

# Fault Diagnosis on Autonomous Robotic Vehicles with RECOVERY: An Integrated Heterogeneous-Knowledge Approach

Kelvin Hamilton, Dave Lane, Nick Taylor and Keith Brown

Ocean Systems Laboratory,  
Department of Computing and Electrical Engineering, Heriot-Watt University, Edinburgh, UK  
Email: Kelvin(dml, nick, keb)@cee.hw.ac.uk Web: [www.cee.hw.ac.uk/oceans](http://www.cee.hw.ac.uk/oceans)

## Abstract

The need for embedding fault diagnosis into goal-orientated autonomous robotic vehicles for increased mission robustness is described. The RECOVERY system, a method for increasing the diagnostic capability by integrating commonly available heterogeneous knowledge is presented. Initial real-water results using the Ocean Systems Laboratory's RAUVER vehicle are given.

## Keywords

Model Based Diagnosis, Autonomous Underwater Vehicles, Subsea, Heterogeneous Knowledge, Robot Vehicles, Mission Robustness

## 1. Introduction

An underlying problem with current generation autonomous robotic vehicles is frailty in the face of unexpected events, such as component failure or environmental interaction. Most mission controllers have specifically coded routines to cope with such events, but for this to be effective all such events must be predicted by the mission programmer. This is clearly impractical. Embedded planners provide the capability to re-plan the mission but their effectiveness varies with the quality and scope of information provided to them. Further, they are infeasible with today's level of embedded computing platform due to the large amount of processing required.

A powerful method of fault diagnosis is Model Based Diagnosis (MBD). This is known as a second-generation diagnosis system [1], where first generation is the familiar rule-based expert system characterized by the use of if-then rules, often in conjunction with fuzzy logic. A model is created of an aspect of the vehicle, often movement, this can then be used for diagnosis by adjusting model variables until the modelled behavior matches the observed behavior. Each parameter generally represents a component, and so component candidates can be generated.

As the model can be built from 'first principles' this is a powerful technique, although far more expensive than a simple rule-base in terms of computing power. Model Based Diagnosis is used in systems such as those developed by Caccio, Takai, Price and NASA's Remote Agent Experiment [5,9,10,14] but all of them concentrate on looking at a single model type and model viewpoint, commonly vehicle movement.

Other vehicle-based fault diagnosis paradigms exist, some use Kalman Filters and other matrix-based analysis methods to detect and diagnose faults e.g. Perrault [4] and Yang [7]. INDOS [3] and Deuker [6] use neuro-symbolic methods.

What is lacking in all of the above is a modular method of integrating the various types of knowledge available on a robotic vehicle. This paper concentrates on the knowledge contained within a realistic autonomous underwater vehicle (AUV) then considers how the available knowledge can be increased and made useful for fault diagnosis.

In this paper, Section 2 considers available vehicle knowledge and discusses the relationship between possible faults, observations and diagnoses. Section 3 shows how the heterogeneous vehicle knowledge may be integrated into a suitable form for diagnosis. Section 4 shows how generalized expert knowledge may be used to guide the diagnosis procedure and Section 5 presents some real-water results of the module described in Section 4.

A system implementing these ideas, known as RECOVERY, is under development at the Ocean Systems Laboratory, Heriot-Watt University, Edinburgh, UK. For more information visit [www.cee.hw.ac.uk/oceans](http://www.cee.hw.ac.uk/oceans).

## 2. Aspects of Vehicle Knowledge

There are three inter-related spaces involved in fault diagnosis: fault space, observation space and diagnosis space. Each space contains the totality of all possibilities in their respective areas. Fault space contains all faults that can occur on the vehicle, observation space contains all observations from the point of view of the vehicle and diagnosis space contains all possible diagnoses available by using the observation space.




Figure 1: Interrelated Spaces

The size of the fault space, or the range of possible faults that can occur on an AUV, when taken together with the unpredictable nature of the environment, tends towards infinity. It is practically impossible to predict (and plan for) all possible faults. This is borne out by practical experience of AUV operations (and almost every other field of engineering).

Further, the range of events that can be planned for is limited by the scope of observation, where an observation may be gained from a sensor such as a temperature or attitude sensor.

