Citation: Brito, M.P. and Griffiths, G., 2016. Autonomy Risk Assessment. Curtin T. ed. Part B: Autonomous Ocean Vehicles, Subsystems and Control. Springer Handbook of Ocean Engineering. Dhanak M.R. and Xiros N.I. (eds). Springer: New York, pp 427-544

Pre-publication version

Chapter 14

Autonomy: Risk Assessment

Summary

Oceanography and ocean observation in general is ever trending toward both automated *in situ* observation and working in extreme environments. These goals can only be met by de-risking the technology and deployment practices to acceptable levels of risks. A number of industries have standardised risk management processes to support the design and development of their systems. The lack of formal risk assessment of autonomous ocean vehicles has hindered the potential for true autonomy, which is required for exploring unstructured and unexplored environments. When discussing risks different stakeholders may have different consequences foremost in mind. For example the vehicle owner may be interested in risk of loss, whereas the user is interested in risk of vehicle unavailability. Other risks, such as legal risks and risk of collision, affect all stakeholders. This chapter presents a risk management process using several methods tailored to autonomous ocean vehicles in which risk assessment is a key component.

14.1 Introduction

Our human curiosity to understand unexplored and hostile environments has led us to develop state of the art automated technology capable of meeting our measurement needs [14.1][14.2][14.3]. Despite having identified the requirements for greater autonomous ocean vehicles (AOV) intelligence and autonomy, to date, physics prohibits the deployment of the truly intelligent autonomous underwater vehicle (AUV) [14.4]. Over the years, for practical implementations, the community has addressed risk by making the vehicles as simple as possible, the missions as modest as possible, and the level of supervision as high as possible [14.4][14.5][14.6].

Nevertheless, over the years there have been a number of vehicle losses. Amongst the most high profile losses are the loss of Autosub2 under the Fimbulisen ice-shelf, Antarctica on 16 February 2005, during mission 383 [14.7] and more recently the loss the Autonomous Benthic Explorer (ABE), during dive 222, off the coast of Chile, on 5 March 2010. Both losses are thought to have been caused by technical failure. In the case of Autosub2 a formal independent inquiry concluded that an Abort Command or a Loss of Power were equally likely to have caused the vehicle loss. There was no formal independent inquiry into the loss of ABE; following the accident, the design and operation team concluded that ABE suffered a catastrophic implosion of a glass sphere used for providing buoyancy causing instant destruction of on-board systems [14.8]. Anecdotal evidence exists of losses of smaller vehicles such as of Remus 100s and undersea gliders but these have never been formally reported.

Prior to the work presented in this chapter the risk of AOV loss may have been estimated by the principal engineer during design and deployment. From anecdotal evidence this was often pursued on an informal basis. Such an unstructured approach is rarely, if ever, found entirely satisfying to all parties. Neither it is likely to be immune to criticism, from one side or another. Estimating the likelihood of loss depends on a number of factors such as the vehicle's intrinsic reliability, the effects of the operational environment, the quality of the maintenance programme and the experience and competence of the deployment team. Assessing this risk requires a formal process that, in addition

to providing a quantitative assessment, is transparent and able to be followed and replicated by others.

In 2007, Griffiths and Trembanis [14.9] introduced a risk management process tailored to the operation of autonomous underwater vehicles. The approach was initially used for managing AUV risk of loss [14.10]. Different stakeholders have different interests in risk, for example the AUV owner is interested in the safe recovery of the vehicle, whereas a scientist is mostly interested in the recovery of data or in the vehicles' availability at a given time. The process for AUV risk management can be applied for managing other risks, not only risk of loss. In this chapter we present a risk management process tailored to autonomous ocean vehicles. Different methods for assessing different risks are discussed in detail.

Most examples are given for propeller based AUVs but these are equally applied to vehicles using buoyancy change engines, typically denoted as gliders. We also look at aspects of AUV design and deployment affecting risk.

14.2 Risk Management Process for autonomous ocean vehicles.

Whilst aspects of reliability had featured within papers on the use of autonomous ocean vehicles, reliability had been the topic of very few specific studies until the mid 2000s. The provocative, anecdotal evidence in Stokey et al. [14.11] gave rise to discussions within the community of users on how to improve reliability and reduce risk at a time when vehicle operations were just starting to be independent of deploying vessels. Spurred by these considerations, and especially by future requirements to operate under ice, Griffiths et al. showed that simple statistical methods could be applied to estimate risk of loss of autonomous vehicles operating in various environments [14.12]. Subsequently, Podder and colleagues extended this approach to look at reliability growth [14.13], which is the main desired outcome. However, these studies looked only at revealed reliability, documenting what problems were emerging; they were not set within a framework for risk management.

Figure 14.1 A flowchart representing the risk management process proposed by Griffiths and Trembanis [14.9].

