

On the Reliability of the Autosub Autonomous Underwater Vehicle

G. GRIFFITHS, N. W. MILLARD, S. D. McPHAIL, P. STEVENSON and P. G. CHALLENGOR

Southampton Oceanography Centre, Empress Dock, Southampton, SO14 3ZH, UK

Submitted to Underwater Technology, July 2001.

Abstract

As autonomous underwater vehicles (AUVs) enter operational service an assessment of their reliability is timely. Using the Autosub AUV as an example, several design issues affecting reliability are discussed, followed by an analysis of recorded faults. Perhaps contrary to expectations, failures rarely involved the autonomous nature of the vehicle. Rather, faults were typical of those that occur with any complex item of marine electromechanical equipment. A statistical analysis showed that the failure rate decreased with distance travelled - an indicator that an AUV underway, submerged, is at less risk of a fault developing than during other phases of a mission.

1. Introduction

Reliability is a prime requirement for an autonomous underwater vehicle (AUV). As a self-contained underwater robot, a submerged AUV has very limited communication bandwidth and range available to it to signal an emergency. Furthermore, contemporary sensor and software systems within AUVs only provide elementary reactions to foreseen problems. As for unforeseen problems, the present generation of operational vehicles has almost no ability to determine and execute an optimum response. As a consequence of this state of affairs, reliability is paramount in the engineering and the human aspects of

the design and operation of AUVs. Yet, the subject has received little attention in the open literature, an exception being an insightful qualitative account of mission-level reliability in the REMUS AUV [1].

For three key reasons, achieving operational reliability has always taken centre stage within the Autosub AUV project. First, with only one vehicle to serve as prototype, demonstrator and operational platform the onus was on the development team to design and construct a robust and dependable vehicle from the outset. Second, Autosub's prime purpose is to undertake science missions; marine scientists need to be convinced that the vehicle can be relied upon to deliver their data. Users naturally apply a credibility threshold to new technology, and an important factor in crossing that threshold is demonstrating reliable operation time after time, without incurring excessive engineering downtime. Third, the AUV needs to demonstrate that it provides value for money to funding agencies and users. Consequently, its short-term operational record and long-term likely working life are important issues. It is perhaps inevitable that any autonomous vehicle may be lost before the end of its useful life, but with care, good design and operational practice the working life of the vehicle should be long and productive.

The structure of the paper is as follows: Section 2 outlines how Autosub was designed with reliability in mind and how operational practice complements that design philosophy. Section 3 identifies, quantifies and classifies the faults that have occurred during four years of use on 240 missions. Section 4 moves on from discussing faults and provides quantitative estimates of the reliability of the vehicle and the probability of loss. The paper concludes with a discussion on the lessons learnt that have most bearing on the design and operation of the next generation of AUVs.

2. Autosub and reliability

2.1 How reliability entered the design process

The design philosophy for Autosub was based firmly on risk reduction. Indeed, the first phase of the project in 1989 comprised 26 work packages that assessed the feasibility and risks of all aspects of the project. Topics covered included support systems such as launch and recovery; legal and procedural issues; data handling; science mission scenarios as well as the technical issues of propulsion, mission control, buoyancy etc.

From 1990-1995 the second phase of the project addressed those tasks and sub-assemblies identified in Phase I as needing the greatest development effort. In several instances, including navigation, propulsion and mission control, the need was for risk reduction. This risk reduction took several different forms. For example: improving the reliability of the propulsion system by removing the need for a shaft seal on the motor [2]; ensuring that the vehicle was able to acquire a GPS fix irrespective of surface wave conditions required developing a GPS receiver with an innovative method of assembling a fragmented ephemeris message [3]; achieving reliability in the vehicle's extensive real-time software (*ca.* 18,100 lines of real-time code, not including embedded software in bought-in systems), off-line planning, compilation and monitoring software (18,000 lines of code) and data processing software (6,000 lines) required setting and adhering to standards [4] as well as making the most appropriate choice of network and software architecture [5]. Finally, mechanical reliability and avoidance of systems integration conflicts were assured through adhering to a thorough design and construction procedure [6,7]. Underlying this approach was the need to detect and correct reliability problems at the earliest stage.