Whilst fault space may tend towards infinity the observation space can be directly controlled by the vehicle designer. The initial size of the observation space is determined by the number of sensors embedded in the vehicle but it can be expanded by combining observations in a similar way to the concept/percept paradigm [13].

The initial size of the diagnosis space is also determined by the number of sensors. In a typical case the ratio of observation space to diagnosis space is 1:1, that is, each sensor is responsible for a specific diagnosis. An example of this is a thermostat placed inside a battery; typically, the thermostat tripping will set a flag somewhere in the mission control software saying 'battery overheating'. Often the ratio will be less than 1:1 because some vehicle sensors do not contribute to fault diagnosis, i.e. vehicle attitude sensors will flag that the vehicle is outside its attitude limits but will not be tied to any components. This poor ratio between observation and diagnosis spaces is a feature of current mission-control technologies.

Ideally, it would be possible to create a model of the vehicle that is correct in every detail, and constructed in such a way that any possible faults could be tracked down by operating on the model. This would mean an effectively perfect model of the vehicle would have to be constructed in software, accurate to possibly sub-atomic level. This level of modelling is outside current capabilities, even if such a model could be constructed the amount of time needed to operate with such a large number of variables would be prohibitive. Available on-vehicle computing platforms require limiting the depth of the model and managing the time available for diagnosis, as in Aldea [1].

An alternative approach to perfect model building is realistic model building. The most common model available to AUV engineers is a vehicle-movement model, often generated at the prototype stage by using towed scale-model techniques. Lately there has been work by Caccia [5] on generating a movement model using the vehicle's own sensors.

A model that covers one aspect of a vehicle, in this case movement, is known as a one-dimensional model, where a dimension is a point of view from which a fault can be seen. Most robotic vehicles have a variety of sensors onboard, using which it is possible to monitor a number of dimensions apart from movement. For instance, temperature sensors monitor the thermal dimension and electric-current sensors monitor the power dimension. Using these multiple

viewpoints it is possible to greatly enhance the diagnosis capability, effectively expanding the diagnosis space without increasing the observation space. By extending this information through time, it is possible to access the temporal dimension, enabling the diagnosis of dynamic faults using a system such as that developed by Shen [2].

There are many types of models available for use in diagnosis. The movement models discussed above tend to be equational, with the level of detail determined by the amount of design effort applied to the modelling process. Other types of model include associations, procedures and black boxes; for a guide to the selection of suitable model based diagnosis techniques see Chantler *et al* [15]. It is possible not only to generate, but also to monitor and use a heterogeneous, multi-dimensional model of a vehicle for fault diagnosis.

Other types of knowledge are inherently available on the vehicle, whilst *a priori* information can be provided pre-mission and modified in mission. The available vehicle knowledge breaks down into five main areas: design, sensor, historical, mission and fault. One of the most vital types of knowledge for in-vehicle fault diagnosis is the design knowledge, or the knowledge that tells the vehicle that it is a vehicle. Research into this area is ongoing; but at present it consists mainly of relationships between components, parameters and sensors. A suggested breakdown is given below:

Design Knowledge: 'I am a vehicle'  
(*a priori*)

Sensor Knowledge: 'What's going on?'  
(Synchronous, Dynamic)  
Internal (Systemic)  
External (Environmental)

Historical Knowledge: 'What's happened?'  
(Asynchronous)  
Sensor logs  
Fault logs

Mission Knowledge: 'What am I to do?'  
(*a priori*/Dynamic)

Fault Knowledge: 'What's wrong?'  
(Various)

Much of this knowledge is automatically generated as part of the vehicle design process, for instance circuit diagrams, mechanical drawings, data flow diagrams, systems schematics, fault log books, software class diagrams. Yet more is commonly generated as part of the mission process, such as parameter history which is often recorded as standard.

There is a growing movement for Failure Modes and Effects Analysis (FMEA) to be incorporated into the design cycle. If this is so then a great deal of extra knowledge in

the form of failure models will become available for diagnosis. Work on automating the generation of FMEA knowledge has been performed by Price and Taylor [10].

