-amp buffer circuits for each of the 16 channels. These allow the evaluation board to be flexibly interfaced with other circuits, specifically a separate microcontroller board from Analog Devices that allows

---

[19] One notable change that was necessary was the creation of an SPI library for each serial port, as Dynamic C would not allow the simultaneous use of a single library by all the serial ports.

one to program and operate the AD7490 evaluation board through a Windows computer. The digital signal connections between the AD7490 evaluation board and the Jackrabbit microcontroller are given in Table 3-4 The evaluation board is shown in Figure 3-32, and its datasheet is given in Appendix H-8.

RCM3100 Pin Name	Pin Function	AD7490 Eval Pin Name	Pin Function
PB0	CLKB	J1-A7, C7	SCLK
PD0	GPIO	J1-A6, C6	/CS
PC4	TXB	J1-A5	Din
PC5	RXB	J1-C5	Dout

Table 3-4 – Digital signal connections between the Jackrabbit RCM3100 microcontroller and the AD7490 evaluation board.

One downside of using the evaluation board, aside from the extra bulk of all those unnecessary circuits and redundant connectors, is that the signal input connectors on the board are Sub-Miniature BNC (SMB) type. This type of connector is typically
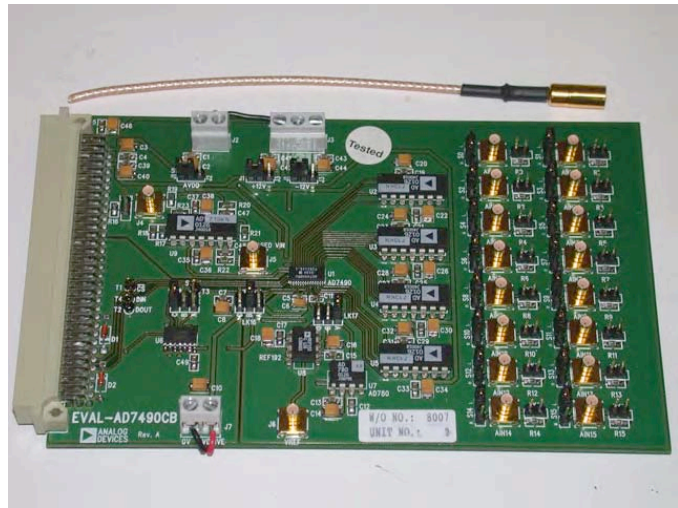
Figure 3-32 – The AD7490 evaluation board, shown with one of the SMB cables used to provide it with analog inputs. The ADC is the TSSOP-28 in the center of the board, with many traces leading to it.

used for radio-frequency and small-signal applications where parameters like cable impedance and noise shielding need to be critically controlled. In this application, connectors of this quality are superfluous, and require the use of costly cables. Another odd design feature of the evaluation board is that the op-amp buffer circuits require a bipolar supply, although the signals being digitized can only be single-ended.[20] As the analog signals going to the ADC evaluation board have all been conditioned through op-amp

---

[20] The reason for this is that the op-amps are general purpose, not rail-to-rail. These may be replaced with rail-to-rail op-amps in the future.

circuits, these buffers are hardly even necessary. Future generations of Cool Robot would benefit from a custom-designed circuit board that has the current and voltage signal conditioning circuits and the ADC on the same board. Not only would this help reduce the size of the measurement circuitry, simplify its construction, and eliminate the unnecessary circuits and connectors of the ADC evaluation board, but would eliminate the costly SMB wires and connectors through increased integration.

As explained earlier in this chapter, the DC/DC boost converters that control the operation of each solar panel receive their duty cycle commands in the form of an analog voltage between 0 and 5 V. That analog voltage is provided by a TLV5630 from Texas Instruments. This is an 8-channel, 12-bit digital to analog converter. It can accept an external reference voltage or use one of two internal reference voltages (1.024 V and 2.048 V). In any event, the output range of the DAC is $0-2V_{ref}$. Please refer to the datasheet in Appendix H-9. In order to have a full 0-5 V output range, the DAC was provided with a 2.5 V external reference.

Like the ADC, the DAC communicates with the Jackrabbit microcontroller via SPI protocol, in this case over serial port D. Table 3-5 gives the digital signal connections between the two devices.

RCM3100 Pin Name	Pin Function	TLV5630 Pin Name
PB2	IA0	/CS
PF0	CLKD	SCLK
PC0	TXD	Din

Table 3-4 – Digital signal connections between the Jackrabbit RCM3100 microcontroller and the TLV5630 DAC.

The operation of this DAC is similar to the 4-ch TLV5614, Which is used to provide analog commands to the four motor controllers. As a result, the code for sending commands to the DAC using SPI on serial port D had already been written [Dietrich and Zettl 2005]. All that was needed to adapt that code for use with the TLV5630 was to establish a library that had

the appropriate bit-level commands and addresses implemented in it. Details on those bit-level commands for addressing and controlling the TLV5630 are thoroughly explained in the datasheet given in Appendix H-9. The bit-level commands and several higher-level API functions are contained in the library `TLV5630.lib` (Appendix I-2)

The DAC is available in an 20-SOIC (wide) package. To connect this surface mount chip to the rest of the circuitry, it was soldered to a 20-SOIC surfboard. In the case of the TLV5614 that provides the motor commands, the 0-5 V output is post-processed so that it can provide a ±5 V signal to the motors, so that they can drive either forwards or backwards.[21] As the output range of the DAC matches the input range of the DC/DC converters, no such post-processing is needed. However, in order to smooth the sharp jumps in the DAC output when new inputs are received, each channel passes through an anti-imaging filter with a bandwidth of approximately 1 kHz. The DAC circuit schematic is given in Appendix G-7. Table 3-6 gives a breakdown of the DAC channel assignments for the power system. As with the ADC, it is recommended that future generations of Cool Robot implement the DAC circuits on a printed circuit board to reduce the size of the circuit, and the labor required to produce it.

DAC Channel #	Corresponding Signal
0-4	DC/DC converters 1-5
5-7	(currently unused)

Table 3-6 – Channel assignments for the TLV5630.

3.5 SUMMARY

This concludes the description of the essential hardware of the Cool Robots power system. In summary, it contains a photovoltaic input source consisting of five solar panels,

---

[21] Due to the arrangement of the motors in the Cool Robot, at any given time two motors are driving forwards while the other two drive backwards.

input power conditioning in the form of five independently-controlled DC/DC boost converters, energy storage in the form of three large Li-ion batteries, a common bus connecting the solar power input to the batteries, with two major breakpoints implemented with MOSFETs, a housekeeping power supply that provides regulated voltage supplies for the rest of the electronics, solid mechanical connections between all the major power components and the load, a host of measurement and conditioning circuits to detect the status of all the parts of the system, and data converters to interface the system to the power microcontroller. Now that these parts have been described, we will now turn our attention to the programming of the power microcontroller, which coordinates and controls the functions of the power system.

4. Control of the Cool Robot Power System

4.1 Review of Photovoltaic Control Techniques

As described earlier in Chapter 2, there are several ways that one can control the operation of a solar panel along its I-V curve. The most commonly used technique, as it allows one to easily vary the operating point, is to connect a solar panel to some power electronics device, such as a DC/DC converter, whose large signal input impedance can be adjusted through modulating the converter's duty cycle.

Let us examine once again the boost converter in Figure 3-10. Let us say that the load is some fixed resistance R. Let us also say that, either by design or by simple measurement, we know the output voltage V_{out}, and that V_{out} is a constant. The output current of the device is then:

$$I_{out} = V_{out}/R \tag{4.1}$$

By the characteristic of a boost converter topology (equation 3.3, reprinted here for convenience):

$$V_{out} = \frac{V_{in}}{1-D} \tag{3.3}$$

where D is the duty cycle. If the efficiency of the converter is ideal, then the input and output power are identical. That is:

$$V_{in}I_{in} = V_{out}I_{out} = \frac{V_{out}^2}{R} = \frac{1}{R}\left(\frac{V_{in}}{1-D}\right)^2 \tag{4.2}$$

so that

$$\frac{V_{in}}{I_{in}} = R(1-D)^2 \tag{4.3}$$

and so the input impedance of the converter, dV_{in}/dI_{in}, is $R(1-D)^2$.

The input impedance of a DC/DC converter appears to the solar panel to simply be a load that varies with the duty cycle (so long as the converter output is indeed a resistance R). Referring back to Figure 2-24 and 2-25, one can see that, with a fixed load and by varying the duty cycle D, one can traverse the whole span of the I-V curve. This is just an example with one particular converter topology, but a similar analysis can be performed for any DC/DC converter circuit that might be used in conjunction with a solar panel.

The primary reason that one would want to have some control over where the solar panel operates along its I-V curve is so that one can extract the maximum amount of power from the panel, and track the location of the maximum power point as the shape of the I-V curve shifts due to changing insolation and/or temperature conditions. Being able to track the maximum power point in real-time requires that the I-V curve be constantly monitored and the duty cycle continually updated to reflect any changes. There are a number of methods for performing maximum power point tracking, and the following discussion will introduce the most often-used ones.

VOLTAGE FEEDBACK METHODS

There are two main ways that one can use the panel voltage in a closed-loop feedback path to produce the optimal duty cycle command for a maximum power point tracker, both of which require *a priori* knowledge of the electrical characteristics of the panel being controlled and also of the power electronics being used. The first method simply measures the panel's open circuit voltage, V_{OC}, at regular time intervals and then assumes that the value of the maximum power point's voltage, V_{mp}, is some fraction of V_{OC}. That value of V_{mp} is the desired panel voltage, which is inserted as the command into a closed-loop

feedback controller. The controller is typically a PI or PID controller, designed using classical control theory techniques, which will exhibit zero steady state error.

While the controller itself is operated in a closed-loop fashion, the determination of the desired panel voltage is done in a predictive, open-loop fashion, and there is no guarantee that one will actually be operating at the maximum power point. Knowing what fraction of V_{OC} to use to find V_{mp} requires a considerable knowledge of the panel's electrical characteristics over a broad range of operating conditions. A typical proportionality constant to use is 0.75 [Enslin 1997]. However, over a broad range of operating conditions, this relationship between V_{OC} and V_{mp} could vary from as little as 0.6 to as much as 0.8. As a result, this method will rarely place the panel at its maximum power point, but merely be in the general vicinity. While most all MPPT methods are able to only operate in the vicinity of the maximum power point (many oscillate about it), this method has no mechanism to bring itself any closer to the maximum power point beyond its initial guess. The primary way to determine V_{OC} is to disconnect the panel entirely from the circuit, which means that, during those times of measurement, the panel can provide no power to the system. An alternate method for determining V_{OC} is to use a reference cell, disconnected from the rest of the array, and always measure its V_{OC}. In the case of the Cool Robot, due to the fact that each panel operates under a different set of operating conditions, one of the 54 cells from each panel would need to be used as a reference. Using this cell as a reference, rather than to provide power, would represent a 2% loss. To mitigate the loss, a fragment of a full-sized cell could perhaps be used as an alternate reference. In any event, this panel voltage method is rarely used except in the most crude systems.

The second voltage-feedback method can allow slightly more accuracy in the determination of V_{mp} over a range of operating conditions, and does not require that the panel be disconnected from the system to find V_{OC}. In this method, a number of sensors measure such quantities as the amount of available sunlight and the ambient temperature. These sensor readings are then fed into a model of the solar panel, which calculates what the location of the maximum power point should be. The model can be in the form of a table in memory, a formula, as the model of the A-300 solar cell was; or something more esoteric, such as a neural network. From this prediction of V_{mp}, one controls the panel's voltage in the same fashion as the previous method. However, as I have shown in section 2.6, developing a numerical model of a solar panel that is accurately able to predict the I-V curve over a wide range of operating conditions (or environments) is quite difficult. In order to execute this method, one needs the sensors, a means to measure them, and some a way of computing the model. For the same number of sensors and computing power, one could utilize more effective methods.

PERTURB AND OBSERVE METHOD

In order to accurately track the location of the maximum power point of a solar panel as it moves through a range of operating conditions, one needs more information than simply the panel voltage. A more typically used control technique for maximum power point tracking seeks to make the derivative of panel power with respect to some controllable variable equal to zero. That derivative may be with respect to the panel's current ($\partial P/\partial I$), panel voltage ($\partial P/\partial V$), duty cycle of the DC/DC converter the panel is connected to ($\partial P/\partial D$), or some other variable that has a strong influence on the panel's power. Sometimes, one does not directly find the derivative of power with respect to I, V, or D, but

rather the derivative of power with respect to time, $\partial P/\partial t$, and then compare that to the time-derivative of *I*, *V*, and *D* by the chain rule.

One problem with trying to force the derivative of power to zero is that, in order to measure $\partial P/\partial x$, where *x* is the variable of choice for the power point tracker, some measurable change in the system is necessary. If the operating conditions of the solar panel, DC/DC converter, and the load are constant in time, there will not be any change to measure. Without a ∂P or ∂x, it is not possible to determine the derivative of power. Therefore, the methods that find $\partial P/\partial x = 0$ to perform maximum power point tracking require a perturbation ∂x to produce a measurable change in power, from which $\partial P/\partial x$ can be determined. The perturbation can be explicit, such as adding a small-amplitude sine wave to the control command, or implicitly coming from the natural operation of the system. These methods are collective called Perturb and Observe (PAO) methods. While PAO methods tend to be straightforward to implement, either as a discrete circuit or in a microcontroller-based system, they have the disadvantage of always requiring some perturbation, and thus operate *around* the MPP, and never stay exactly on it.

In [Sullivan and Powers 1993], the authors present a design for a maximum power point tracker for use in a solar racing vehicle. In that vehicle, as in the Cool Robot, the output of the power electronics circuitry is a connected to a bus, to which the batteries are directly connected. As a result of the batteries being present, the output voltage of the bus is clamped to whatever the battery voltage is, as it is with the Cool Robot. Therefore, because the output voltage is more or less fixed in time, one need measure only the output current, rather than output power. The output current of the MPPT is monitored because the real control goal is to maximize the power delivered to the system, which is the power output of

the MPPT, which may not necessarily correspond exactly to the maximum power point of the solar panel.[1] The control variable in that system is the duty cycle, so the overall control goal is to force $\partial I_{out}/\partial D = 0$. The perturbation on D is produced through what they refer to as a "clocked auto-oscillation method." The control circuitry produces a ramp signal to the FET driver circuitry, which translates the voltage of the ramp into a duty cycle of the PWM signal. In that case, $dD/dt = \pm \mathrm{const}$, which is the perturbation. The direction of the ramp, either positive or negative, is determined by sign of $\partial I_{out}/\partial D$. In their system, when $\partial I_{out}/\partial D > 0$, the desired control action is to increase D so as to drive the system towards the maximum power point. The reverse is true when $\partial I_{out}/\partial D < 0$. In this way, the controller steadily climbs up the power curve, either from the left or the right, until it reaches and passes the maximum power point and the sign of $\partial I_{out}/\partial D$ changes, forcing the controller to switch direction and head back to the MPP. This sort of control, where one heads towards a goal until one passes it, then turns around and heads back towards it, is called "Bang-Bang" control.[2] Another example of Bang-Bang control is a thermostat trying to maintain a certain temperature: the thermostat clicks on and heats the room until the target temperature has been passed, then clicks off and allows the room to cool.

