Failure Detection in an Autonomous Underwater Vehicle

Alec Orrick, Make McDermott, Department of Mechanical Engineering
David M. Barnett, Eric L. Nelson, Glen N. Williams, Department of Computer Science
Texas A&M University
College Station, TX 77843 USA

Abstract—**A system has been developed for failure detection and identification in the depth and heading control of an AUV. A redundancy management technique was implemented using the CLIPS expert system shell. The term redundancy, as used here, does not mean that sensors are duplicated but that independent values of the same quantity can be calculated by combining data from several different sensors. The rules used for failure detection and identification are presented and discussed. This failure detection scheme was implemented and tested on the simulator for the Texas A&M AUV Controller. Failures were introduced and the performance of the system was evaluated based on its accuracy and time response in correctly detecting and identifying failures. All single failures and most multiple failures were detected and identified correctly. False alarms were avoided by requiring several successive occurrences of an aberration before it was recognized as a failure.**

INTRODUCTION

For an autonomous vehicle, failure detection and identification are crucial to enable the vehicle to complete its mission successfully. This paper describes a system for failure detection and identification in the depth and heading control of an Autonomous Underwater Vehicle (AUV). Failure detection is done on all of the heading and depth control surfaces and sensors. Control surfaces considered are the stern planes and rudder. The five sensed values which are checked for failures are depth, pitch (inclinometer), angle of attack, angle of sideslip, and heading (magnetic compass). An expert system processes sensor data using a redundancy management technique. Data from various sensors are used to determine:

1. whether a failure has occurred,
2. what component has failed, and, if possible,
3. how the component failed.

The rules used for failure detection are those which would be used by an expert human operator, i.e., they are based on a thorough understanding of the dynamic response of the AUV and common sense, rather than being based on sophisticated mathematics such as parameter estimation [3].

This failure detection system was designed to operate as part of the Autonomous Underwater Vehicle Controller (AUVC) developed at Texas A&M University. The AUVC consists of a set of cooperating distributed programs, each responsible for some part of the planning, control, diagnosis, and failure recovery functions of the AUVC. When integrated with the rest of the AUVC, the heading and depth control diagnoser discussed here will serve as one of several low–level diagnosers, monitoring various components of the vehicle's equipment and performance. The low–level diagnosers report to a global diagnoser that produces a coherent system–wide diagnosis of the vehicle's status and capabilities. The global diagnoser's report is passed to the AUVC's planning components, which evaluate any failure's impact on the mission and modify the mission plan if necessary.

A simulated vehicle was used in the development and testing of the AUVC, including the diagnoser presented here. The simulated vehicle is large and relatively slow–moving, with control surfaces at the stern. The control surfaces are a vertical rudder and horizontal planes. The planes move as a unit, as does the rudder, so there are only two inputs for heading and depth control: plane position and rudder position.

COORDINATE SYSTEMS AND DIRECTIONS

In analyzing the motions of submarines three coordinate systems are used. The first is an inertial (nonrotating) system XYZ assumed fixed to the earth. The position of the vehicle is expressed in this frame. The second is a water fixed frame. This frame has the same orientation as the inertial frame but is translating with the water at the velocity of the water current. This frame is needed because all hydrodynamic forces and moments are expressed in terms of vehicle linear and angular velocities and accelerations relative to the water. The third frame is a body fixed frame xyz that is fixed to the AUV so that moments and products of inertia are constant in this frame. The origin of this frame is at the geometric center of the AUV, the positive x axis pointing forward, the positive z axis pointing down perpendicular to the floor, and the positive y axis completing a right hand system (pointing to the right). Fig. 1 illustrates the forces, moments, velocities, positions, and angles in the body fixed frame. The xyz components of the vehicle linear velocity expressed in the body frame are denoted by u, v, and w. The xyz components