### 3. Integration of Heterogeneous Knowledge

It is clear that many different types of knowledge exist, or can be usefully added, on an autonomous robotic vehicle. In order for this heterogeneous knowledge to be used for fault diagnosis, some way must be found to link it together and operating methods developed. In particular, the selected methodology should support modularity, machine learning and the ability to use generalized knowledge to aid diagnosis.

Partitioned semantic networks have been chosen as the underlying methodology in the RECOVERY system, modified to reflect the change from natural language to robotic applications. In this case, the nodes can be of many different forms, ranging from equational models to component-specific diagnostic advice. Similarly, the links represent the form of relationship between varying objects, providing a basis for the operational methodology.

Many different forms of objects and links can be specified, with extra types being added as needed to reflect changes and additions to the vehicle's systems. An initial breakdown is suggested below:

Nodes:

- Operational Models
- Components
- Parameters

Links:

- Sensor Equivalence
- Sensor to Parameter
- Component to Parameter
- Transmission
- Power Supply
- Has

### 4. Operating with Generalized Knowledge

An important attribute is the ability to use generalized, domain-independent fault diagnosis knowledge. This knowledge, primarily gathered from human experts, can then be used to guide the diagnosis and model selection process. This realistic approach saves a lot of development time compared with getting the machine to learn these guidelines itself.

The need for diagnosis guidance results from the large solution space generated by the integration of the many and varied types of information. Each model takes time to evaluate, the more detailed models will take substantial time. Dynamic fault models, which are concerned with describing the behavior of a faulty system over time, can take a great deal of processor time to look over a mission log file that could easily be several hundred megabytes long.

Much time can be saved by embedding human

diagnosis knowledge and using this to steer the diagnosis procedure, either by eliminating some components and their associated information or by highlighting others. There are well known methods for the elicitation of knowledge from human experts developed for first generation rule based expert systems [11]. In future, it may be possible to get the robot to learn and enhance diagnosis techniques by itself, but at present this is impractical.

In RECOVERY, one of the main modules that use such knowledge is the correlation detector, or Correlator. At present, this module uses the following abstract knowledge:

1. Parameters that track faulty parameters are likely to be related to the fault.
2. Faults that occur at similar times are likely to be related to the fault.
3. Components that become active (after a period of inactivity) just before the occurrence of a fault are likely to be related.
4. Components that are being used at the time of a fault are more likely to be related to the fault than inactive components.

When a fault occurs the Correlator looks back through the mission log file (temporal dimension) and looks for the correlations described above. If correlations are found between various components and parameters this information is used to modify the semantic network described in section (3). This modification effectively highlights suspicious components or parameters, which are then used to guide the diagnosis engines towards sections of knowledge thought to be most relevant.

An example of this is thermal drift of sensor readings. If a wheeled robot had a speed sensor and a maximum set speed and the measured speed exceeded maximum speed then a snapshot of the system would show only that the robot was travelling too fast. By looking back at the history of the vehicle, it could be seen that the maximum speed of the robot was creeping up at approximately the same rate as the temperature. A correlation of this kind would be extremely useful in diagnosing the actual fault.

This is, perhaps, a rather simplistic example – the robot is assumed to be travelling constantly at maximum speed, which is unlikely to be the case. Even so, the concept is valid as the correlator would have noticed the association between temperature and sensor drift without this knowledge being explicitly represented.

Note that the correlator can effectively generate new semantic links between nodes. This is a powerful ability, the use of which is still under investigation.

### 5. Results

Initial trials on some modules have been successfully

completed in real-water using the Ocean Systems Laboratory's robotic research vessel RAUVER. Rauver is a 2m catamaran robotic submersible with both remote control and full autonomous capability. It is used as a sensor platform, evaluation testbed and support vehicle by the Ocean Systems Laboratory.




Figure 2: RAUVER descending towards the lake bed whilst performing a RECOVERY evaluation mission.