This pragmatic approach to risk reduction and achieving operational reliability did not call for rigorous failure analyses of off-the-shelf components such as electronics hardware, documented in reliability manuals such as US MIL-HDBK-217 [8]. However, the mean time between failures (MTBF) of sub-assemblies was considered at the design stage, e.g. the 50,000 hour MTBF for the critical motion reference unit. Quite simply, the programme budget was not sufficient to delve into the detail of the 'insignificant many' at the expense of dealing thoroughly with the 'significant few'. In following this Pareto principle the development team used their considerable experience of previous marine technology projects, where individual component failures were rarely the cause of reliability problems.

2.2 The role of testing

A substantial fraction of the effort in Phase II of the programme was concerned with testing sub-systems prior to constructing the vehicle. The main sub-systems that benefited from testing before the vehicle construction began in earnest were:

- **Main propulsion motor.** The characteristics and reliability of the motor were established using a purpose-built dynamometer operating within a 600 bar pressure test chamber [2]. The tests showed that the plain ceramic bearings were too brittle, which resulted in unnecessarily high friction losses, neither did they run quietly or smoothly. They were subsequently replaced with a Xylan-coated Aluminium Bronze bearing running against 625 Inconel [9]. Tests indicated an estimated bearing life of 600 hours and showed that periodic cleaning was necessary to deal with the light corrosion [6]. Testing also showed that, when the bearings failed, they failed in a way that would not jeopardise a mission – the motor would remain in a working state [9].

- **Fin actuators.** The critical components within the pressure compensated fin actuators were tested at 700 bar for a total of 110,000 cycles and then left in the compensating oil for 4,500 hours to test for long-term deterioration. These tests showed that the units did not suffer from noise, intermittence or observable deterioration over that period [9].
- **GPS receiver.** Sea trials of the interruption-tolerant GPS receiver were made aboard a full size model of the Autosub vehicle to validate its performance with realistic wave washover.
- **Vehicle hydrodynamic coefficients.** An extensive series of tank tests on a $3/4$ scale model provided the hydrodynamic coefficients for use in the vehicle control system [10].
- **LonWorks nodes.** Testing of individual LonWorks nodes forming the vehicle command and control network took place as they were constructed. Data flows from any nodes not built were simulated by a development system. Partitioning the software into relatively small and simple nodes, each with a single well-defined network interface was a distinct advantage during testing and during subsequent maintenance.

The plan for the post-construction testing of the vehicle called for a series of trials, of increasing difficulty, spanning 11 months, starting in a laboratory tank in May 1996 ending with the vehicle's formal acceptance trials in Scottish coastal waters in April 1997 [7]. Tests in the laboratory tank enabled system integration issues such as ground loops and interconnection errors to be resolved. Such errors were few. Within a month the vehicle was used in a remotely operated mode within the confines of Empress Dock,

Southampton. This 400 m long by 9 m deep dock, immediately adjacent to the Autosub laboratory, proved an excellent test facility. On the very first day of operations in the dock seven surface runs were made to test the heading controller, followed by two dives to 5 m as a test of the depth controller. The next day twelve sub-surface runs were completed with the acoustic tracking system also operating. The major problem was handling the umbilical from a rigid inflatable chase boat.

More elaborate test missions were performed during two campaigns in Portland Harbour (July and November 1996) and one at Oban (April 1997). In all, 88 test missions during these campaigns exercised a number of different vehicle behaviours including: surface navigation; sub-surface dead-reckoning navigation based on time and on waypoints; transects at constant depth and at constant altitude; forward and backward diving; see-saw profiling to pre-set depth and pre-set height above bottom; leaving harbour autonomously; DGPS surface navigation; and event-driven track changes e.g. on reaching a pre-set minimum water depth. The overarching objective for these trials was to establish the robustness of the vehicle and its control systems when carrying out a wide variety of mission segments representative of those that would be carried out during subsequent science campaigns.

