University of Colorado at Boulder

ROBOTIC AUTONOMOUS TRANSPORT

(RAT)

Authored by: Multidisciplinary Engineering Team
University of Colorado at Boulder

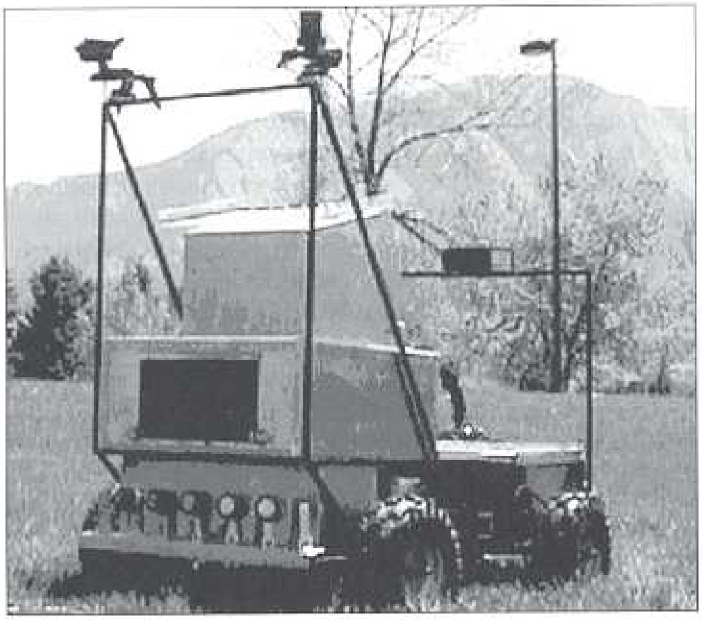

An interdisciplinary engineering team assembled in the Fall of 1997 with the task of designing and building an autonomous robotic ground vehicle to compete in the 1998 International Ground Robotics Vehicle Competition. The design process began by establishing primary design criteria and deriving secondary design requirements. Previous designs from past CU entries were evaluated, and an innovative new design was implemented as discussed in this document. Primary design criteria are shown in Table 1 along with a Criteria No. for referencing throughout this document.

Table 1: Primary Design Criteria and Assigned Subsystem

Primary Criteria	Sub-system responsible	Criteria No.
5-mph capacity	Guidance	G1
Avoid real & interpolated lines	Guidance	G2
Avoid obstacles	Guidance	G3
Detect painted lines	Navigation	N1
Detect obstacles	Navigation	N2
Maneuverability allowing turn radius of 10 ft.	Mobility	M1
Sufficient traction for necessary speed	Mobility	M2
Weatherproofing	Mobility	M3
Size restrictions adjusted to building exit	Mobility	M4
Power to run race course three times	Power	P1
Safety remote & manual stop	Power	P2
120V AC adaptability for testing	Power	P3

Our team then defined subsystem groups to carry out the detailed design. These subgroups were Mobility, Navigation and Guidance. Mobility concerns the physical structure of a vehicle itself as well as its role in propulsion and steering. Navigation determines the vehicle's location relative to its external environment. Guidance decides where to go based on this location information (i.e. turning right or left). Navigation was divided into two subgroups: Vision (responsible for line detection) and Ultrasound (responsible for obstacle detection). A Power subgroup eventually split off from Mobility, focusing on power generation and distribution. Figure 1 shows the team members and leaders of each subsystem. An overview of the system architecture and system interfaces is given below in Section 1, followed by a description of how each of the five subgroups shown in Figure 1 developed a design to meet the defined requirements (Sections 2-6). Finally, Section 7 reviews the integration of all subsystems, covering financial and time expenditures.

System Architecture

The vehicle steers by articulating through a center pivot, has four-wheel drive, two cameras, seven ultrasound sensors and a P166 PC. The system architecture is illustrated in Figure 2 below. Power connections are shown as solid lines, and data connections are shown as dashed lines. Components shown in white indicate their approximate position on the vehicle. The one exception to this is the computer Highlab. The actual position of the computer is shown with a dotted box, and its internal components are shown in a blow-up below the vehicle. Software components are shown in an internal dotted box labeled "Software." Details of this architecture are discussed in later sections of this report.

