



## **Design and Implementation of an Autonomous Underwater Vehicle for the 2004 AUVSI Underwater Competition**

### *Abstract*

*Cornell's entry to the 2004 AUVSI underwater competition is presented. The basic vehicle remains based on the successful platform from the previous year. Design effort has focused on improving and maturing the hardware and software infrastructure in order to add additional modularity and flexibility to the vehicle. Standardization of external connector pinouts and design of an integrated rack/backplane system greatly enhances the flexibility of the platform for both external and internal peripherals. The addition of an Ethernet enabled microcontroller module and development of a multicast Ethernet protocol for state synchronization allows the vehicle to function as a highly distributed system, with high robustness to single point failures.*

### **CUAUV 2004 Members**

116 Phillips Hall  
Ithaca, NY 14853

*Rives Borland, Eli Brown, James Buescher, Chris Caufield, Wei Min Chan, Shawn Chen, Chad Dize, Ben Evans, Marcelo Garza, Vikash Goel, David Hinkes, Kevin Kornegay (Team Advisor), Jayant Kulkarni, Adrian Lai, Chris Lehman, Dennis Leung, Dominick Mancuso, Onimisi Ojeba, Sheng Wuey Ong, Jason Pagnotta, Ming Qiu, Neil Radia, Matt Robins, Karl Schulze, Alex Shih, Phillip Sieh, Bryan Silverthorn, Mike Stanish, Ryan Stenson (Team Leader), Ian Wang, Sean Welch*

## Introduction

The 2004 AUVSI underwater competition sets down a series of tasks for participants to complete during the course of the competition. Competing vehicles must first traverse a validation gate. Vehicles that fail to pass through the gate receive no points for any subsequent mission goals. After the gate, the vehicle must locate a row of bins, marked by a cyan LED. The row of bins will be on an incline, with higher point values for higher altitude bins.

Vehicles may drop up to two markers onto the bin setup, receiving points for dropping markers in the vicinity of the bin setup, and point corresponding to the bin that the marker lands in. In the final phase of the mission, vehicles must find and surface in a recovery zone. The center of the recovery zone is marked with an acoustic pinger. Vehicles must complete the entire mission while fully submerged.

The competition is held in the Transdec arena at the Space and Naval Warfare Systems Center (SSC) in San Diego, California. 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.

## System Overview

Cornell's 2004 entry is a refinement of the 2003 design, retaining much of the external hull and propulsion systems, while focusing design effort on increasing the modularity and flexibility of the vehicle. The 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. The external aluminum frame and all other external aluminum components are hard anodized to protect against corrosion when the vehicle is operating in harsh environments. A CAD model of the submarine is presented in Figure 1.




Figure 1 – CAD model of the 2004 vehicle

## Mechanical Structure and Design

The mechanical design satisfies three major goals:

- Provides a stable vehicle with a range of motion suited to the tasks required of it
- Provides a dry space for components where needed within the vehicle, and


provide sealed high-quality electrical connections to external sensors.

- Mount and accommodate a variety of sensors as required by the onboard Artificial Intelligence (AI) to complete the mission at hand.

Vehicle design for the 2004 competition was augmented by fully modeling all system hardware in the SolidWorks 3D CAD environment. Upgrading to a fully documented CAD design allowed components to be rapidly and efficiently computer machined, in addition to ensuring the compatibility of all vehicle components before manufacture.

### Mechanics – Propulsion

Four custom designed and machined thruster units provide the propulsion for this year's entry. Each unit mounts a single Hacker brushless motor within aluminum housings. The rear of the unit mounts a single three-pin Impulse brand connector designed for submersible use, and a dual O-ring seal ensures watertight integrity. The thruster shafts are fitted with Kort-Nozzel type propellers designed specifically for use in water. A rotating shaft seal provides a watertight rotating connection. An aluminum shroud mounted to the thruster body protects against accidental contact with the propeller, and substantially increases thruster efficiency and maneuvering power. Four threaded holes on the thruster allow for easy mounting with standard hardware.



**Figure 2 – Exploded view of complete thruster assembly**


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 unit for each thruster serves to change the direction when required. Both the relays and the motor controllers are controlled by a NetBurner module located in the lower hull that communicates with the rest of the vehicle through an Ethernet line. This module is capable of maintaining low-level control of the propulsion system even in the event of a complete system failure of the upper hull.