In [Bhide and Bhat 1993], the maximum power point tracking algorithm is implemented using a microcontroller. This PV system, already mentioned in chapter 3, is a modular system where each panel has its own conditioning power electronics, and the microcontroller controls each in turn, similar to the arrangement in the Cool Robot. The algorithm seeks to minimize $\partial P/\partial V$ by varying the duty cycle of the power electronics

---

[1] In truth, the maximum power point of the solar panel and of the solar panel/MPPT combination are close enough such that either power may be used.

[2] The "clocked" label comes from the fact that the "Bangs" can only occur in time to a fixed clock frequency.

module, and so takes measurements of the current and voltage of each panel it is connected to. Because this is a discrete-time system, the derivative $\partial P/\partial V$ cannot be directly measured, but must be estimated as $\Delta P/\Delta V$, where Δ is a difference operator between two time steps. The program runs on a fixed sampling rate of 10 Hz, and has an adaptive step size. The formula used to evaluate and update the step size is not given in this paper, but often is something in proportion to $\Delta P/\Delta V$. In this way, when the algorithm is far from the maximum power point, a large step size in the duty cycle is used. As the algorithm converges on the maximum power point, the step size decreases. The step size is not allowed to go all the way to zero, or else there will be no perturbation from one cycle to the next. One further feature of this microcontroller-based system, which will be reproduced in the Cool Robot, is the use of a data structure in the software to describe the state of each panel and its electronics. Contained within the structure is such information as the current duty cycle, the current step size, the addresses of the ADC and DAC channels associated with that panel, etc.

OTHER METHODS

In [Pacheco 2003], the authors have designed a system that combines not only maximum power point tracking, but the regulation of the output voltage. This two degree of freedom system requires two actuators (i.e., two independently-controlled switches), and some form of energy storage (in this case, a battery). To perform maximum power point tracking, the signs of dP/dt and dV/dt are compared, and the state of the MPPT switch is determined by an exclusive-OR (XOR) operation. This is a form of PAO that does not explicitly seek to minimize the derivative of power. To understand the operation of this system, please refer back to Figure 2-6, which shows the solar panel's power output as a

function of its terminal voltage. Let us suppose for a moment that one is currently operating to the left of the MPP. As one increases the voltage ($dV/dt > 0$), there will also be an increase in power ($dP/dt > 0$). Inversely, as one lowers the voltage, one also lowers the power (both dV/dt and $dP/dt < 0$) If one defines a positive derivative as logic 1 and a negative derivative as logic 0, then the result of an XOR, with the panel operating to the left of the MPP, is logic 0. To the right of the MPP, the situation is different. Now, as one raises the panel voltage ($dV/dt > 0$), the corresponding power will decrease ($dP/dt < 0$). In that case, with the panel operating to the right of the MPP, the signs are opposed, and the result of an XOR will be 1. The output of the XOR gate drives the MPPT switch. In their circuit topology, the panel voltage will increase when the switch is open (logic 0), and decrease when the switch is closed (logic 1), which are the control actions desired to reach the maximum power point. This is another example of Bang-Bang control. The other functionality of their device, regulating the output voltage, is accomplished through a simple PI controller operating a second switch.

There are a host of other, more exotic techniques that can be used for performing maximum power point tracking. The method presented in [Tse 2002] places a perturbation to the frequency of the pulse width modulator, rather than the duty cycle, and evaluates the AC-portion of the resulting panel voltage. In [Midya 1996], the instantaneous ripple in the panel voltage, current, and power created during the switching of the DC/DC converter itself, rather than the coarser average changes produced by perturbing D, are used as a sort of intrinsic perturbation. This is referred to as dynamic MPPT. Some other methods are borrowed from the body of knowledge in nonlinear control theory. Some methods, like the use of trained neural nets, are more academic exercises than practical applications.

4.2 CONTROL CHALLENGES IN THE COOL ROBOT POWER SYSTEM

As should already be apparent, the control of a solar power system rarely lends itself to the tools of control theory. Classical control theory utilizes sensors, feedback loops, root locus methods, Bode plots, and basic analog circuits to bring a system "under control," that is, to make it behave in the way you want. Classical control theory typically is applied to a wide range of systems (mechanical, electrical, thermal, etc.). The control of DC/DC converters, to enable them to produce a steady output voltage, utilizes many of these tools. Classical control theory is generally limited to single-input, single-output (SISO) systems. modern control theory, which uses multiple inputs and multiple outputs (MIMO), is often computer based, and operates using matrix mathematics. Modern control theory is used to control complicated systems like high-performance aircraft.

Most of the well-developed and effective tools of control theory, however, cannot easily be utilized in the Cool Robot power system. The primary reason is that photovoltaic power systems are highly nonlinear. Control theory tends to avail itself to systems that are referred to as linear, time-invariant (LTI).[3] In order to design a compensator using classical control theory, one needs to have a linear input/output model of the system's behavior (or linearized about some operating point), and assume that the model is valid for all time (or at least, very slowly changing). One brief look at the I-V or P-V characteristic of a solar panel (Figures 2-5 and 2-6) will indicate that they are most certainly not linear. As I described in Chapter 2, the electrical characteristics of the solar panels on the Cool Robot are also time-varying, due to the motion of the robot over obstacles and the motion of the sun in the sky. The control of nonlinear, time-variant systems is still a developing field, and the available

---

[3] The author will assume that the reader is familiar with this term and its implications.

control methods are not as general-purpose or have the predictable performance characterized by control methods for LTI systems. The maximum power point tracking schemes described in the previous section are all nonlinear: Bang-Bang control is a method from control theory that works for both linear and nonlinear systems.
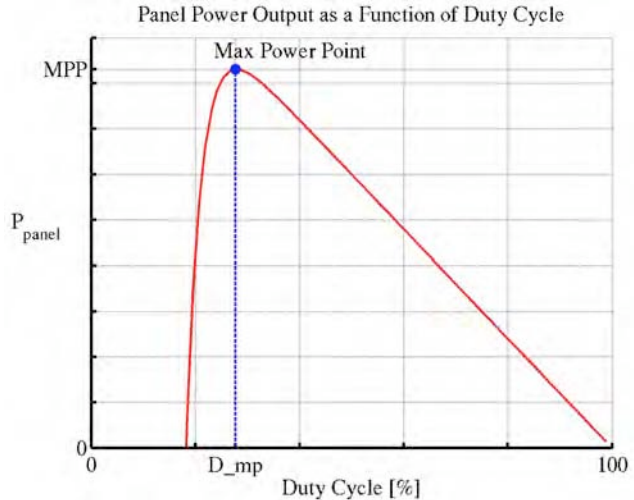
Figure 4-1 – The relationship between panel power and duty cycle, for the case where the panel is connected to a boost converter with a fixed output voltage.

In order to understand the control algorithm that will be used in the Cool Robot power system, it is necessary to have an understanding of the input-output relationship of the system. We begin with the P-V characteristic of the solar panel, which is given in Figure 2-6. This is the input to the DC/DC boost converter, whose input-output characteristic is given by equation 3-3:

$$V_{out} = \frac{V_{in}}{1 - D} \tag{3.3}$$

However, as we noted in section 3.2, V_{out} is a nearly fixed quantity, equal to the battery voltage. Therefore, we rearranged equation 3.3 to produce equation 3.4:

$$V_{panel} = (1 - D)V_{batt} \tag{4.3}$$

Equation 4.3 tells us what the panel voltage will be as a function of the duty cycle input, given that the output voltage is fixed by the batteries. This relation is shown graphically in Figure 4-1. Note that the graph does not extend all the way to a duty cycle of 0%; the graph has an x-intercept. The x-intercept is the duty cycle that is necessary to raise the panel's

140

opern circuit voltage up to the level of the battery bus. For duty cycles between 0% and that x-intercept, the boost converter is unable to raise the panel's voltage up to the level of the bus. In other words, according to equation 3.3, $V_{panel}/(1-D) < V_{batt}$. In this regime, the DC/DC converter operates in what is called
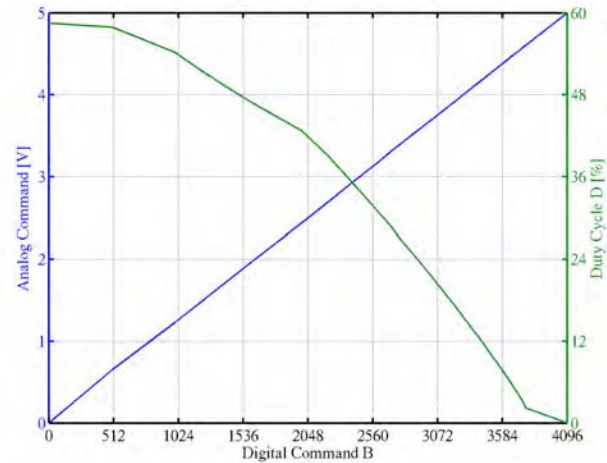
Figure 4-2 – The relationship among the digital command B, the resulting DAC output, and the duty cycle D of the boost converter.

discontinuous conduction mode, which behaves differently than equation 4.3 and will be avoided entirely, as the transition between continuous conduction to discontinuous conduction could confuse the control algorithm.

However, we do not explicitly command the duty cycle *D*. Recall that the power microcontroller sends some 12-bit binary data, *B*, to the TLV5630 digital to analog converter, which produces an analog command in the range of 0-5 V. That analog command is utilized by the DC/DC boost converter's timing circuitry to modulate the duty cycle of the PWM command. So, we are interested in the relationship among the 12-bit number *B*, the analog signal produced from it by the DAC, and the actual duty cycle *D* produced by the DC/DC converter. This relationship, found experimentally, is given in Figure 4-2. The relationship between the digital number, *B*, and resulting duty cycle *D* is closely given by equation 4.4:

$$D = 58.6\% \cdot \cos\left(\frac{\pi}{2} \cdot \frac{B}{3925}\right)$$

So we can further transform the relationship shown in Figure 4-1 to produce panel power as a function of *B*. This relationship is shown in Figure 4-3. The precise shape and location of

this curve will change as the I-V characteristic of the solar panel changes. The gently curving area to the left of the MPP represents the constant current portion of the I-V curve, while the steeper portion to the right of the MPP is the constant voltage portion. As the temperature decreases, the maximum power point
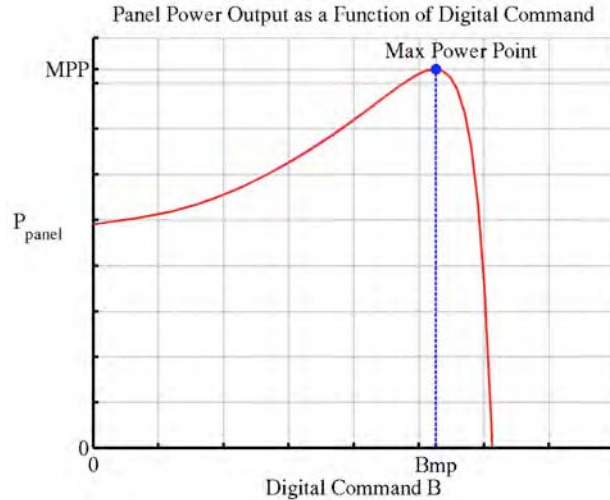
Figure 4-3 – The general relationship between digital command B and panel power output.

will shift up and towards the right. As the insolation increases, the maximum power point will shift up. Using the relationship shown in Figure 4-3, we could design a maximum power point tracking algorithm that increases B when $\partial P/\partial B$ is positive, and decrease B when $\partial P/\partial B$ is negative. This would utilize the Bang-Bang control scheme discussed earlier.

However, the shifts in the curve due to changes in temperature and insolation make this task more difficult. Were this a static curve (i.e., unchanging in time), then finding and tracking the maximum power point would be relatively easy. If the curve was slowly changing, as happens with a solar panel affixed to someone's rooftop, or with the temperature changes in the Cool Robot's solar panels, tracking the maximum power point would still be relatively straightforward. However, during those times when the robot is driving across rugged, uneven terrain and tilting in various ways, the curve in Figure 4-3 could be changing quite rapidly. The tilting of the robot changes the orientation of the five solar panels with respect to the sun, which substantially changes the amount of insolation they each receive. If we assume that the *sastrugi* occur with a spatial period of 1-2 meters,

and the robot is driving over them at its maximum speed of 1 m/s, then the frequency of the robot's motion will be 0.5 Hz to 1 Hz. The speed of this change in the panel's electrical characteristics is much faster than the changes resulting from the sun's motion in the sky or the change in the panel's temperature. This is a unique control challenge for the Cool Robot, compared to stationary PV installations, and one that the power system will be faced with nearly all the time as the robot drives across the uneven snowpack. Depending on the speed of the algorithm, such rapid changes would cause, at the very least, poor tracking of the maximum power point and, at worst, unstable performance.

To illustrate this point, consider the sequence of events in Figure 4-4. The upper curve represents the P-B relationship at time t_k, and the lower curve is at time t_{k+1}. The curves are offset from one another as a result of a shift in the insolation that the solar panel
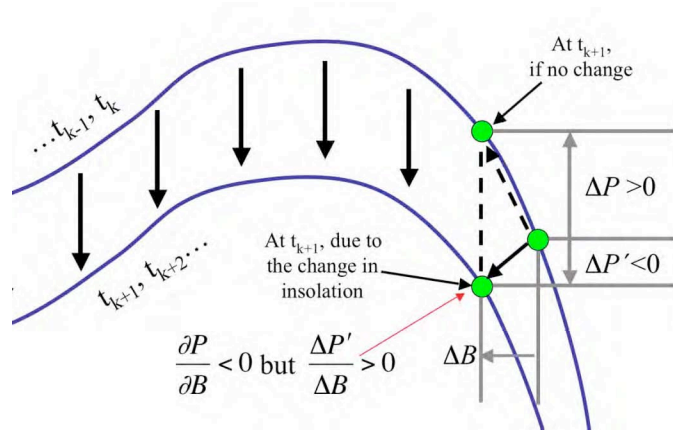
Figure 4-4 – Demonstration of how a shift in the P-B curve can result in an erroneous determination of the slope ∂P/∂B.

receives. In this case, there is a significant decrease in insolation when going from time step t_k to t_{k+1}.[4] Let us suppose that, at time t_k, the slope of $\partial P/\partial B$ has accurately been determined, and a change in duty cycle command, ΔB_{k+1} ($\Delta B_{k+1} < 0$), has been calculated to bring the system closer to the maximum power point. Were there no change in insolation

---

[4] While all of the following examples are demonstrated using rapid decreases in insolation, the same is of course true for rapid increases in insolation. Changes in the curve due to temperature or the movement of the sun about the horizon do not happen nearly as fast, and do not have as dramatic an effect on the shape of the P-B curve, and we will assume that, if the control algorithm can adjust to the rapid changes due to insolation, it will likewise be able to handle the slower changes.

