

ORCA-VIII: An Autonomous Underwater Vehicle



Joshua F. Apgar, Marisa Bober, Winnie W. Cheng, Jonathan S. Edelson, Jason Furtado, Ben Howell, Albert Huang, Samuel H. Kenyon, Ara N. Knaian, David Newburg, Jon Santiago, Dimitrios Tzeranis

Massachusetts Institute of Technology
Project ORCA
77 Massachusetts Avenue, Room 4-405
Cambridge, MA 02139
<http://web.mit.edu/orca/www>

Abstract

The ORCA-VIII is a fully autonomous submarine built to compete in the 2005 International AUV Competition. The ORCA-VIII is 1.27 m long, 0.711 m wide, 0.457 m deep, and masses approximately 45 kg. The vehicle is propelled by a pair of horizontal thrusters mounted on the sides and a pair of vertical thrusters on the bow and stern. Autonomous navigation is aided by a suite of instruments including a water pressure depth sensor, a fluid-bulb inclinometer, a DSP based sonar direction finder, a magnetic compass, an inertial measurement unit, and three underwater video cameras with computer vision software.

The ORCA-VIII follows the basic modular design philosophy of previous incarnations of ORCA. The mechanical frame and core systems are mostly unchanged from the 2004 entry allowing development efforts to be concentrated on building completely new software-hardware systems and reconfiguring/improving existing systems. We have three new cameras—a line scan and two color area scan cameras—with machine vision processing developed for this year's challenging mission.

1. Introduction

ORCA-VIII is designed according to the guidelines of the 8th International Autonomous Underwater Vehicle Competition. The competition arena is located at the SPAWAR TRANSDEC facility in San Diego, California. The arena is an oval 200' wide and 320' long, with a flat bottom 16' deep. There is a 160' diameter semi-spherical depression in the middle that extends to a 38' depth at its center.




Figure 1 Diagram of navigational strategy

In the competition arena there will be a validation gate, a Random Order Light Box, a Docking Station (Station A), a Pipeline Inspection structure (Station B), and a Surface Zone (Station C) marked by an acoustic pinger.

The competing vehicles will have 15 minutes to complete the mission. Each vehicle must demonstrate basic operation by passing through

the validation gate before attempting any other portion of the course. The vehicle can read the Random Order Light Box to determine the randomly-selected order of tasks (order of stations to complete), e.g. Station B followed by Station C followed by Station A, as opposed to doing the tasks sequentially (A-B-C or C-B-A). The vehicle then navigates to and completes each station autonomously.

ORCA-VIII was built to complete this mission reliably, repeatedly, and safely. The modularity of design allowed easy testing and modification throughout the process of construction. An improved 3D simulator also aided in developing and testing new software and hardware. All of the modules were tested to insure safe operation.

2. Mission Strategy

Our mission is directed by a pre-programmed state machine for random order sequence task completion. This section outlines the basic sequence of maneuvers the vehicle will use to complete the mission. Upon activation, the vehicle will dive to a cruising depth of 1.5 meters and will navigate through the validation gate using dead reckoning. Then the vehicle will detect the frequency and color of the random light box to initiate one of the 4 sub-state machines. Each of these is responsible for completing the mission tasks in the indicated order. In this way, the mission is optimized for the particular order of tasks. In general, to get to the docking station, the vehicle will use the forward-facing camera to locate the light source and navigate towards it. To get to the bin, the vehicle will first detect the pipeline with its downward-facing camera. Once detected, it will navigate with feedback information from the vision algorithm that tracks the presence and direction of the pipeline. The vehicle follows the pipeline closely until the image processing algorithm indicates the presence of the bin. It will then center over it and drop the markers. To get to the recovery zone, the vehicle will detect pings and correct its heading to face the acoustic beacon. It then centers over the zone and surfaces.