### Hull Construction

The 2004 vehicle uses a 5 ¾ inch diameter carbon fiber hull with a 0.10 inch wall thickness. Using carbon fiber provides us with a hull that is conservatively depth rated to 2000 feet while only weighing a kilogram. The pressure seal is achieved using a knife and bevel seal with a ring of mounting bolts to provide the necessary sealing pressure. This seal type is not likely to fail in the event of rough handling which could damage the sealing surfaces of

more fragile O-ring type seals. The flange design also provides increased seal integrity as external pressure intensifies.

In order to improve the flexibility of the vehicle configuration, we have standardized our endcap connector scheme using SeaCon All Wet split series “Pie” Connectors.



**Figure 3 – SeaCon “Pie” Connectors**

These connectors allow us to connect 6 devices to a single endcap connector, allowing us to connect many more external devices than was possible in previous years. Additionally, pinouts for many standard signal types has been standardized, allowing for rapid configuration of new devices that use serial, Ethernet or video signals. A pair of blind-mate connector boards that mount on the inside of the endcap and on the end of internal rack, respectively, route power and signal connections for each pie slice to a header group inside the hull for easy connection to the internal backplane.

### **Power Infrastructure**

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 approximately 200Wh each (80% discharged). Increasing battery life was a major design goal this year, which has prompted some changes to the power subsystem. The three battery

packs now work in parallel on the same power grid in order to maximize run time and keep our battery load cycles consistent between packs. Additionally, the NetBurner modules, discussed below, allow us to control power switching for up to 48 devices per module. Taking advantage of this, we are able to shut off power to sensors, peripherals and even computers when not needed. These efficiency gains, combined with the high energy density of lithium batteries, provide us with over eight hours of usable operating time for extended-duration missions.

Placing the batteries in the lower hull gives the additional benefit of improved static stability. This is a direct result of the higher mass of the lower-volume battery hull. Combined with the significantly greater buoyancy of the upper hull, this configuration results in a strong restoring moment to roll and pitch perturbations.

### **Sensor Suite**

The design philosophy of the sensor suite for this year’s vehicle is flexibility. A number of sensor systems on the vehicle are used for vehicle control and navigation. These sensor systems are permanently attached to the vehicle, providing basic support to the onboard AI. To support these tasks we have selected a suite of sensors that provide sufficient environmental data for the AI to operate the vehicle. 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 feedback control system.

The second major sensor system is the collision avoidance system composed of two Tritech 500 kHz sonar altimeters. One altimeter faces forward for wall avoidance and the other faces down to report altitude off the sea floor. A low level collision algorithm runs on the propulsion NetBurner module and seizes control of the motors if a collision is imminent.

In addition to the permanent sensors array, the vehicle provides infrastructure support for a wide variety of mission-specific sensor suites. Using the SeaCon pie connectors, the vehicle is capable of supporting up to 14 external devices. We have outfitted the vehicle with additional sensors specific for this year's competition.

The key mission specific sensor system is our vision subsystem, comprised of two underwater cameras, one forward looking and one directed downward, both mounted on the bow section of the vehicle. The forward looking camera provides for intermediate to short range navigation to the light and the target setup. The downward looking camera provides fine control of bin recognition for dropping markers. The cameras are aligned with the altimeters so that vision data can be combined with altimeter readings in detecting target bins.

### **Upper Rack and Backplane**


A major component in our efforts to improve vehicle flexibility is a new modular rack with integrated backplane units. The rack physically breaks out into 8-inch segments complete with

mounting rails that rigidly bolt together with thumbscrews to form the full rack. Each rack segment has a backplane board mounted to the bow-side bulkhead. The backplane board provides power and signal support for every component that is mounted in the rack segment. The backplane boards are connected to each other by four cable harnesses so that a signal attached to one segment will be routed to every other segment by the backplane network. Power signals are routed using heavy gauge wires capable of handling between 2 and 8 Amps per circuit, depending on configuration. Signal lines are routed using 2 sets of high performance Samtec EQCD impedance matched micro-coaxial ribbon cables.

The backplane network supports 20 power lines, 12 serial lines, 7 Ethernet lines, 3 USB lines, KVM support for 3 computers, 5 video lines and 28 unspecified signal pins for low current applications. The signal distribution was selected not only to support the vehicle's current needs, but to allow for future expansion.

The leading segment of the rack is special in that it contains the blind-mate board used to connect the rack to the external endcap connectors. This segment is home to permanent fixtures of the vehicle's infrastructure, such as DC-DC power supplies and the Ethernet switch. The leading segment is also used to route signals from the header groups that correspond to SeaCon pie connectors onto the backplane network.