Figure 1: Organizational Flow Chart

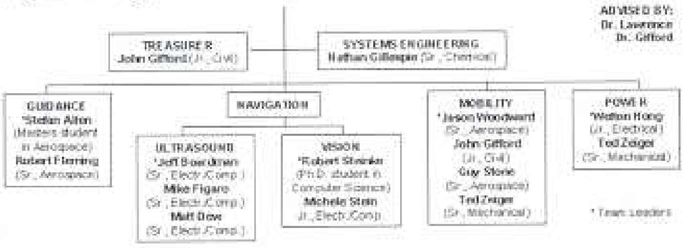

The software and data flow underlying Figure 2 is described in Figure 3. Data flow begins with data acquisition from cameras through frame grabber cards, and from ultrasound sensors through a data multiplexer into the computer's serial port. This data is processed by a navigation subsystem that detects course boundaries and obstacles. This information is passed on to a Guidance subsystem that determines the vehicle's desired speed and steer angle. This data is sent to a PID controller which modulates the voltage applied to the wheel motors until the desired and actual motion data match.

Figure 2: System Architecture

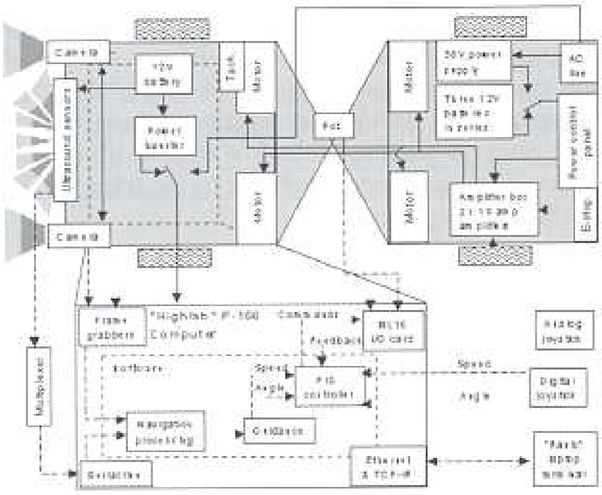

Figure 3; Software Architecture

The Navigation system sends to the Guidance system two kinds of information: boundary data from video cameras, and obstacle data from ultrasonic sensors. Two major design decisions influenced the Navigation to Guidance interface. First, it was necessary to combine Vision and Ultrasound information. Secondly, it was desired that vision data be as condensed as possible before passing the information to Guidance. Therefore, a Navigation software module was created to handle these tasks (see Figure 3). Navigation operates by reading vision data from the two camera frame grabbers and ultrasound data from the serial port to construct a representation of its environment. The Navigation portion of the vehicle control system software processes the images to detect any boundaries within the field of view. The boundary information is then combined with Ultrasound sensor data that detects the presence or lack of obstacles.

The initial design called for asynchronous communication between Navigation and Guidance. The Navigation module would continuously acquire and process new data and the results would be available at any time to the Guidance module through a shared memory structure. Ultimately, an asynchronous system proved inefficient. Vision processing consumes 95% of our CPU time. It made little sense for Guidance to be continuously looking for new data when none was available. To accommodate this dependency, a synchronous system was implemented in which Guidance polls Navigation that in turn polls the devices, processes their data, and sends the results to Guidance. The original data flow architecture was otherwise left unaltered from the design change.

Guidance commands the vehicle with two scalar output values. The first represents the forward/backward speed of the standard reference point, a point midway between the front two wheels. The second is the bend angle between the front and rear vehicle segments. The unique vehicle design requires that these two values be translated into two amplifier control signals. For this, a dedicated closed-loop amplifier control process, the PID controller in Figure 3, runs on the main computer. This process must run at about ten times the frequency of other software subsystems to manage the vehicle's dynamics. Guidance communicates the speed and angle directives to this controller via shared memory.

Guidance subsystem

The responsibility of the Guidance subsystem is to command the vehicle how and when to move, based on sensory input and perhaps a priori knowledge of the competition course. Three fundamental requirements we re imposed by the Official Rules. First, the vehicle must traverse the course from beginning at <5 mph (criteria G1), but in not more than ten minutes. Second, the vehicle must not cross real or interpolated course lines (criteria G2). This meant that Guidance must maintain accurate data on boundary locations. Finally, the vehicle must not displace obstacles (criteria G3). This suggests that Guidance must maintain a safe distance margin from all obstacles the vehicle encounters.