The figure to the right shows the navigation strategy ORCA will use to navigate from point

to point in the arena. Arrows indicate the path of travel and the color/line style indicate which sensor will be used for that leg of the mission. Note that some paths are not directly implemented. For example we navigate from the recovery zone to the docking station by way of the pipeline. This is designed to minimize the amount of dead reckoning.

3. Mechanical and Electrical Systems

The mechanical design of ORCA-VIII is meant




Figure 2 Wireframe drawing of ORCA-VIII

to be simple and modular. The frame was divided into four larger mechanical parts that could be constructed separately: the skeleton, the thruster mounts, the waterproof housing, and the sensor mounts.

The skeleton is made of 80/20 10-series T-slotted aluminum extrusions. Modifications can easily be made to the simple rectangular design. The slotted extrusions allow all of the components to slide into place for easy mounting and vehicle trimming.

Each thruster mount is made of three aluminum pieces, welded together. One rectangular block attaches the mount to the 80/20 frame. A streamlined foil holds the thrusters away from the vehicle, without adding significant drag or blocking flow. A motor mount allows the cylindrical thruster body to be hose-clamped firmly to the end of the mount.

The single dry compartment is a 27-inch long, eight-inch diameter PVC pipe. The hull is mounted on the aluminum frame and closed with PVC end plugs with double O-ring bore seals. Through-hull electrical connections are made

with bulkhead connectors mounted into these end plugs. This compartment holds the computer, radio transceiver, antenna, motor drivers, and power electronics. The electronics are mounted on guide rails supported by parallel discs that slide inside the compartment. This assembly slides in and out of the dry compartment on Teflon rods used to prevent the discs from scraping the hull. This design allows maximum use of the volume of the compartment. The sliding electronics assembly can be removed without disconnecting any cables because a blind-mating multi-pin connector links the last disk to the compartment's PVC end plug.

The sensors and batteries are mounted directly on the aluminum frame. This versatile design allows quick changing of the batteries. The batteries are placed low on the vehicle, to increase its righting moment, making the ORCA-VIII passively stable in pitch and roll.

3.1. Primary Dry Hull

The hull of ORCA-VIII has a single watertight compartment to house the computer and most of the electronics. The batteries and sensors are mounted on the aluminum frame in an open sensor bay below the tube. Two vertical thrusters, mounted at the bow and stern, control the vehicle's depth and pitch. Two horizontal thrusters at the sides control the forward velocity and heading of the submarine.

Electrical connections through the hull are made with hermetically sealed locking multi-pin connectors. The connectors are rated to a depth of 80 m. Each outboard component connects to the vehicle using its own receptacle mounted in the PVC end plug at the stern end of the electronics compartment. In addition to the outboard component connectors, there is a tether connector for development and testing.

3.2. Thrusters

ORCA-VIII uses four maneuvering thrusters, attached via their mounts onto the aluminum frame. Two horizontal thrusters drive the vehicle forward and back and allow control of the vehicle's yaw. Two vertically mounted thrusters allow control of depth and pitch. The thruster

assemblies are provided by Inuktun Services




Figure 3 Wire frame drawing of Inuktun thruster

Limited. Each one draws 7 A at 24 V producing approximately 15 lbs of thrust.

The motor controllers used in the vehicle are the same ones introduced last year. Which have proven suitable and reliable. The firmware was modified this year, primarily to add an autostop timeout. This feature stops the motors if control signals cease for 1 second.

The motor control boards are described in more detail in the ORCA's 2004 journal paper. In short, they are based on OSMC (http://www.robot-power.com/osmc_info/), modified for the tight space requirements and lower current requirements of the vehicle. This speed controller is a simple and robust H bridge amplifier capable of operating over a wide input voltage range and at rather high currents. An HIP4081 driver chip provides gate drive for all of the MOSFETS, of which there are 4 in parallel for each H bridge leg. Various protection components are used to protect the components from possible motor transients. The 12V power supply is provided by a DC-DC converter from the main DC supply voltage. The ORCA motor controllers extend OSMC with an Atmel ATmega32 microcontroller (programmed in C) and communication capabilities to the computer through a serial port, along with two other control modes for testing: servo channel control digital input, and onboard trimmer potentiometer setting.