(and hence, in the shape of the curve) between time steps, this ΔB_{k+1} would result in an increase in power ΔP_{k+1} ($\Delta P_{k+1} > 0$). However, due to the change in insolation, the power actually *decreased* by $\Delta P'_{k+1}$ ($\Delta P'_{k+1} < 0$). Therefore, at time t_{k+1}, when the MPPT algorithm tries to estimate $\partial P / \partial B$ by calculating $\Delta P'_{k+1} / \Delta B_{k+1}$, it will erroneously determine that the slope is positive. The control law states that the next change in duty cycle command, ΔB_{k+2}, will be positive. However, $\Delta B_{k+2} > 0$ will take the system further from the maximum power point.

In truth, even though the P-B curve has translated down from time step t_k to t_{k+1}, the commanded ΔB_{k+1} correctly did move the system towards the new maximum power point. However, the decrease in the curve dominated over any increase ΔP_{k+1} that came about from moving closer to the maximum power point, resulting in the (negative) $\Delta P'_{k+1}$. However, the control algorithm has no way of distinguishing between changes in power brought about by the control algorithm, and those brought about by sudden changes in the panel characteristics. This limitation comes about from the implicit assumption that the shape of the power curve is a function of the command B only, while in fact it is a function of several other variables, which are unobservable to the system.

If the drop in insolation continues beyond time step t_{k+1}, the resulting behavior of the MPPT algorithm is shown in Figure 4-5. As one can see, the control algorithm will alternately make correct and incorrect jumps until the curve stops changing so rapidly. If the curve were to shift up suddenly, similar errors in direction would occur every other time step. If the up and down shifts in the curve occurred with a frequency on the order of the control algorithm's sample rate, it is possible that the algorithm could be repeatedly bamboozled and

144
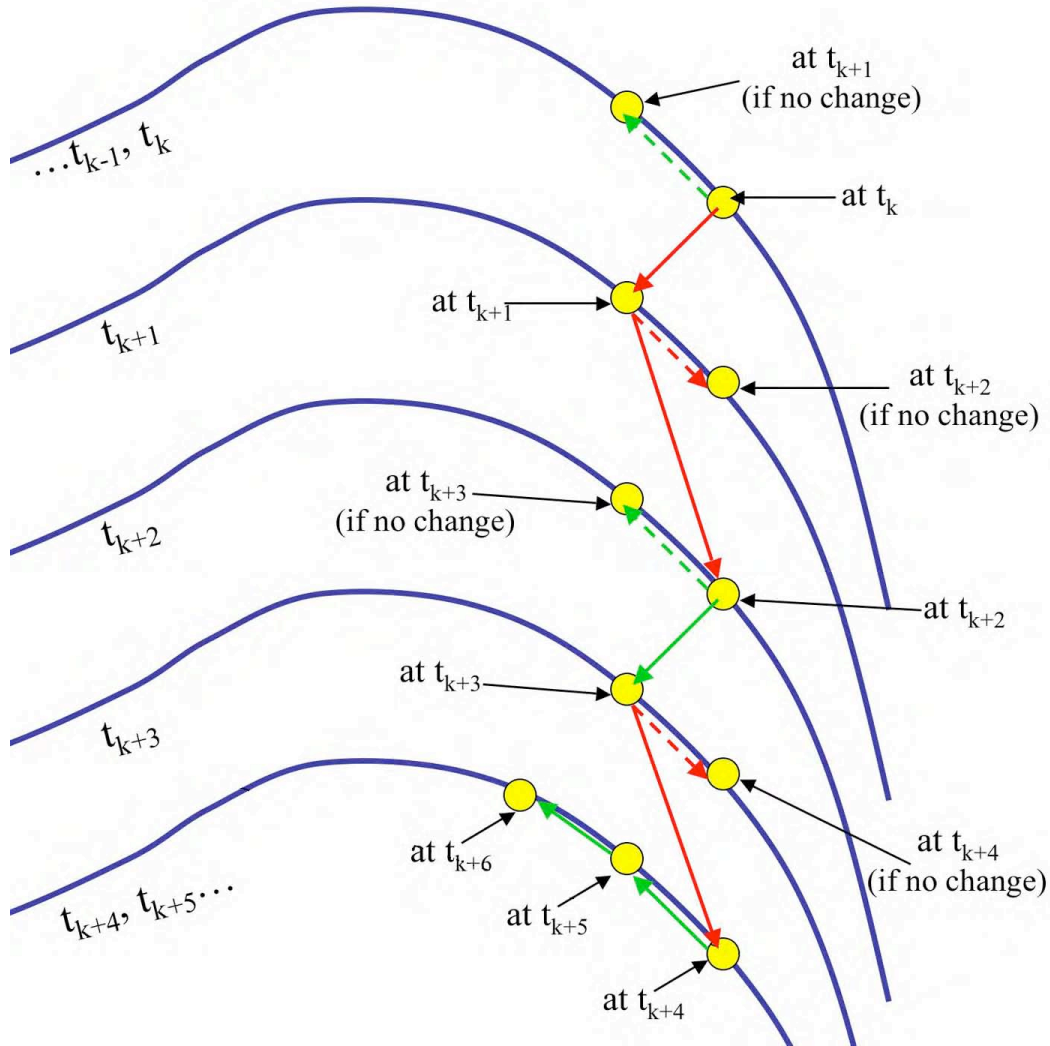
Figure 4-5 – Demonstration of how a continuous shift in the P-B curve can result in alternating correct and incorrect decisions by the MPPT algorithm.

shift completely away from the maximum power point, all the way to the limits on B.[5] Such a situation is extremely unlikely in practice, but does illustrate how the tracking performance of the maximum power point tracking algorithm cannot be guaranteed.

However, if the shift in the curve relative to the speed of the controller is small, then the system can still operate correctly. This is shown in Figure 4-6. In this case, the increase in power caused by the correct duty cycle change ΔB_{k+1} offsets the decrease in power that results from the a shift in the power curve. The resulting calculated slope, $\Delta P'_{k+1}/\Delta B_{k+1}$, will

---

[5] Since B represents the 12-bit data sent to the DAC, the limits are 0 and 4095.

be smaller than it otherwise would
have been, but it will have the correct
sign, so that the controller will move
on correctly in ΔB_{k+2}. In order for
the shifts in the curve to appear small,
the speed of the controller must be
fast.
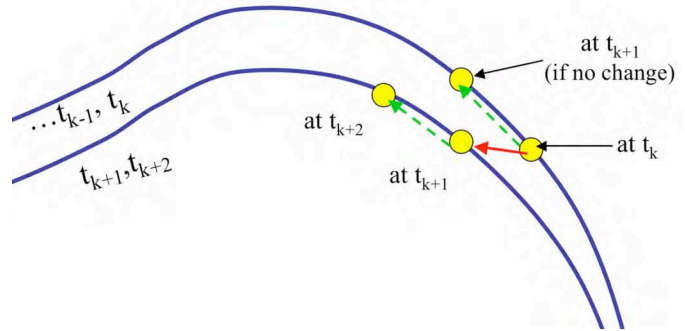
Figure 4-6 – Demonstration of how the MPPT algorithm can still operate correctly if its sample rate is fast compared to the rate of change in the P-B curve.

One may notice that, in each of these examples, I have assumed that the solar panel is operating to the right of the MPP (on both the P-B and P-V curves), where the panel voltage is semiconstant, and the current changes greatly. While theoretically it should not matter which side of the maximum power point we operate on (power is power, after all), the decision to operate to the right of the MPP was intentional. The first reason is that the boost converter topology operates more cleanly when the voltage ratio between input and output is near to unity. Operating the panel and boost converter combination to the left of the maximum power point would mean that the panel voltage would be considerably lower, requiring longer duty cycles to raise that voltage to the level of the bus, and thus reducing converter performance. Related to this is the fact that the boost converter, due to the design of its PWM circuitry, is only capable of producing duty cycles in the range of approximately 2-60%. In this range, only the "constant voltage" side of the power curve is fully accessible (from V_{mp} to V_{OC}). Lastly, as one can see in Figures 4-4 through 4-6, the susceptibility of the MPPT algorithm to changes in insolation is less when the slope of the *P-B* curve is steep. The constant current region of that curve exhibits shallow slope, which means that the magnitude of $\Delta P / \Delta B$ is less there than on the constant voltage side. Therefore, for the same

change in B produced by the control algorithm, the resulting change in P will be less on the constant current side than the constant voltage side. If ΔP is (in part) the "signal" that we measure and base our control decisions on, and the shifts in the power curve are thought of as "noise" (fluctuations in the measurement that we cannot predict or separate from the signal), then the constant current side of the power curve exhibits poorer signal to noise ratio.

As the tracking algorithm oscillates about the maximum power point, and as the power point moves due to changes in the power curve, the operating point will inevitably spend some time to the left of the maximum power point. However, the algorithm should be able to detect when it is operating in that regime, as the estimated slope $\Delta P/\Delta B$ will become positive. In that case, according to the control law, the algorithm will begin to increase B so as to traverse back over to the right side of maximum power point.

4.3 MPPT Algorithm Simulation

In order to investigate more fully the difficulties in performing maximum power point tracking on a time-changing panel characteristic, a MATLAB simulation was created. This simulation would calculate the I-V characteristic of a solar panel using the simple model given in equation A.5 and Figure 2-5. The model approximates the solar panel as a current source in parallel with a diode that obeys the characteristic given in equation 2.1. In the simulation, the current can either be constant, representing a time-invariant panel, or varied in time in a sinusoidal fashion. In the time-invariant case, the current is 5 A. In the time-variant case, the current varies sinusoidally between 3 A and 5 A, with a frequency f_{motion} of 0.5 Hz. The fact that the time-variant case does not extend all the way to 0 A comes about from the fact that, even though the panel's orientation with respect to the sun is changing, it will always receive some sunlight. It should be noted that the significant amount of reflected

147

insolation from the snow provides a large portion of this DC component, and probably helps to mitigate the severity of the changes in insolation from the robot's motion.

From this *I-V-t* data, the transformations given earlier in this chapter are used to create a *P-B-t* dataset. The resulting 3-D surface, showing how the P-B relationship varies in time in response to a time-changing insolation, is shown in Figure 4-7. A contour plot of the same surface, showing how the location of the maximum power point moves, is given in Figure 4-8. The code that generates *P* as a function of *B* and *t* is given in Appendix J-1.
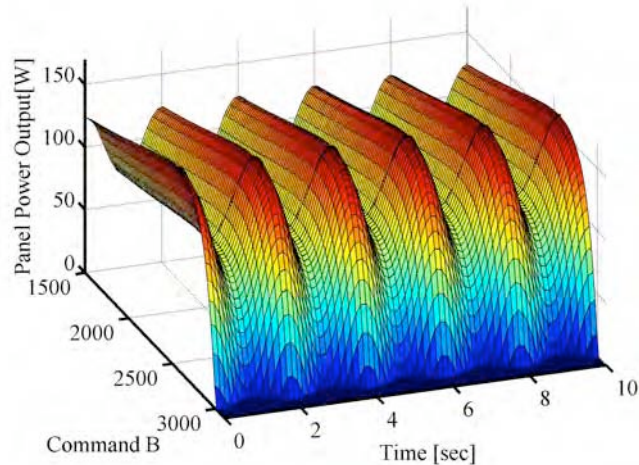
Figure 4-7 – Simulation of how the P-B characteristic changes in time due to a 0.5 Hz sinusoidal variation in the insolation the panel receives.
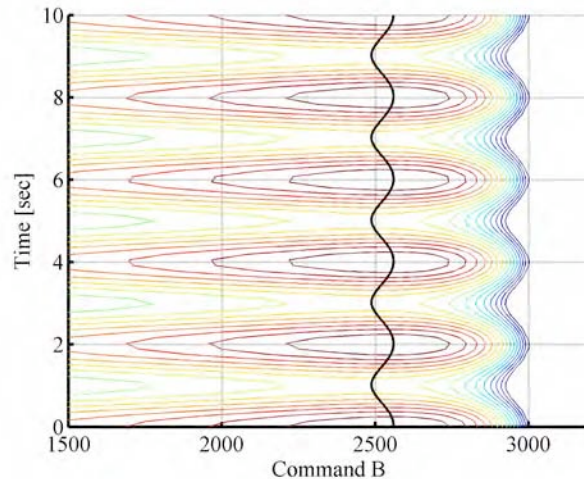
Figure 4-8 - Contour plot of the surface shown in Figure 4-7. The location of the MPP is the black track. Notice that the location of the MPP changes in both P and B.

With the ability to determine approximately what the power output of this simulated panel would be for a given time *t* and duty cycle command *B*, the MPPT algorithm could be tested and tuned. The MATLAB simulation code is provided in Appendix J-2. The code implements the gradient ascent control approach introduced earlier in this chapter. Like in [Bhide and Bhat 1992], the algorithm uses an adaptive step size, μ_k, as it was quickly

determined to be much faster in finding and tracking the power point than a fixed step size. The step size is proportional to the calculated slope $\Delta P/\Delta B$. The sampling rate of the MPPT algorithm was varied from as slow as 10 Hz to as fast as 200. The adaptive step size's constant was adjusted in turn for optimal performance at each frequency. At each time step, the simulated panel power was calculated for the current operating point, from which the slope, step size, and next duty cycle were calculated. In addition to the power for the single operational point, the power for the whole range of B at each time step t_k was calculated, so that the instantaneous value of the maximum power point would be known. At the end of the simulation, these data were used to determine the effectiveness of the maximum power point tracking algorithm with a figure of merit defined by $P_{achieved}(B_k,t_k)/P_{mpp}(t_k)$, the power achieved by the power point tracking algorithm divided by the actual maximum power possible. The average over all t is a good measure of the effectiveness and adaptability of the maximum power point tracking algorithm. In general, the adaptive step size algorithm, operating at sampling frequencies between 20 and 100 Hz, achieved an average tracking effectiveness of over 95%. That is, over 95% of the available power was captured on average by the tacking algorithm. In general, I found that that tracking effectiveness did not appreciably improve above, 40 or 50 Hz

For qualitative visualization and comparison purposes, the simulation code would also overlay the track of the algorithm on the calculated power surface. In one case, the track is plotted in 2-D on a contour map of the power surface. This is shown in Figure 4-9. The other visualization plots the track in three dimensions on the power surface of Figure 4-7. This second visualization is given in Figure 4-10.
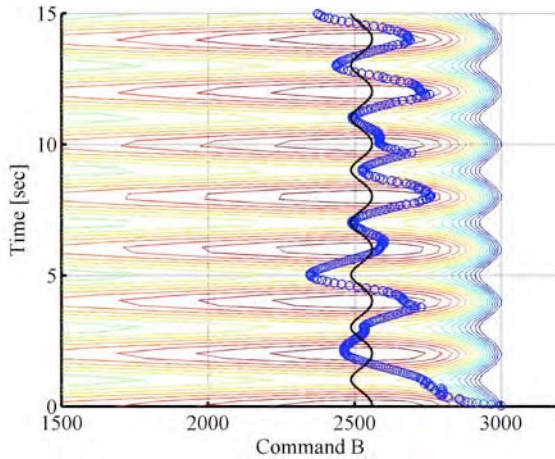
Figure 4-9 – The track of the maximum power point tracking algorithm overlaid on a contour plot of the power curve. The track of the true MPP is the squiggly black track.
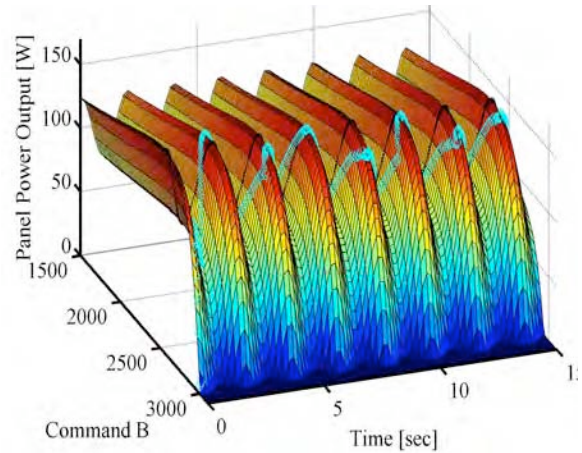
Figure 4-10 – The track of the maximum power point tracking algorithm overlaid on the power surface. The tracking effectiveness for this simulation run was approximately 96%.