The Navigation subsystem communicates boundary data to Guidance as a series of (x,y) vectors, in a reference frame attached to the vehicle, for every vision cycle, i.e. for each processed frame from the cameras. "Boundary" generically refers to the course edge, regardless of whether it is a painted line or the edge of a sand pit. Each ordered pair, (x,y), corresponds to a point on the ground which is at that instant contained by the course boundary. These data are stored by Guidance in one of two First In/First Out queues (one each for the left and right lines), which hold about ten seconds of data, although this is configurable. As these queues naturally contain "historical" data, the data must be transformed during each Guidance cycle to account for vehicle movement. This transformation is composed of translation and rotation, using bend-angle and vehicle-speed measurements. Both the left and right data sets are processed by a least-squares second-degree polynomial fit algorithm, producing two sets of polynomial coefficients. These coefficients constitute the primary navigational aid for the Guidance subsystem. Using curve fitting in this fashion eliminates dependence on a continuous input stream of boundary data, i.e. navigation is possible even when the course boundary is temporarily lost to the sensors.

Obstacle data is provided by Navigation as an array of seven scalar values, representing the distance from the corresponding ultrasonic sensor to the nearest object in its field of view. Guidance relies on instantaneous obstacle data, and does not track individual obstacles over time. This information is used to identify a path free of obstacles.

In the absence of obstacles, Guidance uses a closed loop to center the vehicle between the two boundaries. This is done by issuing bend-angle commands to steer the vehicle toward a point (nominally) ten feet ahead of the standard reference point (defined above), and midway between the calculated boundaries at that distance. This coordinate is given by Ydesired = .5 • (L(10 ft) + R(10 ft)).

In the presence of obstacles, two distinct approaches have been evaluated. The first uses a priori knowledge, essentially pre-programming the vehicle to use the same technique as with no obstacles, but replacing 0.5 (for centering) in the equation above with some other factor, a, where $0 < a < 1$. Such values would be stored in turn for each scenario on the course, based on the team's qualitative visual observations. This method is increasingly susceptible to synchronization errors as the vehicle progresses through the course, depending on accurate odometry aided by ultrasonic input for resynchronization.

The second approach is more truly autonomous. In this case Guidance must recognize obstacle scenarios as they are encountered and appropriately command the vehicle to circumnavigate them. When an obstacle effectively divides the course, Guidance bisects the largest of the two candidate gaps, i.e. left or right of the obstacle. This latter approach was selected for its proven simplicity and because the former system's strong reliance on synchronization was viewed as a weakness.

Two special cases required careful consideration. For the "passable barricade," as defined in Competition Rules, the simple method of bisecting the largest gap is inadequate. This situation will at first mislead Guidance, but when the more distant obstacles comprising the trap are detected, this scenario is identified, and a backup command course change is initiated. The second special case is the ramp. The ramp is detected by Navigation, which notifies Guidance with a special message type. This event triggers the special ramp mode, during which the vehicle uses dead reckoning at a slower velocity until the ramp is behind it; that is, ultrasonic sensory input is suspended during this time because it is considered unreliable.

Vision

The Vision subsystem was designed to satisfy the key requirement of detecting painted lines (criteria N2 in Table 1). To accomplish this we decided to process images from two high-angle video cameras because it was a simple and proven approach.

For implementing an image capture system three options were considered. These were camcorders with frame grabbers, "videoconferencing" serial port cameras, and "security" cameras with frame grabbers. We finally decided on solid state security cameras because they are smaller, lighter, and less complex than camcorders. As long as power is supplied the security cameras transmit a continuous video signal whereas a camcorder has its own set of controls that must be started independently of the computer program. We decided on security cameras over parallel port cameras because of the low bandwidth of the serial port (2 frames/second) relative to the frame grabber boards (30 frames/second).

The image-processing algorithm consists of two steps. The first decides which pixels in the captured image will be considered lines. Then the positions of those pixels are transformed to a ground coordinate system and accumulated in bins aligned with a grid. If the number of pixels detected in a particular bin exceed a given threshold, then that grid cell is indicated as part of a line. This second operation prevents isolated noise pixels from being detected as lines.

Two difficult problems were faced while implementing this system. The first was the determination of which pixels should be considered lines, and the second was the calibration of camera orientation relative to the vehicle.