2.3 Operational procedures and reliability

Planning for the safe return of an AUV begins long before launch. What the user wants may be at variance with what the more experienced engineer/operator is willing to allow. The mission plan, jointly formulated by the user and operator, must try to satisfy the user requirement but must also be within the bounds of a strategy to minimise hazards and establish a safe, properly considered rendezvous for recovery. Hazards may include

excessively rough topography during terrain following missions, strong currents, ice, shipping and fixed structures.

As the scheduled launch time approaches a weather eye must be kept, remembering that launching in rough weather is generally much easier than recovery. Launching on a deteriorating weather forecast could be fool hardy. Shortly before launch a predefined checklist should be completed with no concessions made to non-compliance without very careful consideration of possible consequences.

Once the mission is started, a short-range acoustic link with the vehicle is monitored to ensure that the initial stages of the mission are adhered to. If the AUV's behaviour is not as expected, the support ship may stop tracking. If the behaviour causes concern, the mission should be aborted and if the problem is not diagnosed and fixed through the radio modem link, the vehicle should be recovered.

3. Identification and classification of faults

The mission log sheets for Autosub contain details on all of the problems and faults discovered. The information on these log sheets formed the basis for the list of faults discussed in this section.

The faults were clustered under thirteen headings. Some faults have multiple entries, for example, when a 'Human Error' gives rise to a 'Collision with Vessel'. The Pareto plot, Figure 1, shows that most faults were the result of human error.

Of course, this form of presentation does not show the relative impact of the faults or if the profile of faults changed over the tremendous range of missions (from 90 s duration in Empress Dock to 53 hours on the edge of a hurricane off Bermuda). Neither does the Pareto plot take any account of the changes to the vehicle systems over a 4-year period.

For example, some 30 to 40% of the embedded real-time software has (gradually) changed from June 1996 to June 2000, primarily to add extra functions.

3.1 Time history of Autosub reliability

It is instructive to examine the fault history of Autosub over the following six groups of missions, summarised in Table 1.

Group A comprised the first missions in Empress Dock as an ROV in June 1996; ROV and AUV runs in Portland Harbour in July 1996 and Empress Dock trials in July through October 1996. In all 55 missions covered 34 km.

Group B comprised Portland missions in November 1996; Empress Dock and Oban missions in April 1997 (the acceptance trials). In all 33 missions covering 199 km.

Group C comprised Oban trials in October 1997 and a Florida campaign in December 1997. In all 34 missions covering 242 km.

Group D comprised a Portland campaign in March 1998; Bermuda in August-September 1998 and Oban in April-May 1999. In all 57 missions covering 616 km.

Group E comprised science mission campaigns on FRV *Scotia* in the northern North Sea in July 1999, on RRS *Challenger* in the southern North Sea in August 1999 and at Oban in November 1999. In all 37 missions covering 1034 km.

Group F comprised science missions on MV *Tershelling* at three sites on the west coast of Scotland in March and April 2000, equipment trials at Plymouth and a science campaign in the Strait of Sicily in June 2000. In all 24 missions covering 869 km.

By assigning the recorded faults to these mission groups, retaining any instances of multiple faults on a mission, the average number of faults per mission and faults per km travelled have been calculated, Table 1.

