Design and Implementation of an Autonomous Underwater Vehicle for the 2003 AUVSI Underwater Competition

CUAUV Team
Cornell University
Ithaca, NY 14853

Abstract

The design of a fully autonomous underwater vehicle capable of completing the mission parameters set down by the AUVSI organizers is described. Details of the mechanical subsystems, power and propulsion infrastructure, computing resources and structure, and signal acquisition and processing are described.

Introduction

The 2003 AUVSI underwater competition sets down a series of tasks which competitors are to complete during the course of the competition. Competing vehicles must first traverse a validation gate. Vehicles that fail to traverse the gate receive no points for any future accomplishments. After the gate has been successfully passed, the vehicle must then find the group of arrows located on the bottom and select the one arrow with an active indicator light. The arrow with the active LED array serves as a directional bearing, and its position relative to the other arrows specifies a frequency for an acoustic pinger located at the target.

In the final phase of the mission, competing vehicles are to navigate to a target comprised of three concentric boxes. Up to two markers may be dropped onto the target, with the inner boxes valued at more points than the larger outer boxes. Vehicles must complete the entire mission underwater.

The competition is held in the Transdec arena. Transdec is a US navy test pool, and is comprised of an outer shallow ring 16 feet in depth, and an inner deep bowl that reaches 38 feet in depth. The deeper bowl is not used in this year's competition, and three target arrays are provided in the competition arena, with a fourth target on the other side of the arena for testing and practice.

System Overview

Cornell's 2003 entry is a ground-up re-design based on the 2002 design, with much of the general structure retained. A two-hull solution was retained, with each hull separately mounted to an aluminum exoskeleton that forms the backbone of the submarine. Four custom thrusters are mounted, two mounted vertically for depth control, and two mounted just behind the centerline that provide for both forward motion and yaw control. Environmental sensors also mount to the frame, with all electrical connections routed using sealed submersible connectors.

The larger upper hull contains the main computer systems and support equipment, providing the processing power necessary for autonomous operation. The lower hull contains the batteries and power electronics.

Mechanical Structure and Design

The mechanical design looks to satisfy three major goals. The first is to provide a stable vehicle with a range of motion suited to the tasks required of it. The second is to provide for dry space for components where needed within the vehicle, and provide sealed high-quality electrical connections to external sensors. Third, the system must mount and accommodate the range of sensors required by the onboard Artificial Intelligence (AI) to complete the mission. Finally, where possible, we would like to minimize the weight and volume of the vehicle as much as possible.

Mechanics – Propulsion

Four custom-designed and custom machined thruster units achieve the propulsion for this year's entry. Each unit mounts a single brushless motor within an aluminum housing. The rear of the unit mounts a single three-pin Impulse brand connector designed for submersible use, and a dual O-ring seal ensures that the connector end is watertight. The thruster shaft mounts to a propeller designed specifically for use in water through a rotating shaft seal, allowing a watertight rotating connection. An aluminum shroud mounts to the thruster body to protect against accidental contact with the prop. Two threaded holes on the thruster wall allow for easy mounting with standard hardware.

Independent motor controllers mounted in the lower hull assembly drive each of the four thrusters. Since the brushless motor controllers we selected were not reversible, a relay is present for each thruster that serves to change the direction of the submarine when required. Both the relays and the motor controllers are controlled over a single serial link from the upper hull. This division means that only a low-power serial link connects is needed for motor control, reducing the electrical noise induced by the thrusters in the upper hull.

Hull Construction

In an effort to save weight, the pressure hull wall has moved from a 6" PVC construction to a slightly smaller carbon fiber material. This provides us with a much stronger hull at a significantly reduced weight. The pressure seal is achieved using an edge crush seal similar to previous years, with a ring of mounting bolts providing the necessary torque. This has the dual advantage of both being a simple and robust technique for relatively rough and tumble use as well as providing a better seal as the depth increases. A testament to this method is the fact that our 2001 entry spent over a month on the bottom of the arena without any noticeable leakage.