0-7803-1808-0/94 $4.00 © 1994 IEEE

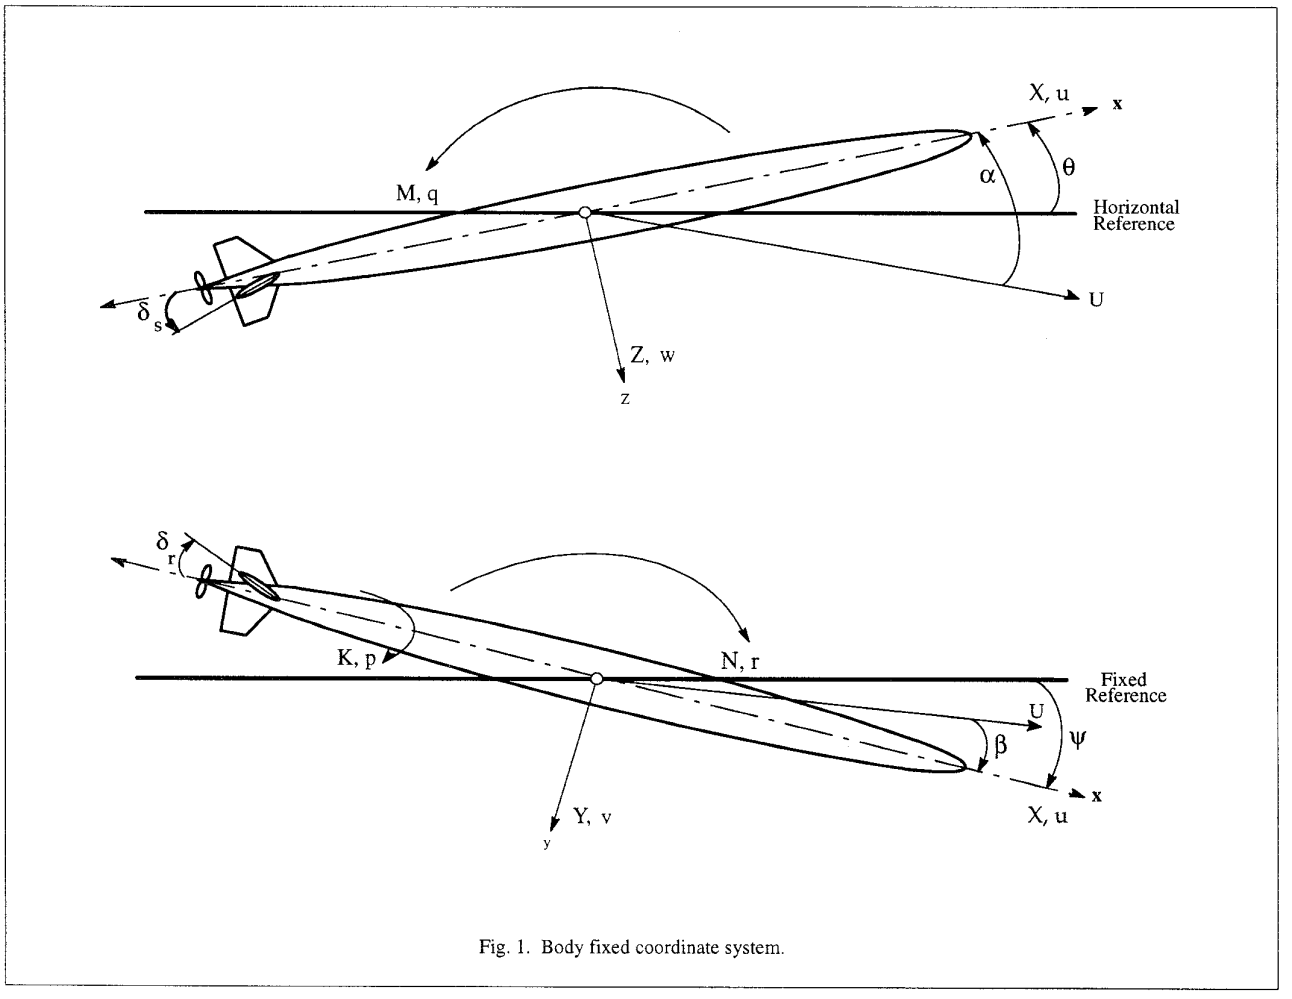

Fig. 1. Body fixed coordinate system.

of the vehicle angular velocity expressed in the body frame are denoted by p, q, and r. The orientation of the body frame with respect to the inertial frame is defined by the roll (ϕ), pitch (θ), and yaw (ψ) of the AUV. These rotations can be taken in sequence (roll, pitch and yaw) to build up the transformation matrix in Fig. 2, which transforms a body fixed vector to the inertial frame. This transformation is needed since some quantities are measured in an inertial frame (e.g., depth) while most others are measured in the body fixed frame. Comparison of two independently determined numerical values of the same quantity requires that both values be expressed in the same coordinate frame.