Figure 2 is a diagram that shows faults per km against faults per mission for the six groups. The arrows show the time progression. Group A missions were very simple, many were with the vehicle configured as an ROV, while the autonomous missions were very straightforward. The average number of faults per mission was very low, probably because the vehicle behaviours that were under test were only one step removed from the laboratory tests. The average length of the mission script was 3 lines. The complexity of the tasks given to the vehicle increased significantly for Group B. Mission scripts were up to 97 lines on the 6.5 hour mission 78. Also, the team generally worked away from the laboratory and missions were on average ten times longer than during Group A. While the number of faults per km remained steady, the average number of faults per mission increased. In part, this increase was expected. The purpose of the engineering trials was to discover and rectify faults. For Group C, while the average number of faults per mission and faults per km showed a further increase, the *rate* of increase had decreased. Given that this group of missions included the first campaign outside of the UK and the first campaign in oceanic rather than coastal waters, this was a credible achievement even if disappointing. Faults per mission and per km dropped significantly with Group D, and dropped again for the missions in Group E as further experience was gained, previous faults were repaired, checklists and operational procedures were improved and automated launch and recovery was introduced. The growth in reliability of the vehicle was not continued into group F. The detailed logs show that the newly introduced recovery-float deployment system had several problems. The emergency abort system also triggered several spontaneous weight drops. The last two missions in this group suffered from collision with the seabed in the Strait of Sicily, both involved terrain following. While the

vehicle was able to continue on its way without any intervention after mission 239, on mission 240 it collided with a cliff and became trapped under an overhang.*

3.3 Impact of faults

Until now, faults have been considered irrespective of their impact on the operational performance of the vehicle. While some faults may make no significant impact (e.g. omission of the 'surface' command at the end of mission 121) others could lead to the loss of the vehicle (e.g. collision, major pressure vessel leak, battery short circuit). In between these two extremes lies a spectrum of the impact of a fault on the operational performance of the AUV. Table 2 splits the impact of faults into twelve categories and shows the number of arisings in each category for the six mission groups A to F. Some faults lead to more than one entry in the impact table, e.g. a mission abort also leads to a direct cost (the loss of the mission abort weight) and a minor delay. It could also mean that the mission objectives were not met, or were only met in part. There is no evidence that the fewer faults in mission groups D and E were more serious than in earlier groups.

From Table 2 we see that 26 faults were classified as High Impact. These 26 faults occurred on 19 different missions. On the past track record, therefore, the probability of a mission suffering one or more high impact faults was $19/240$ or 0.079. Of these 19 missions, the faults on eight occurred either on deployment, during pre-dive checks or on the first dive. On eleven missions the faults occurred when the vehicle was underway. Therefore, the probability of a High Impact fault occurring when underway was 0.046 or once every 22 missions.

* The vehicle was later recovered using an ROV.

3.4 Building a statistical model of Autosub faults

By fitting parametric models of reliability to Autosub fault data it is possible to produce formulae that can be used to predict the probability of a fault occurring on any mission in terms of distance travelled and the operating regime (for example, either mid-water or terrain-following).

For each Autosub mission the data consists of the distance travelled, the time taken (although this was not used in the analysis) and whether a fault developed on the mission. Those that did not have faults were important, as they comprised the vast majority of missions and provide information about failures. Although not as useful as the actual distance to failure, such information is valuable particular for the longer distances. The missions were divided into the same six groups discussed earlier and were analysed by group and as a whole.

Parametric models of reliability use censored* and uncensored data to fit a probability distribution and use this to estimate the probability of a fault occurring over any distance [11]. Six different distributions were fitted to the data: extreme value, normal, logistic and their logarithmic versions: Weibull, log normal and loglogistic. It was clear that the log distributions fitted the data much better than their linear counterparts. The Weibull distribution showed the highest failure rate, and, to err on the conservative, was chosen as the preferred parametric model.

There are two important functions that can be calculated from a parametric model of survival. One is the distribution function. This is the probability that there will have been

* A censored observation is one where the mission ended without a fault developing.

a fault before a certain distance has been reached, i.e. if the distance to a fault is represented by X the distribution function is

$$F(x) = P(X < x)$$

A related function is the survivor function, which is the probability that a fault has not occurred by x . This is related to the distribution function by

$$R(x) = P(X > x) = 1 - P(X < x) = 1 - F(x)$$

The other important function is the hazard rate. This is the probability that a fault occurs in the next instant given that a fault has not occurred up to distance x . This is given by

$$\lambda(x) = \frac{f(x)}{R(x)} = \frac{f(x)}{1 - F(x)}$$

where $f(x)$ is the probability density function

$$f(x) = \frac{d}{dx} F(x)$$

Fitting the Weibull distribution to each of the six groups and the entire data set produced the parameter estimates in Table 3. Note that no estimates could be produced for Group A as there were only two missions in this group where actual failures were recorded, the other fifty-three data points were censored. However, Group A data are included in the 'All missions' group. Figure 3 shows the probability of a mission ending without a fault against distance travelled for the 'All missions' group.