Detection of line pixels is performed by quantitative manipulation of the image, usually a convolution with some detection mask followed by a threshold on the resulting values. For any particular situation it is relatively easy to develop an algorithm that will confidently detect line pixels. The difficulty arises in developing an algorithm that will provide good detection in a variety of situations. Some of these situations can be compensated for before each run, for example using different threshold values for a sunny day versus a cloudy day. However, some

continued on page 28

continued from page 14

Team Name	**Nationality**	**Accrued Qualifying Points**
1. Simon Fraser University	CANADA	445
2. Georgia Institute of Technology	USA	343
3. DeVry-Calgary	CANADA	374
4. Massachusetts Institute of Technology	USA	377.75
5. Oakland University	USA	0
6. Southern Polytechnic State University	USA	387
7. Rose-Hulman Institute of Technology	USA	393
8. University of California, Berkeley	USA	380.25
9. Technische Universitaet Berlin	GERMANY	337.75
10. University of Waterloo	CANADA	403.75
11. University of California-San Diego	USA	326.25
12. University of Central Florida	USA	0
13. University of British Colombia	CANADA	344.5
14. Mesa State College	USA	262.75
15. Instituto Tecnologico y de Estudios Superiores de Monterrey, Guadalajara	MEXICO	0

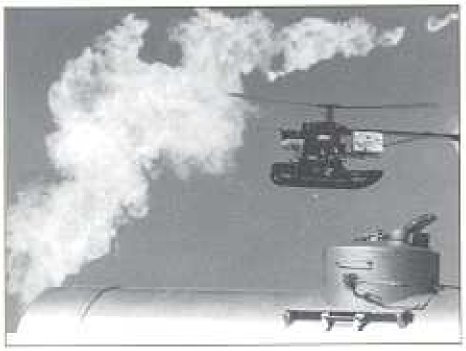

The event was cosponsored by the Association for Unmanned Vehicle Systems International (AUVSI) and the U.S. Department of Energy (DOE). The DOE provided the HAMMER facility and logistics funds to put on the event and pay for the operation of the disaster props. AUVSI is setting aside prize money each year until AD 2000 to create a purse that will be in excess of $30,000 to the winning team. Special technical support is being provided by NovAtel, a Canadian GPS manufacturer, which is supplying all official teams with the free loan of approximately $30,000 in differential GPS equipment.

For further information, consult the world wide web at: http://avdil.gtri.gatech.edu/AUVS/IARCLaunchPoint.html

Here the official rules and Millennial Event description can be found along with links to Frequently Asked Questions, pictures from not only the 1998 qualifier, but also from the past several competitions, and updates regarding the upcoming event.

Colorado
continued from page 26

Figure 5: Images Captured by Cameras (Top) and After Line Detection Processing (Bottom)

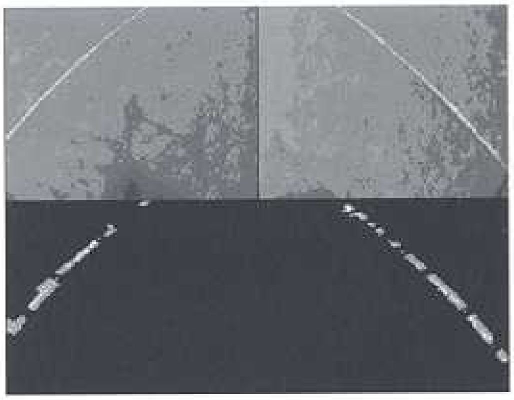

situations cannot be defined before a race, such as the direction of the sun relative to the vehicle that changes as the vehicle goes around the course.

An example of line detection is shown in Figure 5. The top two pictures are the actual images sent by the cameras. The bottom images show the areas detected as lines highlighted in the image plane, and the rest of the image is blacked out.

The second main problem to overcome was one of calibration, or how to correlate a pixel in the image plane with a point on the ground. This consists of two parts: doing a mathematical transformation with assumptions about the camera's field of view, position, orientation, etc., and also verifying those assumptions through measurement. It was found to be too time-consuming to point the camera arbitrarily and then try to calculate pixel positions on the ground. Therefore, it was decided to fix the camera's field of view on the ground, and then point the cameras so that a predefined field of view is imaged in its entirety. A program was developed in MATLAB to perform calibration calculations and create files to transform pixel positions to a ground coordinate system. A template was made of string to be used during calibration by placing on the ground in front of the vehicle, thereby centering each camera's field of view.