Having proven the effectiveness of the maximum power point tracking algorithm in a time-changing situation, it was also applied to the static case. That is, to the much simpler case where the insolation and temperature of the panel are held constant in time. This is the situation the robot would be in when it is stationary or driving on a smooth surface. Because the algorithm did not need to adapt to changes in the curve, its tracking effectiveness was much higher, over 99.5%. In this static case, the algorithm was able to climb up to 95% of the global maximum within about 10 iterations, and settled into a steady-state oscillation about the maximum power point after about 50 iterations. The track of this simulation is shown in Figure 4-11.
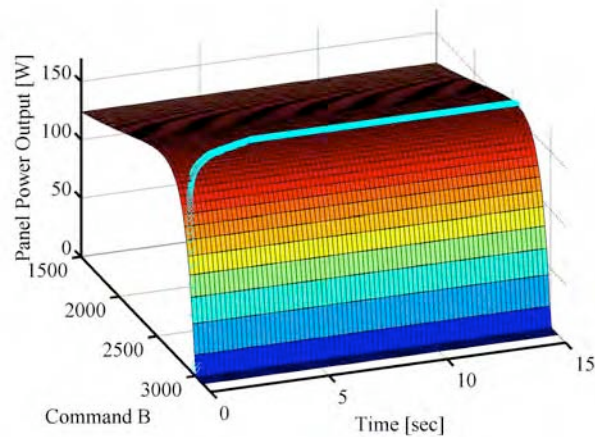
Figure 4-11 - Simulated track of the maximum power point tracking algorithm when presented with a time-invariant power curve.

These two cases, however, assume that the algorithm has perfect knowledge of the power that results from the particular duty cycle command *B*. In reality, however, the

150

algorithm will not have such perfect measurements. Due to the limitations of the analog-to-digital converter and the surrounding circuits, the measurements of the panels' voltage, current, and power will have a certain amount of random noise embedded in them. Therefore, to test the robustness of the maximum power point algorithm in the presence of measurement noise, the simulation was extended. While the power delivered by the solar panel for a given command B was still exactly determined at each time step, Gaussian noise was added to the power measurements before the MPPT algorithm used them. The standard deviation of the noise, which is measure of how large in amplitude it is, became a simulation parameter. A standard deviation of zero would produce no noise on the measurements, resulting in the "perfect knowledge" cases described above. One would ideally measure and record the amount of noise present in the physical system, and then use its standard deviation in the simulation. However, it was not possible to obtain this measure accurately. Instead, standard deviations of 0, 1, 5, and 10 watts were used, with 1 W being a conservative value and 10 W representing the absolute worst case.

While the algorithm did not perform as well in the presence of measurement noise as in the noise-less case, the results of the simulations were nevertheless encouraging. For the case where the power curve varied in time, the tracking effectiveness of the algorithm (averaged over a single run) was generally in the range of 92-94%. In all but one or two of several dozen runs, the tracking effectiveness for each simulation was above 90%. Due to the stochastic nature of the noise in the simulation, the performance varied from one simulation run to the next. The track of one such run is given in Figures 4-12 and 4-13.

The simulation in MATLAB proved essential for developing and debugging the maximum power point tracking algorithm that became an important element for the overall
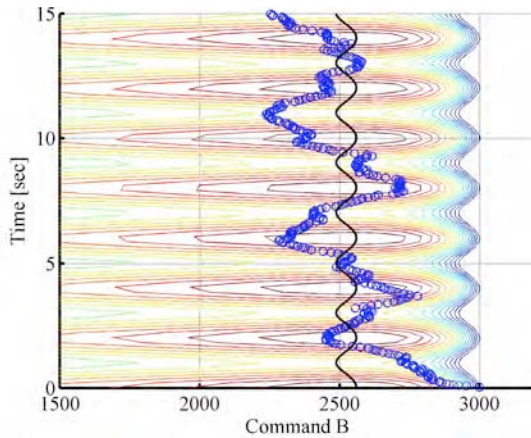
Figure 4-12 - Contour plot of the track of the maximum power point tracking algorithm with measurement noise present. In this run, the standard deviation of noise was 1 W, yet the tracking effectiveness was still 93%.
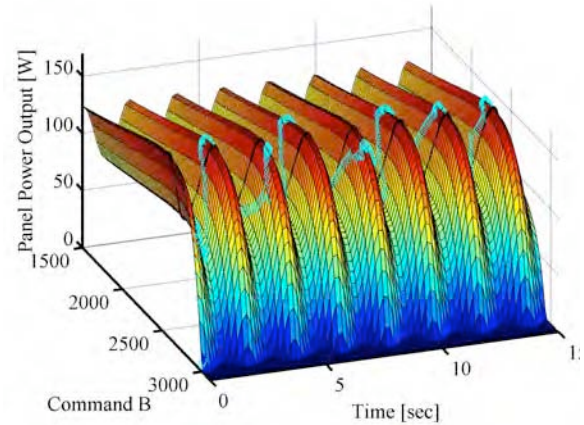
Figure 4-13 - Surface plot of the maximum power point tracking algorithm in the presence of noise for the same case.

power system control algorithm presented in the next section. The resulting maximum power point tracking algorithm is shown as a flowchart in Figure 4-14. These simulation results give us confidence that the maximum power point tracking algorithm will be robust and adaptive enough to handle the cyclic variations in each of the five panels. Additional simulations would be prudent to further characterize the performance and robustness of the MPPT algorithm. For instance, the DC and AC components of the time-variant simulated panel's output were fixed throughout all of these tests. Testing over a much broader range of power outputs would be more representative of the range of possible solar panel operating conditions. Also, the frequency of the changes in the panel's power output was also fixed at $f_{motion} = 0.5$ Hz. What might be more realistic to test the algorithm's performance over a range of sample and motion frequencies, varying f_{sample}/f_{motion} from, say, 10 to 1000. Testing should not merely be limited to simulations with a fixed f_{sample}/f_{motion} ratio, either, but allow it to change (perhaps randomly) over the course of the simulation.

function call

readPanel → Calc $\Delta P_k / \Delta B_k$ → Calc $\mu_k = \text{const} * \Delta P_k / \Delta B_k$

$\mu_k < \texttt{-MAX\_STEP}$? — yes → $\mu_k = \texttt{-MAX\_STEP}$
no → $\mu_k > \texttt{MAX\_STEP}$? — yes → $\mu_k = \texttt{MAX\_STEP}$
no

is_mppt? — yes → $B_{k+1} = B_k + \mu_k$
no → $\Delta P_k / \Delta B_k > 0$? — yes → $B_{k+1} = B_k - 3\mu_i$
no → $B_{k+1} = B_k - \mu_k$

$B_{k+1} <= \texttt{DUTY\_MIN}$? — yes → $B_{k+1} = \texttt{DUTY\_MIN} + 10$
no → $B_{k+1} >= \texttt{DUTY\_MAX}$? — yes → $B_{k+1} = \texttt{DUTY\_MAX} - 10$
no

$|B_k - B_{k+1}| <= \texttt{MIN\_STEP}$? — yes → $B_{k+1} \mathrel{+}= \text{dither}$
no

writeDuty

return
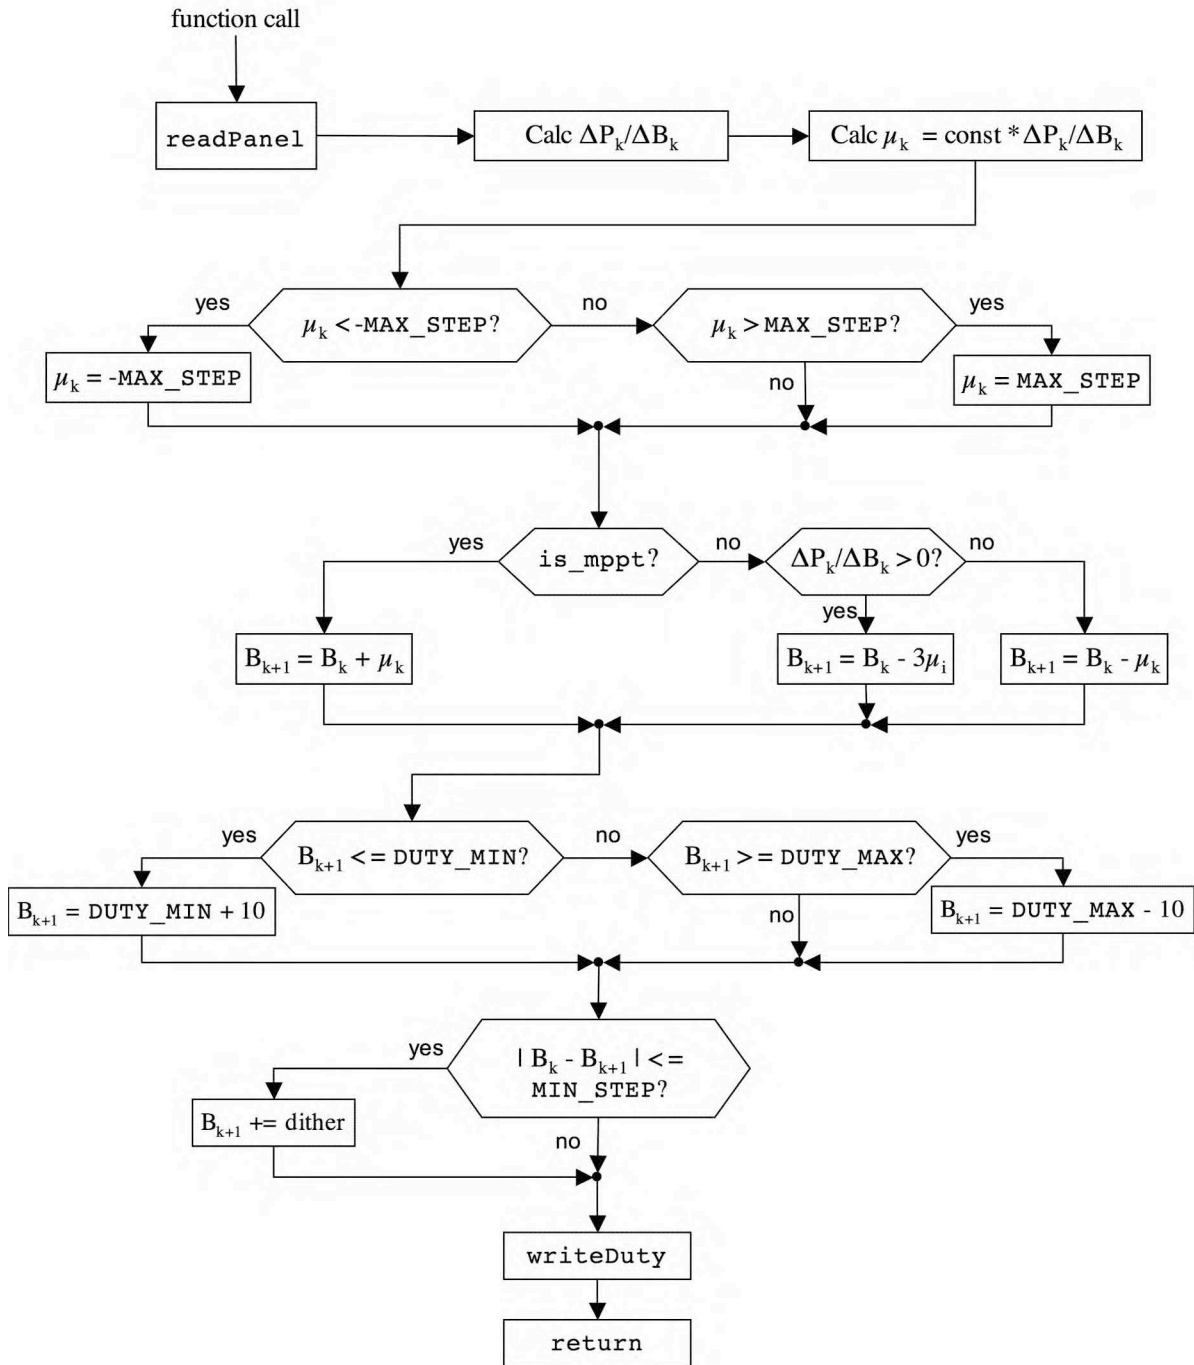
Figure 4-13 – A flowchart documenting the operation of the maximum power point tracking algorithm. The variable `is_mppt` *tells the function whether to ascend or descend the gradient.* `MAX_STEP` *is a constant that prevents the algorithm from making jumps that are too large.* `DUTY_MIN` *and* `DUTY_MAX` *are also constants that place limits on how large or small of a duty cycle command we can make, to force the algorithm to operate within a certain range of commands that avoid the area far too the left of the MPP on the P-B curve and also its x-intercept.* `MIN_STEP` *forces there to be a change in B from one time step to the next.* `writeDuty` *is a function that writes the new duty cycle command to the DAC.*

4.4 The Main Control Algorithm of the Power System

As briefly discussed in Chapter 3, the overall control goal of the power system is not to deliver the maximum amount of possible power to the system, but only as much as is needed at any given time. I refer to this as "power matching" – the power delivered matches the demand (as closely as possible). Usable power that is not extracted from a particular solar panel (because it is not operating at its current maximum power point) is dissipated as heat through its surface. In contrast, were the robot to operate using maximum power point tracking all the time on all of its panels, the excess power that is extracted from the panels, but not used or needed by the system, would have to be disposed of. Generally, such excess power is dissipated through a resistor bank, and then out into the open air. Dissipating unneeded power through the solar panels, which have a considerable surface area, appeared to be a viable option, without the additional weight, expense, and complexity of some other power dissipater.

How is one to determine how much power is needed at any given time? One could explicitly measure the power being delivered to the various subsystems in the robot (the four motor controller and the housekeeping power supply). However, that would require additional sensing circuits. The approach used in this thesis is to minimize the current flowing into or out of the batteries. If the amount of power being provided by the solar panels is greater than what the system currently needs, it has no place else to go other than into the batteries. On the other hand, during those times when the panels cannot deliver enough power for all the robot's needs, the batteries can make up the shortfall. However, when the power delivered by the panels exactly matches that needed by the robot, the batteries will not be utilized at all, and the current through them will be zero. This, then, is

the overall control goal under normal conditions – to minimize the battery current by adjusting the duty cycle commands of the five DC/DC boost converters.