For the Weibull distribution, α is the scale parameter or characteristic life - the life at which there is a 63.2% probability of a fault having developed. The longest characteristic life was achieved in Group D. This group contained the longest mission to date (263 km) and it did not end due to a failure. As a consequence, this censored data point had the effect of increasing the characteristic life of the group. Except for the effect of this single mission on Group D, the characteristic life, in terms of distance travelled, increased

steadily from Group B to Group F. That is, the reliability of the Autosub improved through the team gaining experience and correcting earlier faults.

Now consider the hazard rate. This gives the probability that Autosub will fail in the next ϵ km given that it has not failed in the previous x km. A constant hazard rate would imply that the probability of failure was constant with distance, an increasing hazard rate would imply that Autosub 'aged', i.e. it became more prone to failure as the distance travelled increased. On the other hand a decreasing hazard rate would imply the opposite, that the vehicle became more reliable the greater distance travelled. When the shape parameter β of the Weibull distribution is less than one the hazard rate decreases with distance travelled (time). Conversely, when β is greater than one the hazard rate increases with distance (time). For a Weibull distribution, the hazard rate as a function of distance can be expressed as:

$$\lambda(x) = \frac{\beta}{\alpha} \left(\frac{x}{\alpha} \right)^{\beta-1}$$

The most robust 'All missions' group showed a decreasing hazard rate, irrespective of the statistical model chosen, Figure 4.

The decreasing hazard rate with distance travelled suggests that it would be beneficial to track and monitor the vehicle for the first part of any unescorted mission. The reduction in risk is easily calculated. If the distribution function is given by $F(x)$ then the probability that X is greater than x given that it is greater than y is given by:

$$P(X < x | X > y) = \frac{F(x) - F(y)}{1 - F(y)}$$

where $x > y$. If, for example, no fault occurred during the first 10 km of a 150 km mission, the probability of a fault occurring on the remainder of the mission is reduced from 0.40

to 0.35, a 13% reduction in risk for a penalty of monitoring the vehicle for 7% of its track. In practice, the reduction in risk comes from the opportunity to recall the vehicle if a fault develops while it is being tracked. This requires the vehicle to telemeter enough information to the support ship for a decision on whether the vehicle is carrying out its mission correctly. Achieving the reduction in risk also assumes that the vehicle can be recalled and that it can be recovered.

4. Survivability analysis

Losing an AUV in the open ocean is a rare event. Most vehicles have visual, radio, sonar and satellite tracking systems that enable the operator to find a vehicle if it is lost on the surface or trapped subsurface. For the vast majority of faults, the emergency abort system senses problems and takes appropriate action. As a consequence, for open ocean work, the main interest is in the statistics of faults, not losses. However, for a vehicle operating in ice-covered waters, or in other inaccessible places, the probability of loss becomes important.

In this section we develop the concept of the probability of losing a vehicle when operating in an ocean with a rigid lid. That is, with the vehicle working in an operating environment where there is no method of recovery should it fail to complete its mission. Such an environment could be under sea ice or under shelf ice.

4.1 Survival probability in an ocean with a rigid lid

First, the survival probability is derived using the simplest statistical model based on a constant probability of loss per mission, independent of mission length. This is clearly an oversimplification, given the analysis in section 3.4, nevertheless it is instructive. Second, the survival probability is derived using a hazard rate varying with distance based on the parameters of a Weibull distribution.