Ultrasound

The key requirement for Ultrasound was to provide reliable obstacle detection (criteria N2 in Table 1). Additional design criteria were chosen by Ultrasound to provide the most reliable sub-system to the RAT team. One of these secondary design criteria was that the Ultrasound electronics would be able to read any information from any sensor that adheres to the RS232 serial protocol. This allows for a variety of sensors to be used with the same hardware interface to the "Highlab" PC.

Another secondary design criteria was the ability to add or remove sensors as needed without major changes to the hardware or software. This proved particularly useful when the number of sensors was increased during the design phase from five to seven. The hardware and software allowed for this change to occur without any major difficulties.

The Ultrasound sensor system is very simple. Each sensor is selected by the computer's digital I/O card. The output from the multiplexer is then read in by the serial port of Highlab. The software then determines if the information is valid, or within the desired range, and then selects the next sensor, repeating the process.

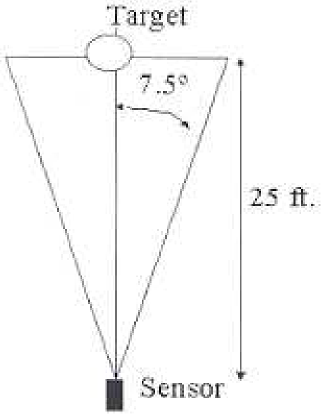

Figure 6: Ultrasound Observation Cone

The sensors that have been used are Senix OEM ultrasound sensors. Five of these were donated to the team by Senix and the additional two sensors were purchased at a discounted rate. The sensors have a range of 5 inches to 25 feet at a resolution of 0.01". The observation cone is 15° total or 7.5° from the center of the cone, shown in Figure 6.

The Ultrasound design changed over the course of the last month to become more reliable and durable. The old design used a Motorola HC11 microcontroller to control the multiplexer. This design required 24 electronic chips and over 100 connections. Such complexity significantly detracted from its reliability and ease of troubleshooting. The number of connections also made it considerably less durable. In contrast, the new design, including the computer-controlled data multiplexer , uses only five chips and only 31 wire connections, making for a cleaner, more rugged design.

Mobility

Four key requirements were specified for the Mobility subsystem. The maneuverability of the system was required to have a turning radius of 10 ft (criteria M1), allowing the vehicle to navigate a greater range of obstacle situat ions. It was also necessary to have sufficient traction no matter what terrain was being traversed, particularly in traversing the sand pit (criteria M2). Another key requirement for both the competition as well as during testing was to have the vehicle weatherproofed (criteria M3) giving sufficient time to get the vehicle back into shelter in the case of rain. The last key criteria for Mobility was to consider where construction and testing would take place, since the vehicle needed to fit through the elevator and exit doors of the building (criteria M4).

The RAT vehicle is an articulated vehicle which uses four electric motors for both propulsion and steering. PID controllers control both the velocity and the direction of the vehicle by sending torque commands to the motors. As shown in Figure 2, the cross-coupled pairs of motors are connected in series with each other and each pair is powered by a single amplifier. The response of the vehicle to commands is measured by a potentiometer (bend angle) and tachometer (forward speed) which feeds this information back to the PID controllers. This articulated vehicle design was based mainly on a desire to develop a highly maneuverable vehicle using as simple a design as possible. As seen in past rovers, such as CU Boulder's Robo car from last year's competition, Ackerman steering is a popular option, but it does not have the maneuvering capabilities of an articulated design such as the one used on the Russian Mars98 rover. One question posed early on in the design process was how to control an articulated design using only motors with no center actuator. This was answered by building a small LEGO model with cross-coupled motors which successfully proved that such a design could work.