The sign conventions associated with the coordinate frames described above provide no easy way to perform the consistency checks required for failure detection based on the signs of the vehicle states and controls. For example, a positive sternplane deflection will cause a negative pitch rate and pitch angle but a positive depth rate. To make it easier to code, debug, and maintain the failure detection rules, directions of up, down, left, and right are defined and used as described below.

Directions of up and down are used for depth control, where up and down correspond to decreasing or increasing depth, respectively. For the stern planes, a positive deflection corresponds to down because a positive sternplane deflection will

$$\begin{bmatrix} \cos\theta\cos\psi & \sin\phi\sin\theta\sin\psi - \cos\phi\sin\psi & \sin\phi\sin\psi + \cos\phi\sin\theta\cos\psi \\ \cos\theta\sin\psi & \cos\phi\cos\psi + \sin\phi\sin\theta\sin\psi & \cos\phi\sin\theta\sin\psi - \sin\phi\cos\psi \\ -\sin\theta & \cos\theta\sin\phi & \cos\theta\cos\phi \end{bmatrix}$$

Fig. 2. Direction cosine matrix from body fixed to water fixed frames.

cause a dive. A positive value from the inclinometer corresponds to a climb; therefore, it has the direction up. A positive angle of attack α also indicates a climb and is given the direction up. The depth slope, defined as the time rate of change of Z, is also classified as up and down with a positive depth slope designated down. These directions are applied to each control surface and sensed value individually but do not necessarily indicate the direction the vehicle is moving. For instance, if the vehicle is in a dive when the sternplanes are deflected to cause a climb, the sternplanes will be up but the depth rate will be down for a few seconds. This condition is not considered an aberration since it is an expected dynamic situation.

The other directions used are left and right, which are applied to the heading sensors and rudder. Those directions are used for the same reason up and down are used. Left is defined as the direction of decreasing heading angle, and right is the direction of increasing heading angle. A positive rudder deflection will cause a turn to the left. A positive sideslip angle β corresponds to a turn to the right.

Left and right are also used to describe the relation between the actual and desired vehicle heading. The actual heading is to the right of the desired heading if the actual heading is between the desired heading and the desired heading plus 180 degrees. The actual heading is to the left of the desired heading if the actual heading is between the desired heading and the desired heading minus 180 degrees (see Fig. 3).

Methods of Failure Detection

Redundancy Management

Redundancy management [4] requires the comparison of two values of a single quantity obtained by two independent methods. A redundancy management technique for this AUV problem was implemented using the CLIPS expert system shell. The term redundancy, as used here, does not mean that sensors are duplicated but that independent values of the same quantity can be calculated by combining data from several different sensors.

For example, only one depth sensor is used but depth rate can be calculated from the depth sensor and from the combination of forward velocity, roll angle, pitch angle, yaw angle, angle of attack and angle of sideslip. If these two independently determined values of depth rate are not the same, a failure of one of the sensors is indicated; a failure is detected but not isolated. By doing this with other quantities which depend on the same sensors, the failure can be isolated. There are many different ways of accomplishing the task of failure detection but with the sensors that were available, redundancy management appeared to be the best method.

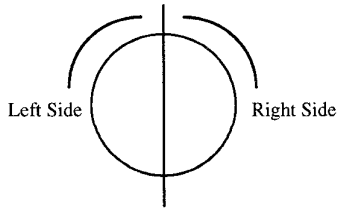

Fig. 3. Definition of left and right side of desired heading angle.

In addition to redundant parameters, error bounds on these parameters are necessary to determine tolerances on the measurements for comparison purposes. Also, part of failure detection is the minimization of false reports of failures, which is influenced by the choice of tolerances and by the algorithms used. For the purposes of this report, an aberration is a condition that indicates a failure but it is not considered a failure until enough occurrences of the aberration warrant that determination. This method is used to ensure that only true failures are reported, with no unnecessary reporting of unusual conditions.

Consistency Checks

This method uses an algorithmic approach to create redundant parameters used in consistency checking. Several instruments could be used to determine the direction that the vehicle is moving without providing directly comparable parameters; consistency checking consists of using those instruments to ensure that one instrument isn't reporting vehicle movement that is inconsistent with reports from the other sensors.

Tolerances

The failure detection rules must decide when a value is "zero" and when two values are close enough to each other to be called identical. These tolerances were determined based on runs with the nonlinear simulation and include allowances for sensor noise. All of the tolerance levels are based on the simulated noise levels of the sensors, as opposed to limitations of the sensors themselves. Once particular sensors have been specified, the sensor precision and accuracy must be factored into those tolerances. Also, limitations in the ability of the computer to access true sensor values must be included in those tolerances (round off and discretization errors).