In order for the batteries to be effective buffers for the power bus, for those times when the battery current is not minimized, they need to be able to accept power as well as provide it. Therefore, the batteries should neither be fully charged or discharged. One could choose a setpoint right in the middle, a 50% state-of-charge, which would correspond to a bus voltage of approximately 47.6 V. However, as the batteries also exist to provide backup power in the event of an emergency, they should be left mostly charged. The operational setpoint chosen for the control algorithm is a battery stack voltage of 48.6 V,[6] which corresponds roughly to an 80% state-of-charge.[7] This setpoint, being in the constant-current phase of its charge cycle, allows the batteries to accept a full 3 A of current. Were the setpoint located much higher, we would be in the constant voltage portion of the charge profile and able to accept less current. Indeed, because the voltage is constant during that phase the charge profile, we would not be able to tell how much current the batteries could accept, and run the risk of damaging the batteries through overcharging, nor would we be able to estimate, even crudely, the batteries' state of charge.

The control law is a sort of nested Bang-Bang controller. The inner loop actually consists of five independent loops, one for each DC/DC boost converter. Each control loop executes a Bang-Bang control law based on $\partial P / \partial B$ as described above. The outer loop monitors the current flowing in an out of the battery, and calls for more or less power from

---

[6] Although the motors are intended to operate up to 48 V, the manufacturer has informed us that operating the motors a little above this level will not damage them.

[7] It should be noted that the battery voltage is only a rough gauge of the battery state-of-charge, particularly with Li-ion batteries.

the inner loops accordingly. When the outer loop calls for more power, the control effort for each DC/DC converter is to seek out more power through maximum power point tracking:

If $\Delta P/\Delta B > 0$ (i.e., to the left of the MPP), then increase B.
If $\Delta P/\Delta B < 0$ (i.e., to the right of the MPP), then decrease B.

When the outer loop calls for less power, the control effort is to migrate each solar panel away from its maximum power point and towards the open-circuit voltage of each panel (i.e., operate in the regime where $\Delta P/\Delta B < 0$).

If $\Delta P/\Delta B > 0$ (i.e., to the left of the MPP), then increase B until past the MPP.
If $\Delta P/\Delta B < 0$ (i.e., to the right of the MPP), then increase B until B is maximized.

There are additional features to this algorithm. For instance, if the power demand is low, the operation of some of the panels will be driven all the way towards their minimum power (high duty cycle command). Once the power being extracted from a panel has been minimized, there is no sense in trying to minimize it further, so the DC/DC converter for that panel will be disabled until there is a greater demand for power. When a panel is about to be reintroduced to the system, it is necessary to start operating the panel near to its open-circuit voltage, so that initially its contribution to the power system is be small. Therefore, before the DC/DC converter for that panel is re-enabled, the initial duty cycle command to use, B_o, needs to be determined using the following formula, a combination of equations 3.4 and 4.3:

$$B_o = 3925 \cdot \frac{2}{\pi} \cos^{-1}\left(\frac{1}{58.6\%}\left(1 - \frac{V_{OC}}{V_{batt}}\right)\right) - 20 \qquad (4.5)$$

The subtraction of 20 at the end of this expression ensures that the initial power is substantial enough to be measured, so that $\Delta P/\Delta B$ can be properly determined. It also ensures that the panel begins operating well away from the discontinuous conduction mode of the DC/DC converter, which would occur for overly high duty cycle commands (low duty cycles). At

the opposite end of the spectrum, when all of the panels are operating near their respective maximum power points and the power produced is still not enough for the current demand, the power microcontroller will inform the master microcontroller that the robot needs to slow down to reduce the power demand.

One must accept that this control procedure will not be perfectly able to minimize the current flowing in and out of the batteries. The best that can be hoped for is that, on average, the net charge through the batteries is zero. However, realizing that this may not be possible, and that the batteries will gradually drift from the desired setpoint of 48.6 V, 0 A, towards a higher or lower state-of-charge, there is yet a further level of control that monitors the battery voltage. If the bus voltage drifts too low, to 46 V, then the batteries need to be charged. In that case, the current through the batteries should no longer be minimized (i.e., driven to zero amperes), but rather towards a charge current of –1.5 A.[8] On the other hand, if the bus voltage has reached its maximum level of 49 V, then the batteries need to be discharged, and the battery current setpoint is raised to +1.5 A. In both cases, once the bus voltage has returned to the desired value of 48.6 V, the battery current setpoint is returned to zero amperes. The operation of this algorithm is given in the flowchart in Figure 4-15.

4.5 Other Modes of Operation

In addition to the normal mode of operation described in the previous section, there are several other modes of operation that will be necessary for the Cool Robot. Due to time constraints, there was not time to implement these modes in software. However, an overview of each in turn is given here.

---

[8] Recall the convention established in Chapter 3 that current leaving the batteries is defined as positive, while a charge current is defined as negative.

periodic interrupt

Load setpoint data and limits
$V_{batt,o} = 48.6V$ $I_{batt,o} = 0$ A
$V_{batt,max} = 49$ V $I_{batt,max} = 12$ A
$V_{batt,min} = 46.5$ V $I_{batt,min} = -3$ A

readBatt

$V_{batt} > V_{batt,max}$? — no — $V_{batt} < V_{batt,min}$? — no

yes

$I_{batt,o} = 1.5$ A
(until V_{batt} sinks to $V_{batt,o}$)

yes

$I_{batt,o} = -1.5$ A
(until V_{batt} rises to $V_{batt,o}$)

$I_{batt} > I_{batt,o}+\varepsilon$? — yes — Need More Power — $I_{batt} > I_{batt,max}$? — yes

no

•Inform master μC
•Disable all DC/DCs
•Disable motor power
• return

no

Nth consecutive time ? — yes — All panels enabled? — yes — Tell master μC to slow down

no

no

For each panel do the following: — panel disabled? — yes — • initPanel
• Enable panel for next iteration

no

mppt(this panel) — next panel or return

$I_{batt} < I_{batt,o}-\varepsilon$? — yes — Need Less Power — $I_{batt} < I_{batt,min}$? — yes — • Disable all panels
• return

no

no

No Change

Nth consecutive time ? — yes — Tell master μC that it can speed up

return

anti- mppt(this panel) — no — panel disabled? — For each panel do the following:

$P_i < P_{min}$ — yes — • Disable DC/DC
• return

no

Next panel or return

Figure 4-14 – Flowchart describing the normal operations algorithm.

Stationary – When the robot is not driving, such as the months when it is on station, its power requirements are considerably lowered. The only power requirement is for the housekeeping supply, which provides power to the internal electronics and scientific payload, which all together total 30 W at most.[9] This is a small enough amount of power that all but one solar panel and DC/DC boost converter can be disabled. The power will be delivered entirely by the top panel, which is assured of having direct insolation all the time. The same power matching algorithm will be used to ensure that the power extracted from the top panel equals that of the instantaneous demand. In addition, because the robot does not need to use its motors when stationary, the MOSFET that connects the motor controllers to the power bus will be turned off, so that they do not consume idle power. The power microcontroller will enter this mode when instructed to do so by the master microcontroller. The master microcontroller will periodically bring the system out of this mode of operation to move the robot, to prevent it from being snowed in. While the payload power supply is currently always on, future generations of Cool Robot may choose to optionally disable some or all of it, since some instruments do not need to record data when the robot is driving, or because the instruments would draw too much power to be able to operate while the robot is driving. When the robot is in stationary mode and its power requirements are much less, the full payload power supply can be enabled.

Quiet – An extension of the Stationary mode, the Quiet mode can be used when the robot is taking scientific data that needs a low level of electromagnetic noise. The motor, motor controller, and DC/DC boost converters, with their significant power handling and substantial magnetic components, can throw off a considerable amount of magnetic noise

---

[9] This is a limitation of the DC/DC converters used in the housekeeping power supply. If additional power is needed for the science payload, additional capacity can be added in future designs.

that could disrupt the measurements of the on-board magnetometer. In addition, being rapidly switching devices power electronic devices, they can cause significant amounts of switching noise to appear in the power and ground supplies of the Cool Robot. Therefore, while in Quiet mode, all five DC/DC boost converters are disabled and the MOSFET directing power to the motor controllers is turned off. The robot will then run on battery power alone. Because the batteries have a significant capacity[10] the robot can operate its electronics on battery power alone for a long while. Eventually, however, the batteries will become discharged enough that the power microcontroller will signal to the master microcontroller that it needs to leave Quiet mode and enter Stationary Charge mode to replenish the batteries.

Stationary Charge – When the robot is in Stationary Charge mode, all five DC/DC boost converters are enabled and power is delivered to the batteries to charge them back up. In this case, the desired battery current is -3 A, until the battery stack voltage raises up to at least 48.6 V, which is an approximately 80% state-of-charge.

Rather than stopping at an 80% charge, it may be advisable to fully charge the batteries using Stationary Charge mode. Aside from allowing the robot to operate in Quiet mode for longer periods of time, it would be healthy for the batteries to be fully cycled once in a while to prevent drift between the cells. Therefore, it even may be desirable to periodically enter Stationary Charge mode while the robot is driving. A determination of whether Stationary Charge mode should extend to a full charge cycle probably won't be possible until more experience with the operation of the robot for prolonged periods can be gained. If a full charge is desired, then the charge should continue until the battery stack voltage reaches 49.2 V. Once there, a different control mode is needed to hold the battery stack voltage at 49.2 V; voltage control rather than current control. The charge current will

---

[10] The robot is able to drive for several hours on a full charge.

trail off of its own accord, following the charge profile given in Figure 3-31. Once the charge current has reduced to 300 mA, the charge cycle will be declared as complete. When they are fully charged, the batteries will be unable to buffer the bus voltage as they do in all other modes of operation, since they cannot safely accept additional power. If the robot exits from Stationary Charge mode to Quiet mode, the batteries will naturally discharge, being the sole source of power to the robot. If the robot exits from Stationary Charge mode to Stationary or Normal Driving mode, the power microcontroller must first set the batteries' current setpoint to -3 A. With a setpoint of -3 A, the batteries will be power sources, rather than sinks or buffers, until the bus voltage has dropped to its normal level of 48.6 V.[11]

Emergency Driving – There may be situations when the robot requires as much power as possible, such as when it is tackling an obstacle or trying to free itself from the snow. In this situation, determined by the master microcontroller, the power system will still desire to deliver as much power as is necessary from the solar panels. But, in the case that the solar panels cannot deliver all the necessary power, supplemental power will flow naturally from the batteries. In the Normal Driving mode, a prolonged power draw from the batteries would cause the power microcontroller to ask the master microcontroller to slow down. However, since the master microcontroller has deemed that there is some mobility emergency, the power microcontroller will simply allow the batteries to be utilized as much as necessary. However, the power microcontroller cannot allow the batteries to be discharged completely, lest the power system become terminally discharged. Therefore, the power system will instruct the master microcontroller to slow down or stop when the battery stack voltage has reached some lower limit, such as 44 V. If the discharge from the batteries

---

[11] The Stationary and Drive modes would do this naturally anyway, as they periodically check the battery stack voltage, and reduce the battery current setpoint if the stack voltage is too high.

continues, the power microcontroller will eventually cut power to the motor controllers when the battery bus voltage has dropped all the way to 42 V, which is the lower limit of the measurement range and an approximately 90% discharge of the batteries' capacity.

4.6 FEATURES OF THE SOFTWARE

The operations described in the previous two sections are implemented in the program that runs on the power microcontroller. The Jackrabbit RCM3100 from Z-World is programmed using Dynamic C, which is a variant on the standard C programming language.[12] There are several differences between Standard C and Dynamic C. These differences come about from the fact that C is intended to be executed from a computer that already has an operating system up and running, helping the C program coordinate tasks and doling out resources to the program. Dynamic C, in contrast, is intended to be run in embedded, real-time systems that have full control of (and responsibility for) the microcontroller. In addition, Dynamic C has the ability to establish variables that can be stored in nonvolatile memory (such as flash memory or battery-backed SRAM), so that they can be retained and reloaded when the device is power cycled. One very important enhancement in the Dynamic C language is the ability to execute several tasks nearly in parallel by switching the processor's attention from one to the next. A device like a digital signal processor, which is the very embodiment of real-time computing, multitasks and performs periodic functions through a coordinated ballet between hardware and software, based primarily on interrupt service routines (ISR). For the Jackrabbit, programmed with Dynamic C, these things are largely software controlled, and relatively easy for the

---

[12] Dynamic C is both the name of the language and the integrated development environment for programming Jackrabbit microcontrollers.

programmer to incorporate into their code. The operation is similar to the execution and handling of process threads in the Java programming language, the major difference being that Java is an object-oriented programming language. Lastly, there are compiler-level differences in how Dynamic C handles things like libraries and other pieces of source code other than `main()`. The interested reader is directed to [Z-World 2004] for further information.

Multitasking is an often a misnomer. A desktop computer appears capable of multitasking (how many applications can one run simultaneously?); in reality, however, the microprocessor at its heart can really only do one thing at a time.[13] Multitasking then is nothing more than rapidly switching between various tasks in such a way as to make them seem to be executing in parallel. If the operating system and program designs are good, the processor will operate on one task while other tasks are idling around and waiting for something to fall into place. The idle time could arise from waiting for data to arrive from memory or the network. In an embedded controller, the idle time could arise from waiting for user input, or in simply waiting for a specified time to pass between periodic tasks. When one process voluntarily surrenders the processor's attention while it is waiting, allowing other tasks to proceed, it is called cooperative multitasking.

Cooperative multitasking in Dynamic C is implemented using costatements and cofunctions. A costatements is a block of code that needs to be executed sequentially, but can collectively be set aside so that other blocks of code and be executed. A costatement can be set aside if, for instance, it is waiting for something to happen, or for a specific length of time to elapse. Costatements are usually contained in an infinite loop. Each time the loop

---

[13] This is also not quite true – each *functional block* in the processor can only handle one operation at a time, but the processor may have a dozen such blocks (an ALU, a floating point math unit, two integer math units, etc). Multicore and multiprocessor systems are capable of true multitasking.

executes, it examines each costatement block in turn and determines whether it is time to begin or resume its execution. If the time is not yet ripe, the loop moves on to the next costatement block and examines it. The following example, taken from the Dynamic C User's Manual [Z-World 2004], should help to illustrate this functionality.

In this example, the following sequence of events needs to occur while other tasks are happening at the same time:

1. Wait for a Pushbutton to be pressed.
2. Turn on Device #1
3. Wait 60 seconds.
4. Turn on Device #2
5. Wait 60 seconds
6. Turn off both devices.