A typical survey and intervention mission scenario has been designed to test RECOVERY operation under realistic conditions. The scenario is that of surveying an object, such as an oil-platform leg, reservoir intake tower or sunken vessel, which must be accomplished under certain constraints. On completion of the survey the vehicle must land on the seabed and take a sample, then return to the surface in a constrained manner. The mission, a temporal sequence of goals, is as follows:

- Goal 0: Surface and hover for GPS fix.
- Goal 1: Traverse to mission start point.
- Goal 2: Dive vertically to survey-end depth.
- Goal 3: Hover to ascertain suitability of landing spot.
- Goal 4: Land gently on seabed.
- Goal 5: Standby during seabed sample.
- Goal 6: Ascend vertically to surface.
- Goal 7: Hover at surface until standby.

A lift-thruster fault is forced at the beginning of Goal 6 when the vehicle tries to lift off, causing a severe roll, breaking constraints and generating a discrepancy between modelled and observed vehicle movement. In the full RECOVERY system this will trigger the diagnostic process, but this paper concentrates on the action of the Correlator module.

Graphs showing relevant vehicle parameters during normal and faulty operation are shown below. Two parameters are displayed from power and movement dimensions respectively, these are the dimensions in which the fault shows up most clearly.




Figure 3 (above): Port Hull Power Consumption during normal operation.




Figure 4 (above). Port Hull Power Consumption during faulty operation. Note that the loss of one thruster has removed the second zone of high power consumption, this is the fault showing in the power dimension.




Figure 5 (above): Vehicle Roll during normal operation.




Figure 6 (above): Vehicle Roll during faulty operation. Note the large roll caused by the loss of a lift thruster, this is the fault showing up in the movement dimension.

During normal operation the vehicle's parameters, specifically power and roll, stay within the mission constraints. In the faulty mission, where one of the lift thrusters fails, the power shows a discrepancy and the roll violates its constraint. In a conventional vehicle with an observation to diagnosis ratio of less than 1:1 there would be no explicit diagnosis capability using this information. No component diagnosis would be tied to this broken constraint, all that the mission controller would see would be a broken constraint and the mission would probably be aborted. By using a multidimensional model based diagnosis system such as RECOVERY it is possible to diagnose not only components with directly related sensors but also indirectly sensed sub-components such as thrusters.

The Correlator looks through the mission history for correlations that will increase diagnostic confidence. When run on the mission data shown above it produced the following output:

RECOVERY: Correlator Output

"78.91" No Correlation1  
 "78.91" No Correlation2  
 "78.91" Got Correlation3  
 PORT\_LIFT\_THRUST\_CONTROL  
 STBD\_LIFT\_THRUST\_CONTROL  
 PORT\_LIFT\_THRUSTER  
 STBD\_LIFT\_THRUSTER  
 PORT\_LIFT\_THRUST\_BRUSH  
 STBD\_LIFT\_THRUST\_BRUSH  
 PORT\_LIFT\_GBOX  
 STBD\_LIFT\_GBOX

Figure 7: Correlator Output

In Figure 7, "78.91" is a timestamp generated by the vehicle showing the time at which the fault becomes apparent, in this case from the extreme roll of the vehicle.

Correlations one, two and three refer to the generalized diagnostic guidelines discussed in section (4). Correlation3, 'components that become active (after a period of inactivity) just before a fault' has been found, and all relevant components have been listed. The output from the Correlator shows both strongly relevant and weakly relevant associations. For this particular fault the Port Lift Thruster components are relevant because this is the failed component, although the fault diagnosis system is unaware of this. The Starboard Lift Thruster components are less relevant because this thruster is actually working. Information from other fault diagnosis systems will be combined to arrive at a final diagnosis.

The relevance problem is generally true for all types of correlations and so a single type of correlation has limited diagnostic usefulness. The strength of the correlator comes from combining the various detected correlations and using this to strengthen or weaken the confidence of a particular

diagnosis. This method will be the subject of an Ocean Systems Laboratory paper due early in 2001.

## 6. Conclusions

Work has been presented that shows the necessity of increasing the diagnosis space in order to increase the robustness of autonomous goal orientated robotic vehicles. Main points are:

The size of the fault space tends towards infinity.

The size of the observation space can be directly controlled by the vehicle designer.

The size of the diagnosis space is dependent on the observation space and can be significantly increased by the use of model based diagnosis techniques.

The size of the diagnosis space can be further and significantly increased by the integration of heterogeneous, multidimensional knowledge.