Both pressure hulls this year feature a blind-mount interface system. The connector end of the hull is semi-permanently mounted, avoiding the need to disconnect all of the submersible connectors in order to access the internal racks. Rather, each hull has an internal tray with a blind mount board that matches the board on the connector end of the hull. The rack may be inserted and seated from the far end while the connector end is sealed safely. The blind mount plates are designed specifically for each hull, and it is physically impossible by design to connect the blind-mounts incorrectly.

Lower Hull – Power

Power for this year's vehicle is provided by three battery packs mounted in the lower hull. Each battery pack consists of seven Yardney 9Ah Li-Ion cells, for a 28V pack capable of supplying about 200Wh each (80% discharged). Two of these cells operate in parallel to supply the thrusters, and a single pack provides an isolated power to the upper hull. Monitoring circuitry is mounted for each cell, and is combined using a wired-or circuit on the top board to provide two lines that warn of a low or empty battery pack.

These battery packs represent approximately an 8x improvement over the lead-acid cells used in previous designs, offering twice the power at a quarter the weight. They also required the development of a custom charging solution to support the Li-Ion chemistry. Despite the added complexity, the batteries have performed well, giving over three hours of operational power in testing.

Placing the batteries in the lower hull gives the additional benefit of a very stable dynamic system. This is a direct result of the higher mass of the batteries in a smaller hull. Thus, the buoyancy of the upper hull is significantly greater than that of the lower hull. This configuration results in a strong restoring moment to roll and pitch perturbations, corresponding to a very stable design.

Sensor Suite

The primary tasks of this year's competition are to pass through a validation gate, detect and identify the presence of a light beacon on one of three arrows, and use this information to navigate to a target. The targets also posses an acoustic pinger, with each target assembly operating on a unique frequency. Once we reach the target, we are to drop one or two markers onto the target and surface the vehicle.

In order to accomplish this goal, we have selected a suite of sensors to provide sufficient environmental data to the AI subsystem to achieve the mission goals. The first of these sensors is the DVL, or Doppler Velocity Log sonar. Our DVL (an RDI systems unit) uses a correlation analysis of four independent sonar beams at known angles to calculate velocity, and thus gives both velocity and distance measures. In addition, our unit contains a magnetic compass for heading data, and was refitted to include a pressure sensor for depth measurements. The DVL feedback forms the core of the sensors responsible for the PID control system (which will be discussed in more detail in the software section).

The second major sensor system are two underwater cameras, one forward looking and one directed downward, both mounted on the bow of the craft. The forward looking camera provides for intermediate to short range navigation to the target, and the downward looking cameras are used to detect the target on the marker drop run.

Two sonar altimeters have been included, one downward looking along with the camera for accurate target detection, and one forward looking altimeter for wall detection and avoidance.

The final major sensor system is the hydrophone array. There are four mounted hydrophone assemblies at the far corners of the frame. Each assembly consists of a hydrophone, an analog front end consisting of a variable gain amplifier and an analogue to digital converter (ADC), and a Texas Instruments Digital Signal Processor (DSP). The ADC samples are synchronized from a single centralized clock, and are fed to the DSP, where filtering algorithms are applied to eliminate out of band noise and select the specific pass-band of interest. Once the basic filtering is done, the data is streamed back to the upper hull over a USB link.

Upper Hull

The upper hull houses the main processing power for the submarine in two independent PC-104 computer stacks, as well as the necessary support equipment to allow them to communicate and interface properly with the rest of the vehicle. Each system operates from an independent supply, so that any failure modes will be limited to a single system. While this was a design concern, no such problems have been noticed thus far in the testing process.

The two-computer solution was prompted by experiences in previous years with vision processing. Effective visual detection depends on the capture and processing of as many frames as possible by the host computer. Unfortunately, operating a computer in this fashion almost universally degrades performance for other applications on the system. While some sensors, such as low-bandwidth serial links are not overly degraded, others, including the hydrophones, operate at higher speeds and require relatively complex analysis of their own. While not as intensive as vision, the other necessary components of the system do not coexist well with such a processor intensive application.