The following bit of Dynamic C code executes these functions, while allowing other tasks to be processed simultaneously:

```
while(1) {
      costate {…}          // some other task
      costate {
            waitfor(button_pushed());    // waiting
            turn_on_device1();
            waitfor(DelaySec(60));       // more waiting
            turn_on_device2();
            waitfor(DelaySec(60));       // more waiting
            turn_off_device1();
            turn_off_device2();
      }
      costate {…}          // some other task
      costate {…}          // some other task
}  // end while
```

In the above example, the `while` loop sequentially looks at each costatement block. When it comes upon the our costatement for the first time, it recognizes that the costatement is waiting for some logical condition `button_pushed()` to become true. Therefore, the program puts a place marker at `waitfor(button_pushed())` and moves on to the next costatement. The next time the `while` loop encounters the costatement in question, it checks to see if `button_pushed()` has occurred. Once it has, the program then executes

turn_on_device1(), and encounters the first 60 second delay. Realizing that this is another wait, the program moves the place marker to waitfor(DelaySec(60)) and continues on to the next costatement. The while loop will keep iterating, checking each time to see if the 60 seconds has elapsed yet. When it finally has, the program will turn_on_device2(), then encounter the second 60 second delay, and so on. In each case, the program remembers where it left off in the execution of a particular costatement, and return to that place each time the while loop iterates. It is important to note that, while the program is processing one task, other tasks are ignored. That is, until one costatement has run through to its finish, or voluntarily yielded the processor's attention using a waitfor() or similar command, the processor will not handle any other tasks. This is a feature of cooperative multitasking that makes it relatively easy to program. If one requires that a task can be forcibly suspended part way through its execution, so that the processor can deal with another, more pressing task, then preemptive multitasking is required. Preemptive multitasking is considerably more difficult to implement in Dynamic C with a Jackrabbit microcontroller, and is not used in the Cool Robots software framework. There are numerous more details and subtleties involved in multitask programming, which are adequately discussed and explained in the Dynamic C User's Manual [Z-World 2004]. This brief introduction and a knowledge of the standard C language should be enough to allow one to follow the power system code given in Appendix I. The main loop of the program has this basic structure:

```
while (1)
{
   costate
   {
      waitfor(DelayMs((int)(1000/FSAMPLE)));
      switch(ops_mode)
      {
            case 1:      NormalOps(&PanelArray[0], &Batt);
```

```
                                    break;
                case 2:     Stationary(&PanelArray[0], &Batt);
                                    break;
                case 3:     Quiet(&PanelArray[0], &Batt);
                                    break;
                case 4:     StationaryCharge(&PanelArray[0], &Batt);
                                    break;
                case 5:     Emergency(&PanelArray[0], &Batt);
                                    break;
                default:    enterNormal();
                                    break;
            } // switch
        } // end costate
    } // end while
```

Following the example from [Bhide and Bhat 1992], our photovoltaic control program establishes a data structure for each panel-DC/DC converter pair. The `panel` datatype is a `struct`, and contains fields that hold information specific to that entity and useful for executing the Power Matching control algorithm (refer to Appendix I-3):

```
typedef struct {
    float voltage;              // the latest sample of panel's operating voltage
    float current;              // the latest sample of the panel's operating current
    float power;                // the latest calculation of the panel's power output
    float prevPower;            // the previous calculation of the panel's power output
    float mu;                   // the latest calculated adaptive step size
    char is_enabled;            // a boolean, whether the DC/DC is operational or not
    int add_sub;                // a boolean, whether a perturbation is added or subtracted
    int dutyCmd;                // the latest value of the duty cycle command
    int prevDuty;               // the previous value of the duty cycle command
    unsigned int Vchannel;      // a binary code denoting the ADC channel for voltage
    unsigned int Ichannel;      // a binary code denoting the ADC channel for current
    unsigned int Dchannel;      // a binary code denoting the DAC channel for duty cycle
    unsigned int Echannel;      // which bit register in PDDR enables the DC/DC conv.
} panel;
```

When `main()` of power system is first started, an array of these `panel` variables is created in memory, `PanelArray`. A pointer to this array, or to specific elements in it, is passed to the various functions that execute the control algorithm and operational modes of the power system. The battery is also declared as type `panel`, as it uses a number of the same fields like `voltage`, `current`, and ADC addresses. Next, the program assigns ADC, DAC, and general purpose I/O addresses to the various "channel" fields in each `panel` struct.

Once the data structures have been initialized, the program enters into the `NormalOps` control mode. Associated with each control mode (`NormalOps`, `Stationary`, etc.) is an initialization function (`enterNormal`, `enterStationary`, etc.). The purpose of these initialization functions is to establish the bounds and setpoints that each particular operating mode will work under. For instance, the `enterNormal` function sets the following global variables to their appropriate values:

```
void enterNormal()
{
        ops_mode = 1;
        Vbatt_max = 49.0;
        Vbatt_setpoint = 48.6;
        Vbatt_min = 46;
        Ibatt_max = 12;
        Ibatt_setpoint = 0;
        Ibatt_min = -3;
        is_mppt = 1;
} // enterNormal
```

Initially, all of the panels are offline. During the first few iterations of the control algorithm, the panels are brought online one at a time as additional power is needed. When a panel and its DC/DC converter are about to be brought online, a call to the function `initPanel`, in the `powerCTRL` library, is made. This function uses the ADC to measure the panel's open circuit voltage. Based on V_{OC}, and a measurement of V_{batt}, an initial value of `dtyCmd` is determined using equation 4.4. The initial value of `prevDuty` is set to a program constant, `MAX_DUTY`, so that the first calculation of ΔB will be negative. The value of `power` is set to 1, and `prevPower` to 0, so that when the panel is brought online and the first real value of `power` is found, the resulting ΔP will be positive. As a result, if all has been done correctly, the initial calculation of $\Delta P/\Delta B$ will accurately point the program towards that panel's maximum power point.

167

The `main()` of the power system program, given in Appendix I-5, is actually quite brief. It serves mostly to initialize the program and its variables, then enters an infinite loop, in which is the costatement that executes at regular intervals, and switches among the various modes of operation. The real execution of the control algorithm in its various flavors consists of the following high-level functions, which are given in the `powerCTRL.lib` library (Appendix I-4) :

```
NormalOps(&PanelArray[0], &Batt)
Stationary(&PanelArray[0], &Batt)
Quiet(&PanelArray[0], &Batt)
StationaryCharge(&PanelArray[0], &Batt)
Emergency(&PanelArray[0], &Batt)
```

The behavior of these various functions has already been described. With each operational mode is a corresponding `enter` function, such as `enterNormal`. At the time of this writing, only the `NormalOps` and `Stationary` are implemented in code. However, as one might imagine from their description, the other operational modes are quite similar aside from a few parameters such as the allowable range of battery states, the state of the MOSFETs on the power bus, and which DC/DC converters are enabled. The `NormalOps` and `Stationary` modes utilize a number of other functions, found in `powerIO.lib` (Appendix I-3), which are shared with the other operational mode functions:

`readPanel(panel* pSinglePanel, int Samples)` – Uses the ADC functions in `AD7490EV.lib` to sample and store the `current` and `voltage` of the panel pointed to by `pSinglePanel`. To allow one to make more accurate readings, the measurements are actually averaged over `Samples` repeated measurements. Once these values have been read and scaled, the panel's `power` is also calculated. As this is a new measurement of `power`, the previous measurement is first moved to `prevPower`.

`readBattery(panel* pBatt, int Samples)` – much like `readPanel`, this function uses the ADC to determine the current and voltage of the battery stack. The battery power is not calculated, however, as it is not necessary for the control algorithm.

`writeDuty(panel *pSinglePanel)` – uses the `Dchannel` address contained in the `panel` struct, and the latest value of `dutyCmd`, to cause the DAC to output the latest duty cycle command.

Because most of the operational modes make use of the maximum power point tracking algorithm described earlier in this chapter and given in Figure 4-14, it would be wasteful to repeatedly embed that code in each operational mode function. Instead, the maximum power point tracking algorithm, for a single panel, is executed in the function `mppt(panel *pSinglePanel)`, given in Appendix I-4 and accessible to all the operational modes.

4.7 TESTING RESULTS

Before the solar panels and boost converters were connected to the system, it was prudent to test the functionality of the `NormalOps` and `Stationary` operating modes in simulation. If there were there some problem or error with the overall control algorithm, serious damage to a number of the components of the power system may have resulted. This is the difference between developing a signal-level electronics system and a power electronics system – the consequences for mistakes are much higher. What's more, it would be much faster to identify and debug problems in the code on the benchtop, rather than with the robot and solar panels sitting outside in the sun.

To spoof the control algorithm into believing that it was operating on a real solar panel, I inserted several lines into the `readPanel` function, which utilizes the ADC to

measure the panel `voltage` and `current`, and calculate the panel `power` from these two measurements. Rather than actually measuring the voltage and current of a panel that wasn't actually present, I simulated the panel's response in the same way that I did when testing the MPPT algorithm in MATLAB. Based on the panel's current `dutyCmd`, I was able to calculate (approximately) what the panel's resulting current and voltage would be, bypassing the actual ADC readings. This simulation code is still contained in the `readPanel` function, though it is commented out in the final implementation.

This addition of a few lines of code allowed me to successfully debug the control algorithm. It has some limitations, however. For instance, the power produced by the panel is not real power that offsets the power supplied by the batteries; it is only a number. As a result, the control algorithm was constant either seeking or avoiding the maximum power points of the solar panels – it could not actually reach the steady state condition where the batteries were supplying or accepting no power. While this, too, could have been simulated in situ by subtracting the panel's (imaginary) power from the measurements of the batteries' (real) power,[14] it would have required a more invasive alteration of the code to do so. By this point, however, I had enough confidence in the control algorithm to test it with the real thing.

At this point, all of the necessary pieces for the power system were in place. The hardware, described in Chapter 3, was constructed and functioning reliably with the microcontroller. The algorithm described in this chapter for maximum power point tracking had been well simulated in MATLAB, and implemented in Dynamic C. Finally, the Dynamic C code to utilize that algorithm within the `NormalOps` and `Stationary` modes had been written and tested in a limited way.

---

[14] Actually, as the variable of interest is the battery voltage, what would have been necessary is to add the collective panel power, divide it by the bus voltage, and then apply that to the battery current measurements.

Testing began with the `Stationary` operational mode. As it only uses one panel, this seemed a straightforward task to implement and monitor. The control of one panel in `Stationary` could then easily be replicated for multiple panels in `NormalOps`. From a more practical standpoint, it would not have been easy to provide the power system with sufficiently large loads warranting more than one panel's output short of actually driving the robot. Driving the robot would have required a second microcontroller and someone else to do the actual driving, which is an unnecessary complication at the start of testing. What's more, it seemed prudent to test the algorithm's control of a panel whose power curve was not changing in time, which is easier to ensure when the robot is not moving.

Testing was done in late May 2005 outside the Thayer School. During the course of several tests conducted over several hours, the power demand of the robot was changed from its quiescent level of less than 6 W to nearly 60 W. The load could be changed in several ways. One of the simplest ways was to alternately connect or disconnect power to the motor controllers, which collectively draw 10-20 W when idling. Another load would be to connect power resistors to the housekeeping power supply, mimicking the power draw of a test payload up to 20 W. A third way to increase the power draw would be to place the robot up off the ground and let its motors run in an uncontrolled fashion.[15]

In order to be able to follow the program and trace its progression, I added a number of `printf` statements to the code. The following sample program output and detailed discussion should illustrate that, for a single panel and moderate loads at least, the control

---

[15] On power up, the DAC circuit that provides the motors with an analog velocity command swings to the negative supply rail, which causes the motors to drive at full reverse speed unless 1) the controllers are disconnected from the DAC circuit or 2) their power supply is disconnected. As the master microcontroller was not being used during these tests, there was no way to actively assert a "stop" command to the motors. When off the ground and spinning in this uncontrolled fashion, each motor draws approximately 15 W.

algorithm presented in this chapter is able to meet its task. Once every two seconds, the program will sample the battery and panel voltage and current and output them, so that one can follow the progress and effectiveness of the control algorithm. Lastly, because the robot is not driving and the panel's orientation to the sun is not rapidly changing, the update rate of the algorithm was slowed down from its typical value of 50 Hz to a more sedate rate of 4 Hz, so that the program output would be readable.

When this particular test began, I was supplying power to the motor controllers, but not driving the wheels, so that the power draw from the batteries was approximately 20 W. When the program enters the Stationary operations mode, it recognizes this power draw from the batteries and brings the sole panel online using the function `initPanel`:

```
Entering Stationary Operations control mode.
Battery Voltage = 46.518184 V
Battery Current = 0.440369 A

***Seeking More Power***
Initializing Panel.
Panel voltage = 31.971184 V
Necessary duty cycle is 0.312704
Calculated dutyCmd is 2598
```

At this point, the panel is online and providing some small amount of power to the system. One limitation of my maximum power point tracking algorithm, when used on the robot, is that it is not very accurate or effective when the panel's power is quite low. It boils down to the fact that the signal-to-noise ratio of the power measurements is rather low, and the algorithm becomes susceptible to the noise-induced errors discussed in section 4.2. Therefore, the overall control algorithm does not use the `mppt` function unless the panel's power is above some minimum threshold, like 15 W.

When more power is needed, the program simply decrements the duty cycle command B by a fixed amount in open-loop until either the power demand is met, or the panel's power contribution becomes large enough that the mppt function will be effective.

```
***Seeking More Power***
Autodecrementing the dutyCmd.

***Seeking More Power***
Autodecrementing the dutyCmd.

***Seeking More Power***
Autodecrementing the dutyCmd.

Battery Voltage = 46.722286 V
Battery Current = 0.144534 A
Panel Voltage = 31.693584 V
Panel Current = 0.422385 A
```

Within the space of just four iterations, the control algorithm has brought the panel online and operated it so that it is now providing 13 of the needed 20 W. Were the update rate of the program at its full speed of 50 Hz, this would have occurred in less than a tenth of a second. A few iterations later, the panel's power has become sufficiently high for the maximum power point tracking algorithm to take over:

```
***Seeking More Power***
Panel voltage = 31.589909,        Panel Current  = 0.548822
P          dP            dB          dPdB          mu            B
17.337248   2.003963    -5          -0.400793    -31.685442   2443
```

The numbers at the bottom of this output block show the mppt algorithm at work. The first number is the panel's instantaneous power output, followed by the change in power from the previous iteration ΔP. Next is the change in duty cycle command ΔB the last iteration caused. From these two, the slope of the power curve is estimated $\Delta P/\Delta B$, and the resulting step size μ is calculated. This step size is applied to the duty cycle command, which results in a new value of B, 2443. This data is then written to the DAC, which implements the change in B to the DC/DC boost converter, which changes its duty cycle and results in an increased power delivered to the system. In this particular case, however, the

173

additional power is a bit more than needed, and so the next iteration causes uses the `mppt`

algorithm to descend the gradient of the power curve and correct the overshoot:

```
***Seeking Less Power***
Panel voltage = 31.411625,       Panel Current  = 0.799700
P          dP           dB        dPdB        mu          B
25.119867   7.782619    -32       -0.243207   -19.227190  2462
```

After this little adjustment in duty cycle to 2462, the battery current has dropped low

enough so that it is within the ±50 mA deadband around the 0 A setpoint. As a result, we

have achieved a sort of steady state, with the panel providing just as much power as is needed

by the system. In this case, the program output is simply *NC*, for "No Change," for the next

five iterations:

```
*NC*   *NC*   *NC*   *NC*   *NC*
```

The next time the program measures and outputs the system state, we see that, in fact, we

have achieved our goal:

```
Battery Voltage = 46.805904 V
Battery Current = 0.011536 A
Panel Voltage = 30.968826 V
Panel Current = 0.612319 A
```

The batteries are providing little power to the system (about 500 mW); the burden has been

assumed by the solar panel.