From section 3.3 the probability of a mission suffering a high impact fault when underway was 0.046 over the set of 240 missions undertaken between 1996-2000 (11 high impact faults occurred when underway). These missions covered 2994 km (Table 1), that is, an average mission length of 12.5 km.

On which of these missions might the vehicle have been lost if the ocean had had a rigid lid? An analysis of the faults shows seven such missions (hereafter referred to as Group L7). For the seven probable instances of loss, had the vehicle been operating under a rigid lid:

$$P(loss)=7/240= 0.029$$

If a constant probability of loss for each mission is assumed, as generated by a Poisson process, then an exponential distribution is appropriate and the survival probability is given by:

$$R(x)=\exp(-\lambda x)$$

where λ is the hazard rate per mission and the mean number of missions to loss is $1/\lambda$.

The survival probability on a per mission basis can be calculated from:

$$R(x)=1- 0.029 = 0.971 =\exp(-\lambda \cdot 1)$$

giving $\lambda = 0.029$ or a mean number of missions to loss (MML) of 35.

For the case of the hazard rate varying with distance the loss hazard rate can be calculated by censoring those missions where the faults would not have resulted in loss. That is, by including as uncensored data only those missions where the high impact, underway faults would have led to loss - Group L7 as above. For this group, the parameters of the vehicle loss Weibull distribution were calculated as $\alpha=296.9$ and $\beta=1.208$. The loss hazard rate with distance travelled is shown in Figure 4. It is clear from Figure 4 that the variation with distance travelled of the loss hazard rate is different

from the 'All missions' fault hazard rate. The implication is that the faults leading to loss were not a random sub-sample of the population of all faults, but had a different statistical distribution. With $\beta > 1$ the loss hazard rate increased with distance travelled, contrary to the all missions fault hazard rate. Further analysis showed that the increasing loss hazard rate with distance arose because two of the seven missions considered were long (at mission 239 at 170 km and mission 240 at 50 km). On both of these missions the vehicle collided with the seabed when following the terrain.

While it would be appropriate to use these values of α and β to predict the probability of survival during terrain-following missions, it would be inappropriate to include the faults during missions 239 and 240 when predicting the probability of survival for mid-water missions. Taking these two missions out of group L7, to form group L5, results in $\alpha = 588.8$ and $\beta = 1.0065$ and a hazard rate almost constant with distance, Figure 4.

In January and February 2001 Autosub undertook a science campaign in the Weddell Sea, which included a number of missions under sea ice and under icebergs. The statistical model derived above was used to estimate the likelihood of survival. For these missions the water depth was well in excess of 1000 m and the vehicle was not following the terrain. The most appropriate parameters for a Weibull distribution are therefore those of group L5 ($\alpha = 588.8$ and $\beta = 1.0065$). Six of the Weddell Sea missions involved transects under sea ice and four were under icebergs. Table 4 lists the distance travelled under ice for each of these missions, together with the probability of survival calculated by multiplying the distance travelled under ice by the hazard rate appertaining to that distance. The probability of the vehicle surviving the campaign was simply the product of the probability of surviving each mission. The predicted probability of surviving the

campaign was 0.636. In practice the vehicle did survive, but not without incident. Loss was a distinct possibility, when, on mission 252, the vehicle collided with the side of an iceberg at a depth of 150 m. While only minor damage was sustained, and the vehicle continued on its track, the damage could have been more extensive.

6. Conclusions

Autonomous underwater vehicles are an amalgam of real-time command and control systems, environmental sensors, electromechanical components and an energy source inside one or more pressure vessels. The initial specification and design of an AUV must recognise the strengths and weaknesses of the individual components and their interrelationships. This paper has shown that by good design and thorough testing of the 'significant few' systems that could pose high risk to the vehicle, the overall reliability of the *autonomous* vehicle is not dominated by the complex assemblies needed to provide that autonomy. This is an encouraging conclusion.