Because of this, the decision was made to split the processing across two systems. One system is relatively slow and lower power, but is more than suitable for the lower-speed processing requirements. The second is a faster system dedicated to vision processing. This division of labor has helped to achieve good system performance and response times without compromising mission critical cababilities.

The first computer is a Technoland 800MHz crusoe that serves as the main I/O and control processor for the submarine. This system is responsible for handling sensor inputs other than the cameras, acting as a centralized, network-accessible repository of system state, and running the control programs (such as PID) necessary for proper operation.

The second computer is a Pentium III based Jaguar platform whose sole purpose is to capture and process video information from the cameras as quickly as possible. It is configured with a Sensoray frame grabber (based on the BT848/878 chipset) capable of accepting two NTSC video inputs and converting them into YUY2 or RGB images suitable for computer analysis. Captured video is processed entirely on the Jaguar, with the only required communications achieved through a state synchronization protocol.

These two systems are connected through an internal Ethernet switch, which also allows for a direct connection through an Ethernet tether for testing and development. The systems are mounted to a custom PC-104 mounting stack that both mounts the stack in a modular fashion to the tray as well as providing for easy removal and replacement of hard drives.

Signals from the systems are routed from custom cable sets through a wiring conduit located below the main tray. Mission specific signals are routed to the blind-mount connector on the rear of the tray. Debugging and configuration switches, including power switches, Keyboard/Mouse/Video (KVM) signals, and an Ethernet link are routed to the front panel, with a 50-pin connector providing access to the KVM outputs from both computers for debugging access.

Software

Low-level structure

Both of our computer systems are configured to run on GenToo Linux. On top of this we run a hybrid system, comprised of a replicated shared state server and an RPC backbone for state synchronization. This guarantees that any system has a local copy of variables that are accessible through a fast shared-memory mechanism, while also allowing for networked operation to support multiple computer systems. While this does potentially introduce some concerns with system response times and the effects of divergent state on different systems, the current breakdown of operations is designed to minimize such problems.

The structure of the system drivers, conversion processes, and other functions is a good illustration of the benefits of this system. A driver interface may be produced as a simple process stub that reads from the interface in question and writes to the shared memory variables, or that reads from shared memory and writes to a device interface. One may also have a process that reads from shared memory and writes back to shared memory.

The attractive feature of this model is that the shared memory system contains all of the important system state information. Thus, the driver stubs may be modified, re-compiled, and re-started dynamically. This feature has greatly simplified software development and fine tuning for the system.

Hydrophone Processing

The hydrophone systems discussed earlier stream data back to the computers over a USB link through an externally mounted USB hub. Once that data is received, the computer performs a cross-correlation analysis of the data from all four hydrophones. This gives a relative delay matrix for the system, from which one can extract the angular bearing of the acoustic signal.

Factors in the ultimate heading accuracy are the synchronization accuracy of the clock signals, the quantization and sampling error of the ADC, the delay between hydrophones, and the band-select filter performance. Careful design and DSP selection helps to keep filtering and quantization noise to a minimum, and careful routing of signals from the source in the upper hull is necessary to achieve good synchronization accuracy. To minimize further the effects of system noise, the hydrophones are placed at the four corners of the submarine, giving the maximum possible separation. Since the speed of sound is far less than the propagation velocity in our cables, this results in an overall better performance despite the increased routing requirements.

The only caveat is that the hydrophones will provide a bearing to the acoustic source, and in the case that the vehicle is sufficiently close that the source and the target are not aligned on the same bearing the system software must be sophisticated enough to detect this condition. Our design assumption has been that the acoustic pinger will be sufficiently close to the actual target array that visual feedback will be a useful mechanism.

Motor Control

The physical motor control mentioned earlier relies on four brushless motor controllers controlled via an SSC. The SSC accepts from the computer a byte-valued setting that is used to generate the control signals for the motor controllers, accepting up to 255 separate channels if necessary.

The motor control software system is performed in a layered approach in the vehicle to build flexibility and utility. At the lowest level, a program stub reads from the shared memory system a byte value that is written to each specific channel (one byte per controller).