This situation would persist, with many *NC* outputs, and occasionally a single

iteration where the operating point is adjusted a little one way or another. What is of greater

interest is the reaction of the algorithm when we suddenly require more power of it. At this

point in the test, an additional demand of about 25 W was added by attaching some power

resistors to the housekeeping power supply. The algorithm responds by resuming its ascent

of the panel's power curve:

```
***Seeking More Power***
Panel voltage = 30.872902,       Panel Current  = 0.794371
P          dP           dB        dPdB        mu          B
```

```
24.524549    5.561756    -15         -0.370784   -29.313028  2381

***Seeking More Power***
Panel voltage = 30.663612,        Panel Current  = 1.059902
P         dP          dB           dPdB        mu          B
32.500415    7.975868    -30         -0.265862   -21.018259  2359

***Seeking More Power***
Panel voltage = 30.527962,        Panel Current  = 1.241066
P         dP          dB           dPdB        mu          B
37.887226    5.386810    -22         -0.244855   -19.357485  2339
```

These three iterations of mppt are sufficient to raise the panel's power output to meet the

demand, as there is no further change to the panel's duty cycle command for the next five

iterations of Stationary:

```
*NC*  *NC*  *NC*  *NC*  *NC*
Battery Voltage = 46.840938 V
Battery Current = -0.011614 A
Panel Voltage = 30.455292 V
Panel Current = 1.439215 A
```

I allow this state to persist for a while, allowing the control algorithm to make small

adjustments as needed over the next 12-14 seconds:

```
 *NC*  *NC*  *NC*  *NC*  *NC*  *NC*  *NC*  *NC*
Battery Voltage = 46.840328 V
Battery Current = 0.011233 A
Panel Voltage = 30.984329 V
Panel Current = 1.402583 A

 *NC*  *NC*  *NC*
***Seeking Less Power**
Panel voltage = 31.016304,        Panel Current  = 1.387596
P         dP          dB           dPdB        mu          B
43.038112   -0.419968   -15          0.027998    2.213424    2404

 *NC*  *NC*  *NC*
Battery Voltage = 46.834712 V
Battery Current = 0.018374 A
Panel Voltage = 31.088005 V
Panel Current = 1.369835 A

 *NC*  *NC*  *NC*  *NC*  *NC*  *NC*  *NC*  *NC*
Battery Voltage = 46.849849 V
Battery Current = -0.045254 A
Panel Voltage = 31.041496 V
Panel Current = 1.404026 A

 *NC*  *NC*  *NC*  *NC*
***Seeking Less Power***
Panel voltage = 31.027931,        Panel Current  = 1.385376
```

```
P           dP          dB          dPdB        mu          B
42.985363   -0.597694   8           -0.074712   -5.906485   2409

***Seeking More Power***
Panel voltage = 31.101570,      Panel Current  = 1.379604
P           dP          dB          dPdB        mu          B
42.907848   -0.077515   5           -0.015503   -1.225614   2411

 *NC*   *NC*
Battery Voltage = 46.821163 V
Battery Current = 0.019167 A
Panel Voltage = 31.107384 V
Panel Current = 1.352185 A
```

During a later stage of testing, I turned applied power to only two of the motor controllers, and allowed the wheels to spin uncontrolled. The total power demand increased to approximately 57 W. As it turned out, the sky conditions at this point had deteriorated considerably, so that there was a thickening blankets of clouds blocking the sun. As a result, the single panel's maximum possible power output was only about 50 W. The next (long) excerpt of program output demonstrates how the control algorithm continues to call for more power while maintaining its operation at the single panel's maximum power point. In trying to decipher this output, pay particular attention to the panel's power output (leftmost field) and its duty cycle command (rightmost field):

```
Battery Voltage = 46.619503 V
Battery Current = 0.150881 A
Panel Voltage = 26.409027 V
Panel Current = 1.842506 A

***Seeking More Power***
Panel voltage = 26.508827,      Panel Current  = 1.854717
49.166370   0.507576    6       0.084596        6.687901
1919

***Seeking More Power***
Panel voltage = 26.450691,      Panel Current  = 1.892460
50.056865   0.890495    6       0.148416        11.733304
1930

***Seeking More Power***
Panel voltage = 26.388680,      Panel Current  = 1.854273
48.931812   -1.125053   11      -0.102278       -8.085752
1921

***Seeking More Power***
Panel voltage = 26.470070,      Panel Current  = 1.864930
```

```
49.364818        0.153267        -9        -0.017030        -1.346312
1915


***Seeking More Power***
Panel voltage = 26.431313,        Panel Current  = 1.864486
49.280803        -0.084015        -6        0.014002        1.106993
1920


***Seeking More Power***
Panel voltage = 26.458443,        Panel Current  = 1.842284
48.743968        -0.536835        5        -0.107367        -8.488101
1911


Battery Voltage = 46.642940 V
Battery Current = 0.124697 A
Panel Voltage = 26.416779 V
Panel Current = 1.867594 A


***Seeking More Power***
Panel voltage = 26.408058,        Panel Current  = 1.887131
49.835468        0.499657        -9        -0.055517        -4.389037
1906


***Seeking More Power***
Panel voltage = 26.357674,        Panel Current  = 1.877807
49.494609        -0.340858        -5        0.068172        5.389445
1911


***Seeking More Power***
Panel voltage = 26.427437,        Panel Current  = 1.872478
49.484798        -0.009811        5        -0.001962        -0.155132
1906


***Seeking More Power***
Panel voltage = 26.361549,        Panel Current  = 1.913329
50.438320        0.671295        -5        -0.134259        -10.614108
1895


***Seeking More Power***
Panel voltage = 26.318916,        Panel Current  = 1.915105
50.403484        -0.034836        -11        0.003167        0.250365
1899


***Seeking More Power***
Panel voltage = 26.318916,        Panel Current  = 1.902672
50.076267        -0.327217        4        -0.081804        -6.467195
1892


***Seeking More Power***
Panel voltage = 26.369301,        Panel Current  = 1.922210
50.687320        0.611053        -7        -0.087293        -6.901145
1885


Battery Voltage = 46.657222 V
Battery Current = 0.109622 A
Panel Voltage = 26.327637 V
Panel Current = 1.932200 A
```

The next logical step in the testing would be to try and satisfy the remaining power demand by switching from the `Stationary` operational mode to `NormalOps`, which can perform the power matching algorithm with all panels, and not just a single one. However, the weather on this particular day of testing was further deteriorating, so it was wiser to bring the robot and all the rest of the equipment back inside. Between that day of testing and the date this thesis was finalized, there was not a sufficiently long period of sunny or even partly cloudy skies to allow even brief testing of multi-panel operation.[16]

4.8 Prospective of Master/Slave Communications

As mentioned in Chapter 1, the RCM3100 is capable of being configured in a master/slave setup with other Jackrabbit microcontrollers. The evaluation board that the RCM3100 module is plugged into has a second port for a slave module, and the 8-bit data bus between has been prewired in the circuit board. The microcontroller plugged into the slave port need not be used as a slave device, however. In fact, if the user does not specifically program both devices for their respective roles, they will operate as two independent devices. It is likely that the field testing in Greenland will be done in this manner, with the master and power microcontrollers each executing their necessary functions and oblivious to the existence of one another. However, for the Cool Robot to have its full desired functionality, the master/slave communications must be implemented. While there was not enough time during this thesis to implement the interdevice communications, this section should serve as a primer on how this functionality will ultimately work.

---

[16] The month of May 2005 in the Upper Valley of New Hampshire and Vermont was unseasonably cold, cloudy, and rainy. According to the statistics available at http://www.wunderground.com/history/airport/KLEB/2005/5/24/MonthlyHistory.html#calendar, only 8 of 31 days that month were partly cloudy or clear, only 3 days had temperatures above average, and only 9 days had temperatures within ±5°F of the average.

The interdevice communication is achieved by the master microcontroller reading or writing to several registers on the slave microcontroller. The slave microcontroller then polls these registers and the messages contained within them. The slave microcontroller can write responses and data to some of these registers, which causes a single control line, `SLAVEATTN`, to be pulled low. The `SLAVEATTN` line can be an interrupt for the master microcontroller, who could then read what the slave microcontroller has written. A software library to coordinate this exchange, `slaveport.lib`, is included in the Dynamic C software from Z-World. This library contains the high-level functions necessary for establishing and controlling this exchange, from both the master and slave point of view. The precise details of the exchange back and forth is not important, and the reader is directed to Chapter 8 of the Dynamic C User's Manual [Z-World 2004] for further information.

The mechanics and code of the exchange of information between the Cool Robot's master microcontroller and power microcontroller may not be very relevant to this thesis, but the content of the exchange certainly is, and a brief discussion of the messages that will be passed back and forth is warranted. By and large, the messages are used to determine which mode each microcontroller should operate in. For instance, the power microcontroller can inform the master microcontroller that there is not enough power available to drive at full speed. Based on this information, the master microcontroller will drop out of its full-speed navigation and driving program, `wp_follow_full()`, and enter a slower, variable speed navigation program, `wp_follow_partial()` [Dietrich and Zettl 2005]. Communicating in the other direction, the master microcontroller can inform the power microcontroller that the robot has arrived on station and can enter Stationary, Quiet, or Stationary Charge modes. Finally, the power microcontroller can periodically send information and statistics about the

health of the power system to the master microcontroller, which can then log that data to the CR1000 datalogger. The master microcontroller itself will periodically send information about its location, heading, and motor status to the datalogger; information from the power system will augment this. This data, such as how much power the robot is using and how effectively the control algorithm is able to match the power needs of the robot, is of vital importance during the Cool Robot's testing and first field deployments. In addition, should the robot ever become stuck or in some way incapacitated, having telemetry about the health of the power system will be of great use to a remote operator.

5. Summary, Conclusion, and Future Work

In this thesis, I have endeavored to give the reader an understanding of the various engineering challenges that had to be understood and addressed in the design of the Cool Robot solar power system. The operation of the Cool Robot power system can be described as a multi-input, multi-output, nonlinear, time-variant control problem. After introducing the Cool Robots project and the context that the power system was been designed in, I provided an overview of the physics behind photovoltaic technologies, their behavior and limitations, and the state-of-the-art. In order to understand the wide range of electrical behavior the solar panels could exhibit over the wide range of operating conditions, a numerical model of the A-300 solar cell was developed. From this beginning, the model was extended to try and predict the electrical characteristics of our constructed panels, with limited success. Meanwhile, based on some modeling and empirical data, an estimate of the total possible power available from the solar panels was made, indicating that there would almost certainly be a power surplus to the Cool Robot.

With this knowledge of the operating conditions and behavior of the solar panel power sources, the rest of the power system was designed. The design was broken down into several smaller components, which I described in turn. For each subcircuit, I tried to provide a thorough description of the relevant design issues and decisions.

After describing the design and construction of the power system, I moved into its control. The Cool Robots power system is presented with several difficult control issues that distinguish it from other photovoltaic applications, primarily that each panel is presented with a different set of operating conditions, and that those operating conditions could be changing

rapidly in time. To aid in the development of the control algorithm, I developed a simulation of the varying input-output relationship between the control command and panel power output. This simulation was used to tune and validate an adaptive maximum power point tracking algorithm that proved robust enough to follow the changes in the panel's power characteristic. The maximum power point tracking algorithm formed the inner loop of the main control algorithm for the power system, which attempts to exactly match the power provided by the solar panels to the instantaneous load. This control algorithm was implemented on a microcontroller that coordinates the whole of the power system. Although there was not enough time during this thesis to implement them, a description of the other modes of operation for the power system was also given, as was an overview of the communications that shall eventually take place between the power and master microcontrollers.

For future work, I would emphasize the need for further testing to validate and improve the operation of the power system's control algorithm. This shall be a priority for my work in the next several weeks, as we prepare to take the robot to Greenland for field testing. There are several other areas for improvement in the power system, which primarily have to do with the construction of the power system itself. Throughout Chapter 3, I tried to indicate those places in the power system that could stand to be modified or redesigned. I would recommend, for the construction of more compact and less labor-intensive circuits than I crafted in this thesis, that the redesigned circuits be implemented as printed circuits. On the whole, it is my hope that one would conclude that I have been successful in my task of designing and controlling a solar power system to operate in the Antarctic. Although there is a large amount of work that can and should still be done, I believe that my work in the preceding pages has created a solid enough foundation to allow that work to continue.

6. Bibliography and Works Cited

[Aceves 2004]   R. Aceves, Lawrence Livermore National Laboratory. In correspondence with the author, May 2004.

[Alchimisti 2004]   P. Alchimisti, Advantage Converting. In correspondence with the author, May 2004.

[Apostolopoulos 2000]   D. Apostolopoulos, et al. "Technology and Field Demonstration of Robotic Search for Antarctic Meteorites." *International Journal of Robotics Research*, vol. 19, no. 11 (November 2000), pages 1015-1032.

[Bhide and Bhat 1992]   P. Bhide and S. R. Bhat. "Modular Power Conditioning Unit for Photovoltaic Applications." *Proceedings of the 23rd IEEE Power Electronics Specialists Conference*, 29 June –3 July 1992, vol 1, pgs 708-713.

[Buchmann I]   I. Buchmann, Cadex Electronics Inc. "Part One – Basics every battery user should know." <http://www.batteryuniversity.com/partone.htm>, accessed Sept. 22, 2004.

[Buchmann I-5]   I. Buchmann, Cadex Electronics Inc. "Is lithium-ion the ideal battery?" <http://www.batteryuniversity.com/partone-5.htm>, accessed Sept. 22, 2004.

[Buchmann I-12]   I. Buchmann, Cadex Electronics Inc. "Charging lithium-ion batteries." <http://www.batteryuniversity.com/partone-12.htm>, accessed Sept. 22, 2004.

[Buchmann II]   I. Buchmann, Cadex Electronics Inc. "Part Two – Getting the most out of your batteries." <http://www.batteryuniversity.com/parttwo.htm>, accessed Sept. 22, 2004.

[Callister 2003]   Callister, W. D. *Material Science and Engineering*, 6th edition. John Wiley & Sons, New York, 2003.

[Carey 1994]   P. G. Carey, R. Aceves, N. J. Colella, and K. A. Williams. "A Solar Module Fabrication Process for HALE Solar Electric UAVs." *Proceedings of the 1st IEEE World Conference on Photovoltaic Energy Conversio*n, 5-9 December 1994, Hawaii, vol. 2, pages 1963-1969.

[Chambouleyron 1991]   L. Chambouleyron. "Chapter 9 - Long Term Prospects for Photovoltaics." In *Generating Electricity from the Sun*, Fred Tremble, editor. Pergamon Press, Oxford, 1991.