Every other rack segment contains a group of configuration specific components, allowing for rapid vehicle configuration changes. A new external sensor component can be mounted, attached to an open SeaCon pie connector, routed along the



backplane and read by an internal device that was snapped into the rack all in a matter of minutes.

### Computers

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 capabilities.

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 Versallogic Jaguar platform whose purpose is to handle processor intensive operations such as vision and acoustic processing in real

time. 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. Each system is mounted to a custom PC-104 stack that both secures the stack in a modular fashion to the rack and provides for easy removal and replacement of hard drives.

Debugging and configuration switches, including power switches, Keyboard/Video/Mouse (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.

### **Hydrophone Processing**

The hydrophone system is a passive acoustic system designed to extract a heading to an acoustic source emitting short pings in the range of 20 kHz to 30 kHz. The heading to the source is calculated based on the time differences between when each hydrophone receives the signal.

Four hydrophones are positioned at the extreme corners of the vehicle and raised above the upper hull, to eliminate interference, by aluminum arms. The basic goal in designing the geometry of the array was to maximize the total distance between all the hydrophones,

producing longer delays and minimizing error due to finite sampling resolution.

Each of the 4 hydrophone channels comes directly into the upper hull as a differential analog signal via balanced lines. Inside the upper hull are 2 custom built circuit boards stacked on top of each other for amplification and analog to digital conversion.

On the first board, each of the channels coming into the upper hull goes through an AD605 variable gain amplifier from Analog Devices, which takes a differential signal as input and outputs a single-ended signal. Analog automatic gain control circuitry, which uses 2 different length integrators and feedback, sets the gain control on the AD605 to always output a ping signal that is just under 2.5 V peak-to-peak.

On the second board, each of the roughly 2.5 V peak-to-peak ping signals is passed through an AD7644 16-bit ADC and then a 4 to 1 16-bit multiplexer. A Netburner module with the Motorola MOD5282 microprocessor core running at 66 MHz provides a clock signal for the ADCs. The Netburner is continuously running an interrupt at 100 kHz. Inside the interrupt signals are sent to the AD7664s to start new conversions and then the 16-bit samples are read into the Netburner through GPIO while sending the appropriate signals to the multiplexers to get all 4 channels. The Netburner stores the samples in an array on length 1000 for each channel and for each full array determines if there may be ping present by basic thresholding. Arrays containing potential pings are streamed via Ethernet packets to the Jaguar PC104 computer.

The Jaguar performs tight digital bandpass filtering around the known pinger frequency, sinc-interpolation, and cross-correlation to extract the delay

times between each channel. Once the delays between each channel are known, the horizontal heading to the pinger is calculated based on simple geometry and the speed of sound in water, and sent to the Crusoe PC104 control computer.

Theoretically, with our array of 4 hydrophones in a rectangular plane, the 3-dimensional heading, 3-dimensional distance, and relative horizontal coordinates can be calculated exactly, and the absolute value of the relative vertical coordinate can be calculated. However, using our signal processing system, the accuracy the horizontal heading measurement is much more accurate than any of the distance calculations and is also the most crucial for the mission.

### **Motor Control**

Physical motor control relies on four brushless motor controllers. Each is controlled directly, through the statecast mechanism, by a NetBurner module in the lower hull. It is responsible for generating the PWM signal driving the motor controller, and for translating the non-linear response of the thrusters into a more linear version suitable for a PID control algorithm. This linearization is performed by creating a percentile scale in which 100% provides the maximum operational thrust, 1% generates the minimum available thrust, and 0% corresponds to full off.

Such a translation eliminates the thruster dead zone (the area in which the applied power is insufficient to create a meaningful response from the thrusters) and decouples the high-level control algorithms running on the mission computer from the specifics of the underlying hardware.

Three PID (Proportional Integral Derivative) algorithms control the thruster hardware, through this abstract interface, to move the vehicle according to the requests made by the mission execution processes. The PID system is actually a combination of three separate PID feedback loops: for depth, for velocity, and for heading. They run the standard PID feedback algorithm using constants derived empirically.

The depth PID loop operates independently, since it is the only software component responsible for control of the fore and aft thrusters. The heading and velocity loops both operate on the lateral thrusters, and must therefore control the same hardware cooperatively. This requirement is satisfied by running the velocity PID as a common mode thruster value, and the heading PID in differential mode. Net thruster settings are simply the sum of the two outputs.

### **Distributed Responsibility**

The CUAUV software suite is composed of multiple processes running on multiple machines, each of which is responsible for a specific subset of overall vehicle operation. Processes communicate using two complementary mechanisms: a custom library that provides performant, reentrant, network-transparent remote procedure call (RPC) functionality, and a set of “shared variables” resident in shared memory on each participating machine.