A second program is responsible for translating the non-linear response of the thrusters into a more linear version suitable for a PID control algorithm. This is done by creating a percentile scale in which 100% is the maximum operational thrust, 1% is enough to just turn the thrusters, and 0% is full off. Such a translation helps by eliminating the thruster dead zone (the area in which the applied power is insufficient to create a meaningful response from the thrusters) and decoupling the high-level algorithms from the specifics of the underlying hardware.

A PID (Proportional Integral Derivative) mechanism operates at the highest level to achieve controlled operation. The PID system is actually a combination of three separate PID systems, one for depth, one for velocity, and one for heading. They all operate on a standard PID algorithm based on empirically derived constants, and output a thrust value for each thruster in the percentile scale.

The depth PID operates independently, as it is the only bit of software responsible for control of the fore and aft thrusters. The heading and velocity PID systems both operate on the forward thrusters, and thus must control the same hardware cooperatively. This is achieved by running the velocity PID as a common mode thruster value, and the heading PID in differential mode. The net thruster settings are simply the sum of the two outputs.

Vision processing

There are three major vision problems that the vision processing subsystem seeks to address. The first is determining the active indicator arrow.

Detecting an active indicator arrow implies the ability to detect the cyan indicator array. First selecting patches of cyan pixels, and then analyzing the patches to locate the most circular ones in the system gives a pixel-based locating for the indicator light.

One an indicator light is determined, a Gaussian blur is applied to reduce image noise. An edge detection algorithm utilizing a Hough transform is then applied to extract the most likely line vectors from the image. If the image is well formed, there should be a single line that contains the cyan lights detected in the first phase. This line containing the cyan light indicates the heading to the target, which is extracted and passed to the navigation system.

The same edge/line detection scheme is re-used during the target drop phase. Once the lines have been extracted, the software extracts four lines that form a square or near-square (accounting for the fact that we may be at a slight angle to the target). It is also possible that concentric squares may be present, in which case we proceed to place the submarine directly over the smallest square detected.

Visual detection of target arrays and gate edges is still a consideration, but is a work in progress.

High Level AI

At the highest level, the operational AI is simply a state machine, with one state for each major mission task. Once a task is completed, we proceed to the next state.

The initial state is very simple, consisting of a dead reckoning through the validation gate. Once we pass through the gate, we begin a search phase to locate the indicator arrows and search for the active light in the array.

Upon locating the indicator arrows and detecting the active light, we extract a heading from the arrow in question, and also use this heading to determine the frequency of interest for the hydrophone array.

Initial navigation to the target is accomplished through the extracted heading. During the navigation phase, additional feedback is also provided by the hydrophone array and the vision processing subsystem. This allows for a robust mechanism that makes use of the acoustic pingers at longer ranges, but switches to a visual system as we approach the target assembly for greater precision.

During the final phase, we use the two cameras to navigate the submarine directly over the target array. The downward firing altimeter is also employed for accurate depth determination, as the beam spread on the DVL makes it possible to miss the target when it is right beneath the submarine. Once we are aligned with the target, a bank of electromagnets drops two markers. The submarine then surfaces, terminating the mission.

The AI structure at the highest level is relatively straightforward, as the shared memory structure allows for the independent development and testing of the various software components necessary, allowing the high-level AI to merely start and coordinate the various components as necessary in any given phase.

Conclusions

The overall structure of Cornell’s 2003 entry is based upon the highly successful 2002 design. The dual hull exoskeleton design has been retained, and the Impulse connectors used in previous years have been retained. The computing infrastructure has been significantly upgraded, integrating two independent computer systems with a fast Ethernet interconnection. Power and propulsion has also been improved, with a better battery chemistry providing a much higher energy density and new brushless thrusters providing a maximum thrust far in excess of that in previous years.

Cornell’s entry provides a robust operating platform that is not only useful for this year’s competition, but is also a useful platform for general AUV research and development. Our hope is that the current platform will remain a useful operational platform for several years to come.