3.3. Power Distribution and Monitoring

Two separate voltage busses power ORCA VIII—one for the motors and motor drivers, another for the computer, electronics and sensors. Devices on these busses are individually fused to prevent an isolated failure from bringing the whole system down. Efficient boost and step down converters provide all voltages other than the battery voltages.

Power from the batteries is switched through a set of mechanical relays. Two waterproof magnetic kill switches with colored ripcords can be used to power down the motors or the computer.

The batteries are sealed in external battery pods that attach to the vehicle via through hull connectors. The electronics battery pods have 20 NiMH 4/3A cells wired in series giving a nominal voltage of 24 V and a capacity of 4 Ah. A single electronics battery pod powers the vehicle for four hours on a single charge. The motor battery pods have 20 NiMH SCU3300 cells wired in series giving a nominal voltage of 24V and a capacity of 3.3 Ah. A pair of motor battery pods powers the motors for a half hour to an hour depending on use.

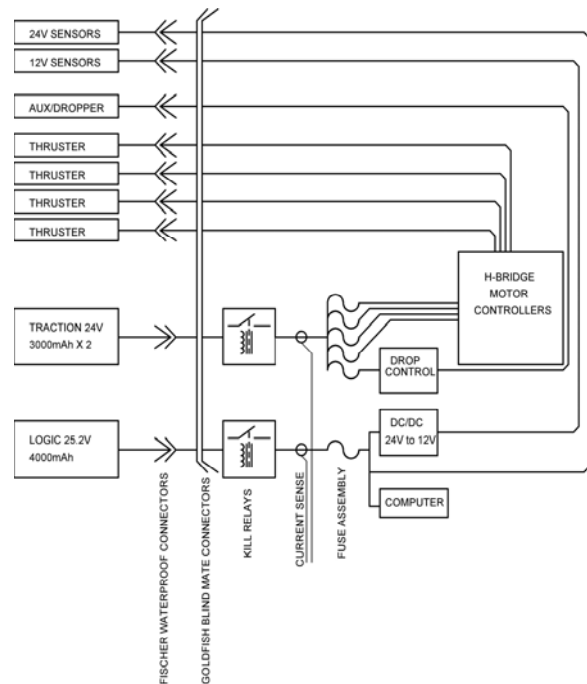


Figure 4 Vehicle power distribution




Figure 6 Wire frame drawing of endcap machined for through hull connectors

The use of external battery pods has the advantage of allowing the batteries to be changed much more quickly than if they were housed in the dry-hull.

The vehicle is equipped with a power monitoring system. The voltage and current through the batteries and motors are read, and a PIC microcontroller reports values to the computer using a RS-232 serial interface. The monitoring system also has a temperature sensor to alert the operator to over temperature conditions. Blown fuses are signaled with LEDs for fast identification.

3.4. Marker Drop Mechanism

The marker drop mechanism for ORCA-VIII is comprised of two permanent electromagnets that each hold one 1.5 inch steel ball bearing marker. When off, the device is a simple magnet that holds the marker in place. When on the electromagnet creates a field that cancels the permanent magnet's field, allowing the marker to drop. The permanent electromagnets operate on 24V from the motor bus, and weigh about 1 lb. Each electromagnet can be actuated independently. The permanent electromagnet provides power savings over an electromagnet since holding the markers in place, which is the state of the dropper for the majority of the time, requires no power.

3.5. Onboard PC104+ Computer

All navigation and control code is run under Linux on a Pentium-based PC/104+ embedded computer. This computing platform provides a

stable and familiar programming environment, is amenable to remote operation, has modular standard peripherals and has a small install footprint.