Figure 7: Control of Articulated Vehicle

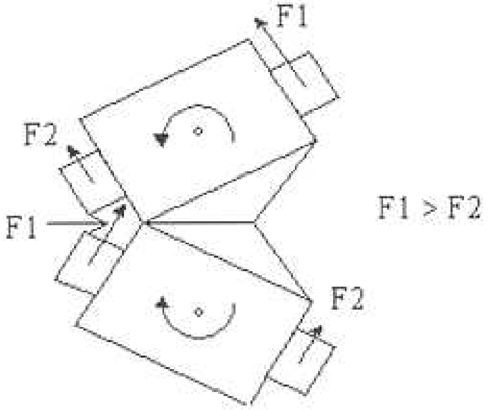

Figure 7 shows how the control arrangement operates. If force command F1 is higher than the command F2, these forces will generate moments around the center points of each body section as shown. These moments act about the center pivot point causing the vehicle to bend and thus turn. Keeping the commands F1 and F2 equal will keep the vehicle from turning, propelling the vehicle if the command is large enough to overcome rolling friction.

The PID controllers implemented in the final design were developed using a computer model implemented in MATLAB's Simulink. This model aided in setting the controller gains, and in understanding the dynamics of the model. It was also used as a test bed for adding items such as low pass filters to the control software to reduce response to noise. The PID controller stands alone as a separate own software module, and also accepts commands from other modules to include the navigation code and the joystick code for manual driving.

One major problem found during development of the PID controller was that the vehicle could be hard to control even with high derivative gains because of its free center pivot. This problem later displayed itself as the vehicle hardware was built and tested and was solved by adding dampers to the center pivot. The AVM Corp. dampers chosen to solve this problem were sized based upon the turn rate desired to allow the vehicle, while traveling at 5 mph, to miss an obstacle detected from 9' away. Adding these dampers greatly increased the ease of driving of the vehicle and helped to make the control system more robust.

Dayton 1Z833 electric motors were used to propel and steer the vehicle. The sizing for these motors was based upon the self-imposed requirement of driving a 250lb vehicle up a 15% slope at 5 mph. Rolling-friction was also incorporated into the sizing o f the motors, but was later determined to be too small an estimate. Adding to this problem was an overweight vehicle relative to initial weight projections. It was determined that the motors could provide extra torque to help overcome these problems, but the amplifiers could not provide this extra power. Two options could fix this problem; one was to increase the gear ratio and add a battery, and the other was to acquire more capable amplifiers. The former option implied changing all of our drive electronics to deal with increased voltage from an extra battery and was therefore discarded in favor of switching out the amplifiers. The Dayton motors were ultimately selected for not only their power capabilities, but also for their all weather casing and industrially proven design. Advanced Motion Controls 50A8 PWM amplifiers replaced the 25A8 amplifiers to power the motors.

Figure 8: Current and Voltage needs at 5 mph up a 15% slope at 328 lbs

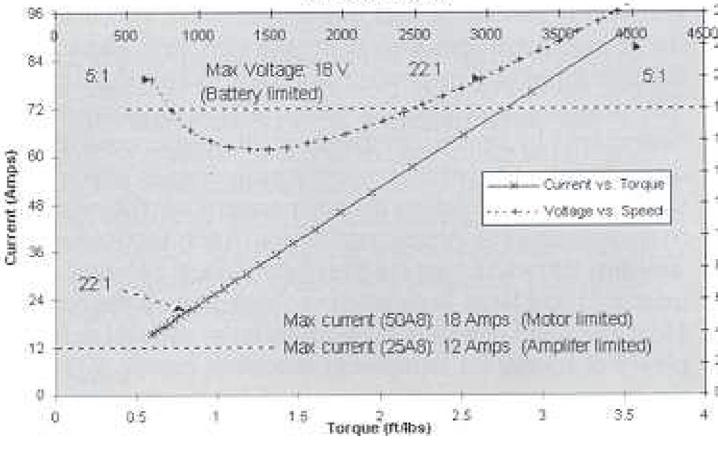

Figure 8 shows the relationship between the voltage and current needs of the vehicle with respect to the torque and speed provided by the Dayton 1Z833 gearmotors. Each tick on each of the lines represents gear ratios from 5:1 up to 30:1. Each 5:1 gear ratio is marked as well as the gear ratio of 22:1 implemented on the RAT.

The drive system imposed stringent requirements related to weight distribution. In order for our steering system to work, the traction on all four wheels needs to be approximately equal. This was accomplished primarily by the free roll motion in the center pivot, which keeps all four wheels in contact with the ground. In addition, weight was evenly distributed throughout the vehicle so that no one wheel has significantly more traction than another. Also, the center of gravity in each of the vehicle's two segments had to be kept over the axle as closely as possible in order to minimize any excess torque on the center pivot.