Aberration Count

The occurrence of step changes in sensor output will be reported after only one occurrence, but all other failures must occur at least N times in a row before they are reported. A threshold value of N=6 was found to work well, minimizing false reports, while producing timely and accurate detection of real failures. Without the six occurrence rule, some failures were being reported that were caused by one failure making a second one seem possible for a second or two, until the failure detection software could catch up to what is actually taking place. For instance, a depth sensor failure to zero depth may be misidentified for two or three seconds until the calculated depth slope reaches zero to concur with the actual failure.

Classification of Failures

The way that a sensor or control surface fails provides information about what could have caused that failure; therefore, the position of the surface or output of the sensor is desired at the point of failure. The system classifies failures into the following types: full deflection, zero deflection, opposite direction, constant, or step. Full deflection is defined as being at the physical extreme for the control surfaces or at the measurement extreme for any of the sensors. Zero deflection is defined as the

neutral point for the control surfaces or zero output for the sensors. Opposite direction is defined as the condition in which a sensor or actuator is reporting a change in the direction opposite that of several other sensors. Constant failures are defined by the sensor being in an unchanging state but not at full or zero. Step failures occur when a sensor indicates a change between sequential time steps which is beyond the physical capability of the vehicle.

IMPLEMENTATION

This diagnostic system was implemented using two computer languages, C and CLIPS (C Language Integrated Production System). CLIPS is a rule–based language that is designed to be embedded in a C program [1] [2]. Rule–based languages take one or more facts and reach a conclusion based on the application of rules to these facts. The C portion of the program was used for numerical computations, whose results were asserted into the fact base used by the rules written in CLIPS.

The rules used for failure detection are described below. The rules and the order in which they are applied ensure that the conditions that indicate a possible aberration actually are indicating an aberration, rather than a dynamic condition or an aberration of a different instrument.

Rules for Depth Sensors and Controls

Depth rate is calculated via two methods and the resulting values are compared. The first method is to use depth values recorded at successive times and calculate the time rate of change of those values, denoted by $(\frac{dZ}{dt})_1$. Because the dynamics of the AUV are slow, a linear curve fit is adequate to estimate $(\frac{dZ}{dt})_1$. The other method of calculating the depth rate is to transform the body fixed velocities to the inertial frame and evaluate the Z component of the inertial velocity. That transformation results in

$$(\frac{dZ}{dt})_2 = (-U\sin\theta + U\tan\beta\cos\theta\sin\phi + U\tan\alpha\cos\theta\cos\phi) \quad (1)$$

A linear curve fit is most accurate near the midpoint; therefore, the values of $(\frac{dZ}{dt})_1$ and $(\frac{dZ}{dt})_2$ must both be calculated at this point to provide a valid comparison. This means that the values of $(\frac{dZ}{dt})_2$ must be stored for several cycles until there are enough values of depth to calculate $(\frac{dZ}{dt})_1$ and make a comparison. For this project five values of depth were used to calculate $(\frac{dZ}{dt})_1$ and the value of $(\frac{dZ}{dt})_2$ at the third point was used for comparison to ensure maximum correlation.

Sternplane failure is checked next. An indication that the sternplanes have failed is that the vehicle is past its commanded depth and diverging. It is possible to be past the commanded depth if the controller overshoots that depth; however, if this is the case, the sternplanes will be changing as the controller tries to correct the problem. Therefore, an additional test condition is that the sternplanes be in a constant position. The next action is to ensure that the indication of failure is not due to a sensor error in the depth gage or inclinometer. This can be done with one check between the depth slope and the inertial Z velocity described previously, since it checks all the sensors at once. The last item to be checked is the forward speed, which ensures that the apparent divergence is not due to the forward speed being too slow for the planes to exercise adequate control authority. If that is the case, that condition should be reported so that it can be rectified if possible.

The previous checks are for a sternplane failure in the current or full deflection condition. A different set of circumstances indicates a failure in the neutral position. First, the condition must be established that the vehicle is neither climbing nor diving and yet not at its commanded depth. Sensor errors must also be ruled out to ensure that this condition is not due to a sensor aberration. Finally, a check on the commanded deflection is made to ensure that the controller did not command the planes to that position; this is to prevent the detection of an aberration if the controller halts the climb or dive for any reason.