The ORCA-VIII PC/104+ stack consists of a CPU card, a switching power supply, two eight serial port expansion cards, and a Cardbus module. Most sensors and actuators interface to the computer using the RS-232 serial protocol, although the three cameras use Firewire to a multi-port 1394 FireWire CardBus PCMCIA card (Adaptec FireConnect).

Data can be sent to and from the vehicle via a 100Mb Ethernet link. During development this high speed data link allows for allows team members to inspect mission data in real time including streaming video while vehicle is submerged. In an autonomous run, data is recorded to the 60GB hard drive for post processing. This combination has proven invaluable when debugging complex autonomous maneuvers.

3.6. PCBs

In addition to many off the shelf components ORCA-VIII uses custom printed circuit boards extensively. These allow the use of small surface mount components, reduces fragile point to point wiring, and make for a neater, more robust submarine.



Figure 5 Vehicle Dropper Markers. These markers have proven to be robust and reliable.

4. Optical Systems

The ORCA-VIII has three optical sensors. Two color digital video cameras pointed forward and downward, and one grayscale line scan camera pointed forward.

4.1. Pipeline Detection/Tracking

The goal of the pipeline vision system is to capture and analyze images of the pond bottom and determine the existence, position, and orientation of both the pipeline and the pipeline "break." This break is a bin, which appears from above basically as a black rectangle within an amorphous shape white shape.

The ORCA VIII is equipped with two color area scan video CMOS cameras (Prosilica EC750C). They are mounted in individual underwater housings made for easy field reconfigurability. These cameras allow the exposure time and gain to be set in software, making them particularly suited to machine vision applications. Images from these cameras are sent via Firewire to a multi-port 1394 FireWire CardBus PCMCIA card (Adaptec FireConnect) which is plugged into a PC/104+ Cardbus module (PCM 3794) in the onboard computer. Machine vision algorithms running on the main computer analyze the images and provide real-time targeting information to higher-level control programs.

One part of the vision algorithm software is the scanner, which checks pixels such as to minimize the number of pixels scanned and still recognize objects. If a pixel is within the range of an object's color, the algorithm then finds all the contiguous pixels around the seed, to leaving the object. Each type of object has valid conditions that must be met to ensure it is the object. We calculate moments of the object to find the heading the vehicle must take to move toward the object. Each frame is processed in this way to allow access to object data for the mission software to use to make decisions.

The vision algorithms are based largely on the fact that our interesting objects are connected blobs of color in some range. We implemented HSVColor and HSVColorRange classes to encapsulate a single HSV color and a range of colors between two such colors. Another class, Massachusetts Institute of Technology

HSVMultipleColorRange, was added to create ranges which contain many disconnected ranges (white/black bin). It can then be used to scan parts of a picture for areas of interest containing these colors, using a fill algorithm to mark related pixels once a single interesting pixel is found, and never scanning pixels which have already been marked. Each interesting area is created by initializing an AreaOfInterest object. This object performs the fill algorithm and collects a lot of data about the area, such as the leftmost, rightmost, lowest and highest point, and information necessary to compute the centroid. Depending on what we are scanning for, these areas of interest can be analyzed to classify them as bins, lights, pipelines, etc.

4.2. Blinking Light Detection

The ORCA-VIII includes a modulated light source detection system. This system detects and tracks the docking station light, and is used in conjunction with the forward color video camera to read the Random Light Display signpost. The docking station has a red light source modulated at 5 kHz and 5 Hz. Once it is struck, its modulation changes to 7 kHz and 7 Hz.

The ORCA-VIII light sensor is constructed using an off the shelf line-scan camera (ISG LW-ELIS-1024a-13194, Imaging Solutions Group, Fairport, NY). This camera returns 512-pixel grayscale lines to the main computer via Firewire at 20 kHz and at 14 bits-per-pixel, values sufficient to demodulate the docking station's high and low frequency modulation in software. A custom optics system using spherical and cylindrical lenses is used to focus light from the area onto the linear sensor.