Another major concern in the placement of vehicle components was their easy accessibility without interfering with access to other system components. A prime example of this would be the placement of the laptop, which is used for interface with the vehicle's computer. It was decided that this should be placed in such a way as to enable someone to use it without interfering with the cameras' or the ultrasonic sensors' field of views. Each of the vehicle's systems also had to be easily accessible in order to facilitate working on them. The placement of the vehicle's systems was also designed to minimize the complexity of the wiring harness, and to reduce the number of wires that cross through the center pivot. Most of the information processing components of the vehicle were therefore placed in the front segment, while most of the drive components were placed in the rear of the vehicle. Sensitive components had to be isolated as well as possible from vibrations. Those components most susceptible to damage from impacts and vibrations were identified in order to incorporate dampening into our design.

To ensure safety and reliability of all components, they had to be weatherproof. Therefore, each of the components was packaged in such a way as to minimize the risk that it could be damaged by the elements.

During the process of selecting the tires and wheels for the vehicle there were two key issues that were addressed: the mounting of the sprocket and the tread on the tire. The tires had to provide excellent traction capabilities determined by the requirements of the torque control drive design. The initial brainstorming designs provided the desired size based on the requirement for the vehicle to maneuver a complete circle with in a ten foot radius. The tire chosen is used on heavy-duty snow blowing machines and large lawn tractors. Nick-named the Snow Hog, the tread design is for non-highway use and delivers plenty of torque for equipment that must handle well in the snow or on undulating terrain. A key feature in the Snow Hog tire and wheel were the press fitted 5/8" sealed bearings. The sealed bearings insured that debris would not enter the bearings providing excellent reliability. A 3/4" solid axle was selected to ensure the weight of the vehicle and its components would not produce any noticeable deflection. The bearings simplified the axle assembly and allowed the axle to be welded directly to the frame. Each end of the 3/4" axles was machined for a minimum clearance fit enabling a simple and effective mounting of the wheels. The mounting of the sprocket to the wheel required a 1" aluminum spacer be manufactured in house, which was mounted to the rim, and the drive sprocket was mounted to its opposite side (see Figure 9).

Figure 9: Sprocket Design

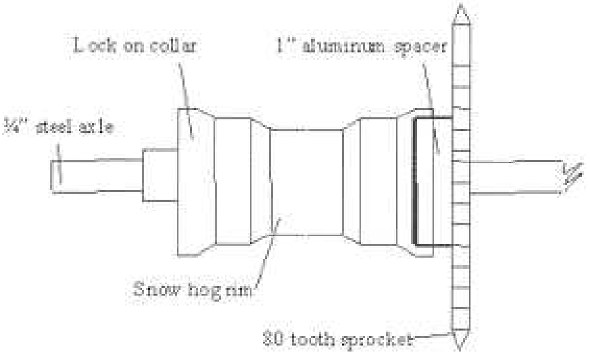

POWER

Three key requirements were identified for the vehicle's power sub-system. The most prominent was the ability to provide enough power to navigate the competition course three times (criteria P1). A safety concern was the ability to stop the vehicle using either the radio-controlled emergency stop (E-stop) or the manual button located directly on the vehicle (criteria P2). For flexibility in testing, it was also decided that the vehicle should allow for conventional 120V AC power source (criteria P3), permitting extended testing without having to recharge the batteries.

The primary function of the power system is to provide electrical power to all vehicle systems. There must be enough onboard storage to power the vehicle for the duration of the contest including practice and testing times. To this end the onboard power is sufficient to run the vehicle for at least 1 hour, plenty of time to make three laps around the course.

The derivation of the exact storage capacity needed for the vehicle to perform had several stages. First the power consumption by the RAT's four motors was estimated (see Figure 8) needed to be found. Second, the current draw of the onboard electronics, including the main computer, sensors and cameras, was measured. Lastly, the weight and size of available batteries needed to be considered. At this point, a perplexing problem was encountered: how to provide enough battery storage without causing the vehicle to be grossly overweight.

Gel lead acid batteries were the optimal solution. They offered good storage capacity while retaining small size and weight. Furthermore their sealed design gave additional safety.