An inclinometer failure at the zero position is another error to check. If the inclinometer does indicate zero, this test can be done by checking the depth rate against the inertial Z velocity. Consistency between the depth slope and the plane deflection is necessary to rule out a failure of the planes or the depth slope. Zero–position failures in angle of attack are detected in a similar way as the zero–position inclinometer failures. After establishing that the angle of attack is zero, a check of the inertial Z velocity is done to ensure that the vehicle is not expected to be at an angle of attack of zero. Again a consistency check between the depth rate and the plane deflection is done to ensure that the failure is not due to a failure of the planes.

Depth sensor failures also need to be checked. Failure at full depth of the sensor indicates a condition that the vehicle cannot achieve; therefore, a depth sensor output of full depth indicates a depth sensor failure. An opposite–direction failure of the depth sensor is possible if the sign of the depth slope is inconsistent with the plane deflection and inclination. If both the inclination and the plane deflection are in the same direction then the depth sensor is assumed to be incorrect. Detecting depth sensor failure at a constant depth requires a different set of conditions. The depth slope being zero establishes that the depth sensor indicates constant depth, but it must be determined whether that condition is true or not. A nonzero value of the inertial Z velocity indicates a failure, but does not definitely establish that the depth sensor is at fault. Checking the inclination and the plane deflection for consistency indicates that two consistent instruments are indicating a change in depth when the depth sensor is not, thereby indicating failure of the depth sensor.

Rules for Directional Sensors and Controls

Detecting directional sensor failures begins with the angle of sideslip sensor. If the velocity angle of the inertial measurement unit (IMU) is the same as $(\psi - \beta)$ then the values of ψ and β are both correct. The IMU is part of an Integrated Navigation System (INS) that is monitored by another diagnostic component of the AUVC. Since the IMU is already redundant with a voting

algorithm and filtering methods to ensure its accuracy, checks for failure on that sensor are not repeated in this diagnoser. When the angles don't match, another check to ensure that the heading angle sensor has not failed must be completed in order to conclude that the sideslip sensor failed at a zero angle.

The magnetic compass failures are particularly difficult to determine since it is impossible to determine the failure state of the compass. First, the check is made between the IMU velocity angle and $(\psi - \beta)$. If they are not identical then either the magnetic compass or the sideslip sensor failed. The IMU heading angle is then checked. Actually, only the IMU heading angle is needed to determine failure of the compass, but a second confirmation helps reduce false detection.

The last check on the directional control list are the rudder failures. To establish the condition of rudder failure at a constant (non–zero) position, first the fact that the vehicle has turned past its desired heading and is diverging must be established. If so, the heading must be checked to ensure that the divergence is not due to an error in the magnetic compass. The last check must be to determine whether the controller is trying to correct an overshoot problem. This check is on the rudder to see if it is changing its position. Rudder failure at zero deflection is the last potential failure that needs to be examined. This failure is characterized by a zero turn rate when not at the desired heading. The magnetic compass must also be checked to ensure that the indication of not being at the desired heading is not due to the magnetic compass having failed.

There are more failures than are described above but the other failures are self explanatory. Failures of some instruments in the full position represent conditions that the vehicle will not be capable of, and thus no checks are required to determine if the failure is a secondary effect of another problem. Detecting step errors involves a check on the last value and the current value to determine whether there was a change bigger than physically possible for the vehicle.

TESTING

Testing of the Individual Rules

All of the above rules were tested individually to verify that the failures were detected, and to determine the amount of time needed to detect the failure. The steps in the angle of attack sensor, inclinometer, depth sensor and plane deflection angle were detected at the time they occurred, as expected. The inclinometer failure at full deflection and the angle of attack failure at full deflection were both detected at 6 s after occurrence, because of the required six occurrences before being reported. These failures were simulated by setting the value given to the control function to the maximum at the time of 76 s into the simulation run. The plane failure at the current or full position was not detected at the time it occurred—it took approximately 20 s. The reason for this is because the rules require that the vehicle overshoot the commanded depth before the possible occurrence of this failure is noticed; therefore, the time required to detect this failure depends on the difference between the desired depth and commanded depth at the time the failure occurs. Failure of the planes at zero deflection took 45 s to detect. This is due to the fact that it takes time for the hydrodynamic forces to act on the vehicle to level it out such that the characteristics of zero plane deflection are evident; mainly zero depth rate. The time it takes to detect this failure, after a depth change is commanded, will vary based on the inclination at the time of failure. Inclinometer failure at zero will take from 6 to 10 s to detect. This detection time depends on the current inclination at the time of the failure. It will take longer to detect if the inclination is small, because it may be within the noise band of the sensor and some of the aberrant readings will be classified as noise. Angle of attack sensor failure at zero took 6 to 10 s to detect for the same reasons as the inclinometer. Detecting opposite–direction failures of the depth sensor takes 10 s: 4 s for it to affect the calculated slope and 6 s until the failure is reported. Detecting depth sensor failure at zero or full requires only 6 s to be reported. Detecting depth sensor failure at constant depth took about 10 s to detect. This is because it takes almost 5 s for the erroneous depth values to alter the calculated slope, and 6 s of failures before it will be reported. Some of these failures are unlikely to ever be reported due to the nature of the sensors being used. For instance, if the depth sensor fails to zero when the vehicle is at 100 m, the rule that detects steps will report a failure before the required six occurrences for reporting depth gage failure at zero.