By detecting the location of the center-of-mass of modulated light on the linear image, the main computer can compute a bearing to the docking station. By determining the magnitude of modulated light, the computer and determine whether the light is within range. By determining the modulation frequency of the light, the computer can determine whether the docking station has been knocked over.

The purpose of the blinking light detection subsystem is twofold: 1) to find/track the Docking Station's omni-directional red LED and output a bearing to it, and 2) detect and analyze the Random Light Display signals at the beginning of the match.

ORCA VIII has been equipped with a forward-mounted line scan camera (ISG LightWise) dedicated to blinking light detection. This can output a full greyscale frame (1x1024 pixels at 14bpp) at a high enough sample rate for us to detect the high frequency components of the Docking Station and the Random Order Light Display. The line scan camera is mounted on the front of the vehicle in a custom housing with cylindrical lenses and a red-pass filter. The forward-looking color area scan camera can optionally provide redundant data to be fused with the frequency detection.

The image processing detects the presence of certain frequencies to find and track the Docking Station. The problem is to detect dual-modulated 50% duty cycle square waves, as the blinking lights in this competition have high-frequency (3 kHz) carriers gated at a low frequency (3 Hz). We have the capability to also detect the 7 kHz signal modulated at 7 Hz, which the docking station emits when it has been "docked" (tipped over). To demodulate the high-frequency carrier signal of the docking station LED, we wrote software using a sliding Goertzel DFT (Discrete Fourier Transform) algorithm. This filter computes coefficients for specific DFT bins centered on N-spaced frequencies of interest. The N-point windows for each pixel are re-processed every input sample, making use of the already-processed spectral components through DFT shifting. This allows us to get coefficients at any time, although we don't need the coefficients every sample so we recompute the needed coefficients (optimized for magnitude) every 1-10 frames. There is an initial small startup period of N frames establish the window, which can be reset in the case of saturation.

The filter gives us magnitudes at the frequency of interest, e.g. 3 kHz, which is compared to a threshold to determine if the frequency is indeed present. This software module outputs a

heading of the docking station based on the pixel index of a detected signal. Optionally, we can estimate the distance to the Docking Station from its width in pixels.

These software modules are also used with a slight automatic configuration change to read and interpret the Random Light Display at the beginning of the match. This has green and red lights at signals of either 2 kHz modulated at 2 Hz, or 5 kHz modulated at 5 Hz. Since we have already made the system capable of demodulating 7 kHz, we don't have to increase the pixel sample rate or seriously modify the image-processing.

5. Acoustic Systems

The ORCA-VIII employs three acoustic elements. Two sonar range finders for navigation and obstacle avoidance, and one passive sonar that can be used to locate acoustic sources such as the pinger in the recovery zone.

5.1. Sonar Range Finders

ORCA-VIII uses two Tritech PA500 sonar range finders. One is used to measure the distance to the floor of the arena. The other is mounted on the starboard side and measures the distance to the sidewall for collision avoidance. The PA500 measures the distance to the bottom by actively pinging at 500 kHz and measuring the time delay to the echo return. It returns the measured distance over an RS-232 serial port. The units are operable from 0.1 to 10 m distance, suitable for the size of the competition arena.

5.2. Passive Sonar System

The ORCA-VIII includes a passive sonar system to determine a bearing to the Surface Zone acoustic pinger. The passive sonar unit is mounted to the 80/20 frame with a neoprene cover for acoustic decoupling from the frame.

The system detects pings using four hydrophones mounted in a pyramidal array. The hydrophones are mounted to the bottom of a waterproof enclosure, which contains processing electronics. The passive sonar system communicates with the ORCA-VIII main computer using an RS-232 serial port. For each ping received, the unit transmits the bearing and

elevation angle to the transmitter in degrees, the frequency of the ping, and the time in milliseconds since the last ping.