The shared variable scheme is the primary mechanism for interprocess communication (IPC). It is best suited to periodic updates of simple variables; it is used, for example, to communicate the current state of battery charge, the desires of the high-level mission process,

readings from external sensors, and other, similar pieces of information.

The RPC mechanism is used to create a set of call-based interfaces to various vehicle system services: the “varshare” service, for example, provides the shared-variable management API to every process which requires access to the shared variables resident on a machine, and the “broker” service allows processes to locate services (such as “varshare”) which have registered themselves with the broker. The services architecture allows IPC to take place which may not fit easily within the shared variable scheme.

These custom, performant IPC mechanisms provide us with a significant amount of implementation flexibility for each software component, especially in terms of language and platform used. IPC library bindings are maintained for the C, Ruby, C#, and C++ languages, each of which is used for tasks for which it is well suited: C and C++ are used primarily for control and vision processing, Ruby is used for scripting and debugging, and C# is used for high-level mission code. The distributed architecture also makes it easy to develop and debug applications: component processes can be run entirely on dockside development machines, shared variable state is logged and visible to developers at all times, and component processes can often be updated and restarted—possibly on an entirely different machine—without affecting other vehicle software components.

### **Universal Ethernet**

The NetBurner Mod5282 device is our solution to the embedded Ethernet problem. The majority of our onboard

sensors communicate via serial links. However, serial is not easily configurable and requires extensive configuration to be set up. To this end, we use microcontrollers that have both Ethernet and serial capabilities to serve as a translator for the serial based devices. The NetBurner is able take serial data, perform some low level analysis or processing, and pass it on to the higher level computers via Ethernet protocols. As part of the modularity effort, also seen in the upper hull rack, we value Ethernet communication because of its ubiquitous nature.

The NetBurner is more than a mere translator. It is based on the Motorola ColdFire 5282 microprocessor. Boasting onboard ADCs, PWM generation capabilities, timers, GPIOs and various serial protocols, the NetBurner is a general-purpose tool. ADCs are used to monitor internal vehicle status – temperature, current draw, battery voltage levels. PWM generated by the NetBurners drive all four motors. The GPIOs are capable of switching relays (used to conserve power by shutting down devices when not in use, allowing for the vehicle to enter a more passive state.

### **Mission Strategy**

High-level mission code is very simple: it proceeds through a sequence of actions, each waiting for a specific trigger condition before proceeding to the next.

We first perform point-to-point navigation, using integrated position information from the DVL, to travel through the validation gate. Once through, the forward-looking camera is used to keep the vehicle oriented toward

the light, and the downward-looking camera is used to find the target array. When the vision processing unit identifies the target below the vehicle, the AUV will dive to hover above the target, using the downward-facing altimeter to maintain its altitude above the top bin. When the vehicle is positioned above the top bin, it releases its markers, and uses the passive sonar array to maneuver to the recovery zone and surface.

### **Vision Processing**

Vision processing in the CUAUV 2004 system is performed on the Pentium III 850 MHz PC104+ onboard computer. Two external cameras, one forward- and one downward-looking, are used; both cameras interface with the computer via a Sensoray PC104+ 311 framegrabber.

The mission calls for two vision processing tasks. The system must first find the blinking light via the forward looking camera, and extract a bearing to the light. The vehicle must then use the downward-looking camera to locate both the target array and the top bin.

A clustering algorithm is used to search for the grouped cyan pixels corresponding to the blinking horizontal light. This cluster is used to identify the light in the image, and the Cartesian coordinate of the light is generated by averaging the pixel locations within the cluster.

The target finding algorithm uses a different approach. The Hough Transform algorithm locates lines within an image. Because the target appears as multiple boxes, an algorithm operating on the output of the transform can identify the target as multiple parallel lines. This recognition will signal to the

high level AI that the submarine is over the target.

Both vision tasks use general-purpose routines implemented in a custom CUAUV visionprocessing library. The library supports many vision processing algorithms for various situations.

### **Acknowledgements**

CUAUV would like to thank Cornell University, the College of Engineering, and the schools of Electrical and Computer Engineering, Computer Science and Mechanical and Aerospace Engineering for their continuing support of the project. We also owe a debt of gratitude to our corporate sponsors, who make this project possible. They are: Shell Oil Company, NetBurner, Samtec, SeaCon, and Check Yourself Inc. Finally, a special thanks to our advisor Dr. Kevin Kornegay, the staff of the ECE accounting office, and Cornell Aquatics.