As with the depth control failures, the heading control failures were tested to determine whether the failures were detected and to determine the amount of time needed to detect each failure. The spikes in the heading angle, angle of sideslip and rudder position indicator were detected at the time they occurred. Angle of sideslip failure at maximum was reported 6 s after it occurred. Angle of sideslip failure at zero was also reported 6 s after it occurred. Magnetic compass failures held no surprises, being detected after 6 s. Although, if it fails by giving random output, apparent aberrations in other systems would intermittently show up, but none would occur more than the required six times and were (correctly) not reported as failures. Rudder failures at current or full deflection took some time to be detected, 48 s plus the required 6 s for confirmation before reporting. Failure detection took that amount of time because of the heading change required to pass through the desired heading. This failure was set up to occur while the vehicle was turning; therefore, it needed to pass through its desired heading before the possible failure of the rudder was noticed. That 48 s could be higher or lower depending on where in the turn the vehicle is when the failure occurs. Angle of sideslip failure at the current position required only 6 s to detect and report. The rudder failure at zero deflection needed some time for the vehicle to stop turning, 12 s, before the failure was detected. The 12 s to stop turning is dependent upon the turning rate at the time of the failure but 12 s is near the highest time required. If the rudder failure occurs while the vehicle is traveling straight, the failure cannot be detected until the controller commands a course change.

Testing of the System

Aside from the tests to determine how long it would take to detect a failure of a certain component, tests were run to ensure that nominal runs would not give false reports of failures. Four test scenarios were made for this purpose:

1: 180 degree heading change with no depth change
2: 300 m dive with no heading change
3: 300 m dive with 180 degree heading change
4: 300 m dive with 180 degree heading change, after 60 s command original heading and depth

These four cases represent the extreme maneuver changes that the vehicle will undergo during a mission. The first three runs were allowed to continue well after reaching the desired heading and depth, to test the system's performance during constant depth and heading. No false reports of failures were made during these runs.

The first three test scenarios were run repeatedly with failures introduced to ensure that once the system was integrated, failures of one component did not cause false reports of failures of another component. All of the imposed failures were detected during these tests and there was no false reporting of failures.

CONCLUSIONS

Failures were promptly detected when the vehicle was changing depth, heading or both at the time of failure. During a straight and level run, certain failures did not manifest themselves until maneuvering was attempted, because many of the instruments are at the zero position and a failure at that position is sometimes not discernable from noise. As soon as a change from straight and level was commanded, the failures were detected and reported.

The rule–based language CLIPS worked well as a development tool for this diagnoser. The rule structure readily expressed the conditions and conclusions that the redundancy management technique dictated.

This diagnoser handles one part of the task of failure monitoring in the AUV. When integrated with the rest of the AUVC, it will report its conclusions to the global diagnoser for recovery action.

REFERENCES

[1] Chris Culbert, *CLIPS Reference Manual*, Artificial Intelligence Section Johnson Space Center, April 1988.

[2] Joseph C. Giarrantano, *CLIPS Users Guide*, Artificial Intelligence Section Johnson Space Center, September 1987.

[3] R. Isermann, "Process fault diagnosis based on process model knowledge – part 1: principles for fault diagnosis with parameter estimation," *ASME Journal of Dynamic Systems, Measurement and Control*, pp. 620–626, December 1991.

[4] Asok Ray, "A redundancy management procedure for fault detection and isolation," *ASME Journal of Dynamic Systems, Measurement and Control*, pp. 248–253, September 1986.