A lot of suitable knowledge either exists on the vehicle or is generated as a normal part of the design process.

Generalized human fault diagnosis heuristics can be used to guide the diagnosis process.

Initial real-water results have been presented from one of the RECOVERY modules, showing that generalized human knowledge can generate useful information. RECOVERY modules capable of using this information are under development at the Ocean Systems Laboratory.

## Acknowledgments

The UK Engineering and Physical Sciences Research Council (EPSRC) funded this research. East of Scotland Water has been extremely helpful in making available their large selection of reservoirs and lochs for testing. The United States Office of Naval Research financed visits to Florida Atlantic University, Woods Hole Oceanographic Institute, the Autonomous Undersea Systems Institute and the Unmanned Undersea Vehicles Symposium 2000, during which many useful discussions took place.

Ian Chalmers, Ron Lynch and Dave Haldane were responsible for making the mechanical components of the RAUVER vehicle, whilst Len McLean built, modified, maintains and upgrades RAUVER systems. Without their high level of skill and effort this research would be purely theoretical. Thanks also to the whole of the Ocean Systems Laboratory team for invaluable contributions and suggestions.

## References

1. "The Use of Scheduling and Hierarchical Modelling Techniques for Time-Limited Model-Based Diagnosis" A. Aldea, PhD Thesis, 1994, Heriot-Watt University, Edinburgh.
2. "Diagnosing Continuous Systems with Qualitative Dynamic Models", Qiang Shen and Roy Leitch, Intelligent Systems Laboratory, Heriot-Watt University, UK, 1995, Artificial Intelligence in Engineering (9)
3. "Intelligent AUV On-Board Health Monitoring Software (INDOS)", Willi Hornfeld and Eberhard Frenzel, STN Atlas Elektronik, Germany, 1998, ICRA
4. "Fault-Tolerant Control of an Autonomous Underwater Vehicle", Douglas Perrault and Meyer Nahon, Space and Subsea Laboratory, Canada, 1998 ICRA
5. "Modelling and Identification of Open-Frame Variable Configuration Unmanned Underwater Vehicles", Massimo Caccia, Giovanni Indiveri and Gianmarco Veruggio, 2000, IEEE Journal of Oceanic Engineering
6. "Fault Diagnosis of Subsea Robots Using Neuro-Symbolic Hybrid Systems", B.Deuker, M.Perrier and B.Amy, IFREMER, France, 1998 ICRA
7. "Experimental Study of Fault-Tolerant System Design for Underwater Robots", Keith Yang, J.Yuh and S.K.Choi, Autonomous Systems Laboratory, University of Hawaii, Honolulu, 1998 ICRA
8. "Improving Fault Handling in Marine Vehicle Course Keeping Systems", Zoran Vukic, Henrieta Ozbolt and Dean Pavlekovic, University of Zagreb, Croatia, 1999 IEEE Robotics and Automation Magazine
9. "Development of a System to Diagnose Autonomous Underwater Vehicles", Motoyuki Takai and Tamaki Ura, University of Tokyo, Japan, 1999 International Journal of Systems Science
10. "Multiple Fault Diagnosis from FMEA", Chris Price and Neil Taylor, University of Wales, Aberystwyth, UK
11. "Artificial Intelligence, A Modern Approach", Stuart Russell and Peter Norvig, 1995, Prentice Hall International, ISBN 0-13-360124-2
12. "An Introduction to Artificial Intelligence", Janet Finlay and Alan Dix, 1996, UCI Press, ISBN 1-85728-399-6
13. "Reactive Task Planning for Unstructured Environments", P. Knightsbridge, 1997, PhD Thesis, Heriot-Watt University, Edinburgh, UK
14. "Design of the Remote Agent Experiment for Spacecraft Autonomy", Bernard et al, JPL/ NASA Ames Research Laboratory, USA, 1998, IEEE Conference on Aerospace.
15. "Selecting Tools and Techniques for Model Based Diagnosis", M.Chantler, G.Coghill, Q.Shen and R.Leitch, Heriot-Watt University, Edinburgh, UK, 1998, Artificial Intelligence in Engineering (12)