The system computes the angle to the pinger by measuring the time delay between the ping signal as received at each of the four hydrophones. The data acquisition and embedded computing hardware used are described in detail in the ORCA journal paper written for the 2002 AUVSI competition. Each hydrophone signal is digitized and input to a DSP microcomputer. The DSP bandpass filters and thresholds the signal from one hydrophone to find the start of each ping. The system captures the next 2 ms of signal from each hydrophone for further processing.




Figure 7 Signals from a two element array showing the angle dependent phase delay

To determine the ping frequency, the system calculates a 2048-point FFT of one of the signals, and finds the maximum energy bin. The system then determines the time delay between each pairwise combination of hydrophones, using the method of generalized cross-correlation as described in Underwater Signal and Data Processing by Joseph C. Hassab. The system uses only the first 150 microseconds of ping energy to determine the hydrophone pair delays, to reject multipath echoes.

To find the bearing and elevation angle to the pinger, the calculated delays are used to solve a system of simultaneous equations that gives the required angles in terms of the delays. The

system of equations is overconstrained, so the program computes a least-squares optimal solution. The equations are derived from the geometry of the hydrophone array and the plane wave approximation.

The calculations for each ping take about 80 ms. Once processing is complete, the bearing, elevation angle and frequency for the ping are transmitted to the main computer in an ASCII string for navigational use.

6. Navigation Sensors

The ORCA-VIII has three main navigational sensors: a magnetic compass, an inertial measurement unit, and a depth sensor.

6.1. Depth Sensor

A Sensotec TJE series analog output pressure sensor measures the depth of the vehicle. A PIC microcontroller provides analog to digital conversion and communications to the main computer.

6.2. Compass

A Honeywell HMR3000 magnetic compass and inclinometer provides the submarine with heading roll and pitch measurements at 10 Hz.

6.3. Inertial Measurement

A custom build inertial measurement unit measures the vehicle's six axis acceleration using a set of Analog Devices MEMs accelerometers and rate gyros. The angular rate measurement is integrated to give a second measure of heading, roll and pitch. This information is combined with the compass data using an Extended Kalman Filter.

7. Software Development

For development and testing purposes a tether can be attached to the vehicle to make communication with the computer possible. The computer uses Java RMI (Remote Method Invocation) to communicate with multiple on-shore computers. From each station the vehicle can be remotely operated with a joystick, and all variables and sensor values can be inspected and modified with a graphical user interface. In addition, the main control program can be remotely modified and recompiled. All of this

can be done while the submarine is submerged and operational.

7.1. Physics Simulator

The control program has a simulation mode that uses a 3D Java visualization and rigid body physics provided by ODEJava (<http://odejava.org>) to simulate the vehicle. The simulator is able to simulate all of the sensor data—including the cameras—and effectors of ORCA-VIII. The simulation mode allows control code and algorithms to be developed and debugged in the lab before being tested on the actual vehicle and/or in the water, and has significantly decreased the time needed to bring new features online.

7.2. Mission Control Software

The mission control software is implemented as a multithreaded Java program. Each sensor has an associated driver thread that communicates with the device and scales its data into engineering units. An autopilot thread keeps the vehicle's depth, heading, pitch, and speed at desired setpoints. The autopilot uses PID control on the four thrusters to servo the values returned by the pressure sensor, compass, inclinometer, and DVL to the desired set points.

The ORCA-VIII controls its depth using feedback from a pressure sensor or sonar altimeter. Heading is controlled using feedback from a magnetic compass. The submarine can navigate in relation to an acoustic beacon with its four hydrophones. To recognize visual cues and aim for the objects of interest, ORCA-VIII uses three cameras and machine vision software.

8. Conclusion

The new mission for 2005 provided challenges that prompted several improvements on previous years' vehicles. After the modifications and testing, we look forward to participating in the 2005 International AUV Competition.