Most of the problems encountered with the Autosub vehicle during its 240 missions might well have been encountered in any non-autonomous oceanographic vehicle such as a deep-towed sonar platform or a remotely operated vehicle. Instances of human error were few when tackling difficult algorithm development problems. Rather, the human errors were mostly commonplace, often related to operations rather than design, and not necessarily related to the autonomous nature of the vehicle.

While treating the history of faults with Autosub using reliability statistics has proved useful, although the results are dependent on value judgements. In particular, the predictions for the mean number of missions to loss when operating in an ocean with a rigid lid should be treated with some caution. Unfortunately, this is far from an academic problem. Autosub has already been used under sea ice and icebergs and will be operating

long missions under polar ice shelves from 2002-2004. A full fault effect analysis could help remove the subjective assessment of the faults that could lead to loss, but a fully objective analysis scheme is unlikely, there will remain the need for informed engineering judgement. Another complicating factor is that Autosub is a working prototype. Its configuration is not always stable. It is altered to suit the user needs on each campaign. Changes to systems, or adding new systems, have had the effect of temporarily reducing the reliability.

Future AUV designs should bear in mind the proven reliability of many of the sub-systems, including the mission controller and the distributed network nodes; the propulsion motor; the actuators and control surfaces and the relocation aids. Wherever possible, human actions should be independently checked, for example by simulation. This is especially important in planning missions and in setting vehicle behaviours. New sub-systems should be thoroughly tested before they need to be part of the operational vehicle. All of these are merely examples of good engineering practice. Autonomy in the open ocean is not overly risky; after all, the main risks, and most faults, happen on or before launch or on or after recovery.

Acknowledgements

We thank all of the members of the Autosub team for their dedication and good humour through a decade of research, development and operation. Our thanks also go to all the Masters and crew of the support vessels that we have used over the last four years. Peter Collar and anonymous colleagues at the Woods Hole Oceanographic Institution provided helpful comments on a comprehensive internal report from which this paper has been distilled. This work was funded by the UK Natural Environment Research Council as part of the Autosub Science Missions programme under contract F3/G12/51/01.

References

1. Stokey, R., Austin, T., von Alt, C., Purcell, M., Goldsborough, R., Forrester, N. and Allen, B., 1999. AUV bloopers or why Murphy must have been an optimist. *Proceedings of the 11th International Symposium on Unmanned Untethered Submersible Technology*, New Hampshire, 23-25 August 1999, AUSI, pp. 32-40.
2. Stevenson, P. and Hunter, C. A., 1994, Development of an efficient propulsion motor and driver for use in the deep ocean. *Proceedings IEE Conference on Electronic Engineering in Oceanography*, 19-21 July 1994, Cambridge, UK. Institution of Electrical Engineers, London, Publication No.394, pp. 51-55.
3. Meldrum, D. T. and Haddrell, T., 1994. GPS in autonomous underwater vehicles. *Proceedings of the Sixth International Conference on Electronic Engineering in Oceanography*, 19-21 July 1994, Cambridge, UK. Institution of Electrical Engineers, London, Publication No.394, pp. 11-17.
4. McPhail, S. D., 1993. Development of a navigation system for the Autosub autonomous underwater vehicle. *Proceedings of Oceans '93*, 18-21 October 1993, Victoria, Canada. Institute of Electrical and Electronics Engineers, Piscataway, USA, pp. II: 504-509.
5. McPhail, S.D. and Pebody, M., 1998. Navigation and control of an autonomous underwater vehicle using a distributed, networked, control architecture. *Underwater Technology*, **23**(1), pp. 19-30.
6. Stevenson, P. Graham, D. and Clayson, C.H., 1998. The mechanical design and implementation of an autonomous submersible. *Underwater Technology*, **23**(1), pp. 31-41.
7. Millard, N.W., Griffiths, G., Finnegan, G., McPhail, S.D., Meldrum, D.T., Pebody, M., Perrett, J.R., Stevenson, P. and Webb, A.T., 1998. Versatile autonomous submersibles - the realising and testing of a practical vehicle. *Underwater Technology*, **23**(1), pp. 7-17.
8. US MIL-HDBK-217: *Reliability prediction for electronic systems*. Available from the National Technical Information Service, Springfield, Virginia.
9. Stevenson, P., 1996. Development of reliable sub-systems for Autosub. *Proceedings of Oceans '96, Fort Lauderdale*. The Marine Technology Society, Washington D.C., USA. Vol. 2, pp. 711-716.
10. Kimber, N.L. and Scrimshaw, K.H., 1994. Hydrodynamic testing of a $3/4$ scale Autosub model. *Proceedings Oceanology International '94*. Spearhead Exhibitions, New Malden, UK. Vol. 4, unpaginated.