[CMDL 2004]   Climate Monitoring and Diagnostic Laboratory (CMDL). <http://www.cmdl.noaa.gov/info/ftpdata.html>, accessed March 2004.

[CRREL1995]   P. W. Richmond, S. A. Shoop, and G. L. Blaisdell. "Cold Regions Mobility Models." CRREL Report 95-1, February 1995.

[Dietrich and Zettl 2005]	G. Dietrich and T. Zettl. "Communication, Navigation and Control of an Autonomous Mobile Robot for Arctic and Antarctic Science." Diplomarbeit, Thayer School of Engineering at Dartmouth College and Üniversität der Bundeswehr München, 2005.
[EERE 2005]	United States Department of Energy, Energy Efficiency and Renewable Energy, "Solar Energy Technologies Program: Bandgap Energies of Semiconductors and Light." <http://www.eere.energy.gov/solar/bandgap_energies.html>, accessed March 27, 2005.
[Eisberg and Resnick 1985]	R. Eisberg, R. Resnick. *Quantum Physics of Atoms, Molecules, Solids, Nuclei, and Particles*, 2nd edition. John Wiley & Sons, New York, 1985.
[Enslin 1997]	J. H. R. Enslin, et al. "Integrated Photovoltaic Maximum Power Point Tracking Converter." *IEEE Transactions of Industrial Electronics*, vol. 44, no. 6 (Dec. 1997), pgs 769-773
[Fahrenbruch and Bube 1983]	A. Fahrenbruch and R. Bube. *Fundamentals of Solar Cells: Photovoltaic solar energy conversion*. Academic Press, Inc., New York, 1983.
[Gay 2004]	Dr. Charles Gay, SunPower Corportation, in correspondence with the author.
[Gravenkötter and Hamman 2004]	G. Gravenkötter and G. Hamman. "Development of a Cool Robot for the Antarctic." Diplomarbeit, Thayer School of Engineering at Dartmouth College and Üniversität der Bundeswehr München, 2004.
[Grenfell 1994]	T. Grenfell, S. Warren, and P. Mullen. "Reflection of solar radiation by the Antarctic snow surface at ultraviolet, visible, and near-infrared wavelengths." *Journal of Geophysical Research*, vol. 99, no. D9 (September 1994), pages 18,669-18,684.
[Gronbeck]	C. Gronbeck. "SunAngle." <http://www.susdesign.com/sunangle/> accessed Oct. 26, 2004
[Guinness 2004]	Guinness World Records. *Guinness Book of World Records 2004*. Guinness & Company, Ireland, 2003.
[Hexcel I]	Hexcell Composites, Inc. "Sandwich Panel Construction – Cut and Fold Fabrication." <http://www.hexcelcomposites.com/Markets/Products/Panels/Sandpanel_fab/Spft_p10.htm>, accessed May 2004.
[Hexcel II]	Hexcell Composites, Inc. "Sandwich Panel Construction – Typical Panel Fixings." <http://www.hexcelcomposites.com/Markets/Products/Panels/Sandpanel_fab/Spft_p09.htm>, accessed May 2004.
[Horenstein 1990]	M. Horenstein. *Microelectronic Circuits and Devices*, chapter 3. Prentice Hall, Inc., Englewood Hills, NJ, 1990.
[IRF 2003]	International Rectifier Corporation. "HV Floating MOS-Gate Driver ICs," Application Note AN978, Rev. B. International Rectifier, El Segundo, California, June 2003.

[JPL 1976]   TRW Systems Group, Power Sources Engineering Department. "Solar Cell Array Design Handbook." H. S. Rauchenbach, editor. Prepared under contract no. 953913 with the California Institute of Technology and Jet Propulsion Laboratories. National Aeronautics and Space Administration, Washington D.C., 1976.

[Kahn et al. 2004]   I. Kahn, B. Kasdon, X. Weng. "Solar Power Design of a Antarctic Robot." Bachelor of Engineering Project Final Report. Thayer School of Engineering at Dartmouth College, Hanover, N.H., March 2004.

[Komp 1987]   R. Komp. *Practical Photovoltaics: electricity from solar cells, 2nd edition*. Aatec Publications, Ann Arbor, 1987.

[Krummann 2001]   W. Krummann. "Next Generation GOES Satellite Power Subsystem." *IEEE Aerospace Electronics Systems Magazine*, vol. 16, no 4. (April 2001), pgs. 23-28.

[Lanzerotti 1999]   L. Lanzerotti, A. Shona, H. Fukunishi, and C.G. Maclennan. "Long Period Hydromagnetic Waves at Very High Geomagnetic Latitudes." *Journal of Geophysical Research*, vol. 104, no. A12, pg. 28,423 1999.

[Lever]   Dr. James Lever of USACE CRREL. During meetings and in correspondence with the author 2004-2005.

[Linacre and Geerts 1999]   E. Linacre and B. Geerts, June 1999. "Surface Temperatures in Antarctica." <http://www-das.uwyo.edu/~geerts/cwx/notes/chap03/antarctica.html>, accessed 4/5/2005.

[Maxim 189]   Maxim/Dallas Semiconductor. "Advanced Rechargeable Lithium-Ion Battery Pack Solution." Application Note #AN189, February 14 2002.

[Maxim 201]   Maxim/Dallas Semiconductor. "DS2770-based Charger Reference Design." Application Note #AN201, April 15 2002.

[Maxim 344]   Maxim/Dallas Semiconductor. "Input Current Limiting inBattery Chargersfor Mobile Computing Applications." Application Note #AN344, December 2000.

[Maxim 680]   Maxim/Dallas Semiconductor. "How to Design Battery Charger Applications that Require External Microcontrollerss and Related System-Level Issues." Application Note #AN680, March 15 2000.

[Maxim 913]   Maxim/Dallas Semiconductor. "Switch-Mode, Linear, and Pulse Charging Techniques for Li+ Battery in Mobile Phones and PDAs." Application Note #AN913, December 27 2001.

[McIntosh 2003]   K. McIntosh, et al. "The choice of silicon wafer for the production of low-cost rear-contact solar cells." *Proceedings of the 3rd IEEE World Conference on Photovoltaic Energy Conversion*, 11-18 May 2003, vol. 1, pages 971-974.

[McIntosh 2004]   K. McIntosh, et al., June, 2004. "Light Trapping in SunPower's A-300 Cells." <http://www.sunpowercorp.com/html/Technical%20Papers/pdf/mcpaper2.pdf>,

June, 2004.

[Merrigan 1975]   J. Merrigan. *Sunlight to Electricity*. The MIT Press, Cambridge, 1975.

[Midya 1996]   P. Midya, et al. "Dynamic Maximum Power Point Tracker for Photovoltaic Applications." *Proceedings of the 27th IEEE Power Electronics Specialists Conference*, 23-27 June 1996, vol. 2, pgs 1710-1716.

[Mulligan 2004]   W. Mulligan, et al., June 2004. "Manufacture of Solar Cells with 21% Efficiency." <http://www.sunpowercorp.com/html/Technical%20Papers/pdf/bmpaper.pdf> SunPower Corp, Sunnyvale, California, June 2004.

[Mulvaney 2001]   K. Mulvaney. *At the Ends of the Earth: A History of the Polar Regions*. Shearwater Books, Washington, D.C., 2001.

[NRC 2002]   National Resaerch Council. "Executive Summary: The Sun to the Earth and Beyond – A Decadal Strategy in Solar and Space Physics." National Academies Press, Washington, D.C., 2002.

[NREL I]   United States Department of Energy, National Renewable Energies Laboratory, "Solar Spectra: Air Mass 0." ASTM standard #E-490-00. <http://rredc.nrel.gov/solar/spectra/am0/>, accessed March 30, 2005.

[NREL II]   United States Department of Energy, National Renewable Energies Laboratory, "Solar Position Algorithm (SPA)." <http://rredc.nrel.gov/solar/codes_algs/spa/>, accessed September 30, 2004.

[NREL III]   United States Department of Energy, National Renewable Energies Laboratory, "Reference Solar Spectral Irradiance: Air Mass 1.5." ASTM standard # G173-03. <http://rredc.nrel.gov/solar/spectra/am1.5/>, accessed March 30, 2005.

[NREL IV]   United States Department of Energy, National Renewable Energies Laboratory, National Center for Photovoltaics, "High Efficiency Concepts and Concentrators." <http://www.nrel.gov/ncpv/higheff.html>, accessed March 31, 2005.

[Pacheco 2003]   V. M. Pacheco, et al. "Stand Alone Solar Energy Storage System with Maximum Power Point Tracking." *Proceedings of the 18th IEEE Applied Power Electronics Conference*, 9-13 February 2003, vol. 1, pgs 97-102.

[PhotonTek I]   Photon Technologies, Inc., 2002. "Solar Car Racing '01 Solar Panel Construction Techniques." <http://members.aol.com/photontek/photon/encap.html>, accessed June 3, 2004.

[PhotonTek II]   Photon Technologies, Inc., 2002. "Solar Panel Encapsulants & Lamination." <http://members.aol.com/photontek/photon/encap2.html> accessed June 3, 2004.

[Plagge 2005]   A. M. Plagge. In personal correspondence and discussions with the author February 2004 through May 2005.

[Price 2004]   A. D. Price. "Cool Robot: Mechanical Design of a Solar-Powered Antarctic Robot." Honors Thesis. Thayer School of Engineering at Dartmouth College, Hanover, N.H., June 2004. Also in conversations and meetings with the author.

[Rodgers 2005]   T.J. Rodgers. "The Time Has Come for Solar Power." Jones Seminar lecture at the Thayer School of Engineering, Dartmouth College, March 4, 2005.

[Sayigh 1991]   A. A. M. Sayigh. "Chapter 2 – Photovoltaic and Solar Radiation." In *Generating Electricity from the Sun*, Fred Tremble, editor. Pergamon Press, Oxford, 1991.

[Shirbacheh   M. Shirbacheh. "Power and Pyro Subsystems for Mars Pathfinder."
1997]   *Proceedings of the 32nd IEEE Intersociety Energy Conversion and Engineering Conference*, 27 July – 1 August 1997, vol 4., pgs 2231-2236.

[Sinton 2004]   R. Sinton, Sinton Consulting. In correspondence with the author, May 2004.

[Sullivan and   C. Sullivan and M. Powers. "A High Efficiency Maximum Power Point Tracker
Powers 1993]   for Photovoltaic Arrays in a Solar-Powered Racing Vehicle." *Proceedings of the 24th IEEE Power Electronics Specialists Conference*, 20-24 June 1993, pgs 574-580.

[SunPower I]   SunPower Corporation, November 2003. A-300 Solar Cell Datasheet. <http://www.sunpowercorp.com/html/Products/datasheet/A-300/A-300.pdf>, accessed February 20, 2004.

[SunPower II]   SunPower Corporation, June 2004. SPR-210 solar panel datasheet. <http://www.sunpowercorp.com/html/Products/datasheet/Modules/SPR-210.pdf>, accessed August 12, 2004.

[SunPower III]   SunPower Corporation. "Draft – Cell Interconnect Concept." September 19, 2003. Provided from the manufacturer.

[Tan and Tseng   B. L. Tan and K. J. Tseng. "Intelligent and Reliable Power Supply System for
2003]   Small Satellites." *Proceedings of the 25th IEEE International Telecommunications Energy Conference*, 19-23 Oct 2003, vol. 1, pgs 249-255.

[TI 555]   Texas Instruments. "NE555, SA555, SE555 Precision Timers." Datasheet literature number SLFS022E, revision March 2004.

[Tremble I]   F. Tremble. "Chapter 1 – Introduction." In *Generating Electricity from the Sun*, Fred Tremble, editor. Pergamon Press, Oxford, 1991.

[Tremble II]   F. Tremble. "Chapter 3 – Crystalline Silicon Photovoltaic Modules." In *Generating Electricity from the Sun*, Fred Tremble, editor. Pergamon Press, Oxford, 1991.

[Tse 2002]   K. K. Tse, et al. "A Novel Maximum Power Point Tracker for PV Panels Using Switching Frequency Modulation." *IEEE Transactions on Power Electronics*, vol. 17, no. 6 (November 2002), pgs 980-989.

[UBI 314]   Ultralife Batteries, Inc. "Charging a Li-Ion Battery with a Fixed Voltage Source"   FAQ   #314,   http://ultralifebatteries.custhelp.com/cgi-bin/ultralifebatteries.cfg/php/enduser/std_adp.php?p_sid=FQymvAsh&p_lva=&
p_faqid=314&p_created=1099858059&p_sp=cF9zcmNoPTEmcF9ncmlkc29yd
D0mcF9yb3dfY250PTIxJnBfc2VhcmNoX3RleHQ9JnBfc2VhcmNoX3R5cGU9
MyZwX3Byb2RfbHZsMT0xMSZwX2NhdF9sdmwxPX5hbnl_JnBfc29ydF9ieeT
1kZmx0JnBfcGFnZT0x&p_li=>, accessed January 2005.

[UBI 5112]   Ultralife Batteries, Inc. "Safety Precautions for Lithium Ion and Lithium Polymer Batteries." Literature #UBM-5112, revision Dec 2004, <http://www.ultralifebatteries.com/whitepapers/UBM-5112_Li-ion_Li-Poly_Precautions.pdf>, accessed Feb 2005.

[UBI 5133]   Ultralife Batteries, Inc. "Instructions: CH0002 Battery Charger." Literature #UBM-5133,   revision   September   25   2004, <http://www.ultralifebatteries.com/techsheets/UBI-5133_CH0002_Instructions.pdf>, accessed April 2004.

[UBI 5134]   Ultralife Batteries, Inc. "Technical Data: CH0002 Charger." Literature #UBM-5134,   revision   August   11   2004, <http://www.ultralifebatteries.com/techsheets/UBI-5134_CH0002.pdf>, accessed April 2004.

[USNO 2003]   United States Naval Observatory, Astronomical Applications Department, October 2003. "The Seasons and the Earth's Orbit – Milankovitch Cycles." <http://aa.usno.navy.mil/faq/docs/seasons_orbit.html>, accessed April 5, 2005.

[Valenziano and   L. Valenziano and G. Dall'Oglio. "Millemetre Astronomy from the High
Dall'Ogllio   Antarctic Plateau: Site Testing at Dome C." *Publication of the Astronomical*
1999]   *Society of Australia*. 1999, 16, pages 167-174.

[Ware 2000]   R. H. Ware et al. "SuomiNet: A Real-Time National GPS Network for Atmostpheric Research and Education." *Bulletin of the American Meteorological Society*, vol. 81, no. 4 (April 2000), pages 677-694.

[Warren 1998]   S. Warren, R. Brandt, and P Hinton. "Effect of surface roughness on bidirectional reflectance of Antarctic snow." *Journal of Geophysical Research*, vol. 103, no. E11 (October 1998), pages 25,789-25,807.

[Z-World 2004]   Z-World, Inc. *Dynamic C User's Manual*, document number 019-0125-C, revision August 31, 2004. Z-World, Inc., Davis, California, 2004.

[Z-World 2005]   Z-World, Inc. *Rabbit Core RCM3100 User's Manual*, document number 019-0115, revision Jan 31, 2005. Z-World, Inc., Davis, California, 2005.