It became evident that the original design for a 12-volt system would not meet the needs of the mobility system. For the RAT to reach five mph, a higher voltage was needed to overcome the back EMF of the motor. Thus, two separate power systems were put into place: a 12-volt system for powering the computer, sensors and cameras, and a 36-volt system for the motors. A single battery supplies the 12-volt power, while three 12-volt batteries in series supply the 36-volt power. Figures 10 and 11 below show the rear power layout and front power layout.

Figure 10: Rear Power Layout

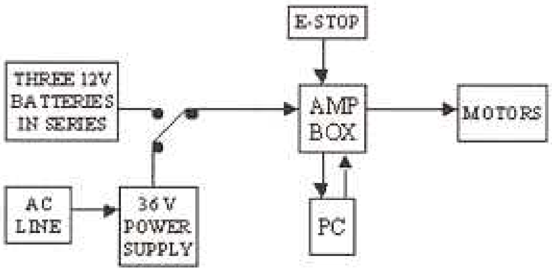

Figure 11: Front Power Layout

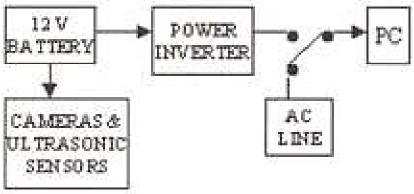

For convenience purposes, the RAT can be run from onboard battery DC power or external AC power. A 36V DC power supply was built to replace the three 12-volt series batteries powering the motors. To take the main load off the single 12V battery, the power inverter can be bypassed and the computer supplied with AC line power.

Safety is one of the main design issues, thus the RAT can be stopped by either a radio emergency stop control or an emergency stop button located in the rear of the RAT. These two switches are wired in series, so if one of the two buttons is engaged, the vehicle will be halted. The principle behind the E-stop is first to inhibit the motor amplifiers and second to short the motor windings. Inhibiting the amplifiers cuts the power to the motors and shorting the windings uses the back EMF to provide breaking force. A standard two-channel remote control was adapted to be used for the radio E-stop. The pulse width modulated (PWM) signal (used to control servomotors) is filtered to provide a DC voltage. This voltage is then used to control a relay via an op-amp and power transistor.

Providing a failsafe emergency stop system and proper shielding of sensitive electronics were the hardest tasks. Ensuring that the vehicle could be stopped in the event of computer failure was a must. A separate system of amplifier control is coupled directly to the E-stop switches. The radio E-stop only allows the vehicle to run only when it is receiving a good signal. If the vehicle goes out of range, it is shut down automatically. With two PWM amplifiers emitting high amounts of RF, in conjunction with analog signal lines susceptible to small amounts of noise, shielding of all lines was necessary. A design was developed that provided electrical shielding, yet removed amplifier heat to the vehicle exterior.

Table 2: Major Resource Distribution

Resources Summary		
Group	Material Cost	Labor Hours
Mobility	$ 1,875.00	400
Power	$ 883.00	300
Vision	$ 300.00	350
Ultrasound	$ 200.00	225
Guidance	$ -	400
Misc.	796	100
Totals	**$ 4,054.00**	**1,775**

System Integration

Major milestones included having a functional chassis by early in Spring 1998 and an operational vehicle ready for testing by the middle of the same semester. Both milestones were reached by the projected deadlines. Subsequently, though, computer hardware problems, including a faulty motherboard, caused significant delay. Having sufficient buffer time in the project timeline then became very significant. The aggressive scheduling strategy also proved useful when difficulties arose in power and control systems. Difficulties in power systems arose in creating a failsafe E-stop, and in shielding sensitive electronic components. Despite these and other obstacles all major milestones were met within two weeks of their original date.

Table 2 illustrates the allocation of the major resources - money and labor — in this project. Of special note is the large portion of resources that went into power and mobility. In addition, much of the computer equipment was already available from previous competition teams, so this cost was not included, but is conservatively estimated at $2500. Money to pay for the vehicle material costs was generously donated by OmniTech Robotics, located in Englewood, Colorado.

Conclusion

This project was unique for all of us involved in the CU team. None of us had ever worked in a project of this magnitude in a multidisciplinary (6 engineering departments) environment before. This project was started essentially from scratch and through the time period of one academic year was designed, implemented and tested in preparation for the competition.