11. O'Connor, P.D.T., 1995. *Practical Reliability Engineering*. Chichester, Wiley.

Table 1 Faults per mission and per km for the six groups of missions 1996-2000

Group	Missions	No. of missions	No. of faults	Faults per mission	km travelled	Faults per km	Average km per mission
A	1-55	55	2	0.036	34	0.059	0.62
B	56-88	33	13	0.39	199	0.065	6.03
C	89-122	34	17	0.50	242	0.070	7.11
D	123-179	57	13	0.23	616	0.021	10.81
E	180-216	37	5	0.13	1034	0.005	27.95
F	217-240	24	14	0.58	869	0.016	36.2
TOTAL		240	64	0.27	2994	.021	12.5

Table 2 Impact of faults on the operational performance of the vehicle over the five mission groups A to F.

	Impact Rating (H,M,L)	Group A	Group B	Group C	Group D	Group E	Group F	Total
Vehicle loss	H	0	0	0	0	0	0	0
Actual damage to vehicle	H	0	2	3	1	0	3	9
Emergency Abort triggered	H	0	0	1	0	1	1	3
Potential damage to vehicle	H	0	0	1	1	0	1	3
Potential Emergency Abort	H	0	3	2	0	1	0	6
Failure to meet prime objectives	H	0	1	2	1	0	1	5
Partial failure to meet prime objectives	M	0	2	4	0	1	1	8
Data lost or corrupted	M	0	0	0	1	0	0	1
Incurs additional cost	M	1	2	5	0	0	1	9
Delay > half a day	M	0	3	4	1	1	1	10
Delay < half a day	L	1	4	2	2	1	3	13
No significant impact	L	1	0	4	2	0	6	13

Note: The total number of entries in this table is greater than the number of faults listed in Table 1 because of multiple entries, for example an entry under 'delay' would also have an entry in one of the other categories .

Table 3 The estimated parameters for the Weibull distribution for each group of missions and for the 'All missions' group. No estimates were possible for Group A because of a lack of data (only two faults occurred).

Group	α	β
A	-	-
B	47.9	1.821
C	108.8	0.576
D	1032.7	0.592
E	275.8	1.133
F	317.5	0.890
All missions	403.9	0.678

Table 4 Overall lengths and under ice distances of missions in the Weddell Sea, 2001 with the probability of survival for each mission and for the whole campaign.

Mission	Mission length (km)	Under ice (km)	R(x)
247	50	40	0.935
248	53	38	0.938
249	53	42	0.932
250	5	0.1	0.999
252	60	45	0.927
253	63	45	0.927
255	10	0.3	0.999
257	9	2	0.996
259	15	2	0.996
262	80	61	0.903
TOTAL	398	275.4	0.636




Figure 1 Pareto diagram of the failure modes and their arisings during missions 1-240.




Figure 2 Time history of the faults per km and the faults per mission for the six groups of missions. The arrowheads show the progression in time from Group A to Group F.




Figure 3 Probability plot for the 'All missions' data set with that estimated from a Weibull distribution with $\alpha = 403.9$ and $\beta = 0.678$. The right hand scale shows the probability of survival.




Figure 4 Hazard rate as a function of distance travelled from the Weibull model for the 'All missions' group, all faults (solid line), for the seven high impact underway faults that would have led to loss in a covered ocean (dashed line) and for the five non-terrain following underway faults (dotted line).