The loss of the Autosub2 vehicle under the Fimbulisen in 2005, and the recommendations of a subsequent Board of Inquiry led to the development of a holistic risk management process for AUVs,

Figure 14.1 [14.9]. By simple extension of differently described operating environments it is applicable to the wider class of autonomous ocean vehicles. This remains the only published structured procedure for risk management of these vehicles.

Taking the steps in sequence, first a responsible owner is identified and they state an acceptable probability of loss for the campaign under consideration. The owner may factor in the importance of the vehicle's mission, the value of the data it would return, the future programme for the vehicle, and other considerations, into the acceptable probability of loss. Independently, the principal investigator, or user, sets out the requirements of the campaign in terms of number of missions, their duration, and the environment characteristics, such as surface or sub-surface, coastal, open ocean, under ice, or in areas of high traffic. The next step is for the technical team to assess the probability of loss in light of the campaign just described. This is a difficult step, and much of the remainder of this chapter is given over to methods used to make this estimate.

Accepting for now that this estimate can be made, it is then compared with the responsible owner's acceptable risk. If the estimated risk of loss is less than the owner is willing to accept, there is a need to demonstrate that the estimated risk is realistic before the campaign can proceed. This demonstration could take the form of reliability trials, where the vehicle would be run in a benign environment over a duration or range commensurate with the planned missions and the outcomes evaluated. If the estimated risk of loss is greater than the owner's acceptable risk several feedback paths are followed.

The first activity is to rank the risk factors identified by the analysis of fault history, or fault projections. Those that can be mitigated, without question, are dealt with. This demands that the causes of failure are fully understood and the corrective measures are known, able to be implemented, and tested, before the campaign. Belief in ability to mitigate faults, or assertion alone, are not sufficient.

In parallel, the user may reassess the missions required and the responsible owner may reassess the acceptable risk, following which, the test for acceptable risk is repeated. Two or more iterations may be needed, or indeed, the decision may be to postpone or cancel the proposed campaign.

Primarily, this process has been used for campaigns with a significant risk of vehicle loss in the Polar Regions. Brito et al. described its use for a campaign of six Autosub3 missions under Pine Island Glacier, Antarctica [14.10]. Here, the initial estimated risk was sufficiently high that reliability-proving trials in a Norwegian Fiord were an essential precursor. While there were incidents on the Pine Island Glacier missions, the vehicle survived the campaign. More formal estimates of the likelihood of successful fault mitigation were made when the process was applied to an International Submarine Engineering (ISE) Explorer AUV for use in under sea ice missions in the Arctic [14.14]. The process has also been used in high-risk open water areas, such as for Autosub3 missions in the high traffic Bosphorus Strait.

In practice, the methods used for quantifying the actual risk are directly linked with the acceptable level of risk. Risk is a measure of uncertainty. If the acceptable risk is high, for example, if the acceptable risk of MAS loss is 90%, there is little reward in applying formal methods for quantifying the actual risk. This is the case for platforms that are considered expendables, which may be deployed to study very rare events, such as a hurricane or the eruption of an underwater volcano. If the acceptable risk of MAS loss is low, then it becomes important to quantify the actual risk. This will give the user more confidence that the acceptable risk is met. This is the case where the platform is financed for long-term use.

14.3 Risk of failure

The discipline of estimating a system's probability of failure, or a system's reliability is well understood in many engineering fields. Reliability is defined as the probability that failure will not occur in the period of interest [14.15]. This probability figure can be estimated from system or component usage. However, in some applications it is not possible to obtain system or component failure history, particularly if the operating environment cannot be recreated in the laboratory or if a new system or component is being used. In such situations reliability estimates can be obtained from expert subjective judgment – this is discussed in detail in section 14.6.1.

Probabilistic modelling is key for effective reliability estimation of a system composed of several components or sub-systems. To facilitate the development of such models graphical methods such as fault tree analysis can be applied. In this section we present a summary of the techniques used for estimating AUV reliability. In section 14.3.1 we present methods for estimating the reliability and reliability growth from operational data. In section 14.3.2 we present how fault trees can be used for estimating the reliability of a complex system comprised of many other components.

14.3.1 Reliability Estimation

The number of failures that emerge during a vehicle, system, or component test can be used to estimate the mean time to failure for the vehicle. However, the mean time to failure alone is not sufficient to support decision-making with regards to AUV deployment. By fitting parametric models of reliability to AUV fault data it is possible to produce formulae that can be used to predict the probability of success of any mission in terms of its length. A number of distributions have been used for modelling AUV failure distribution with time. The first study on AUV reliability estimation was conducted by Griffiths et al. [14.16]. Failures of Autosub1 autonomous underwater vehicles were collected from its first mission within Empress Dock, in Southampton in June 1996, to a science campaign on board the RV *Calanus* at Oban in November 1999. During this time, the vehicle conducted 216 missions covering 2125 km. Fifty failures were recorded, giving a probability of failure per mission of 0.231. Pareto analysis showed that human error, acoustic telemetry failure and failure to dive were the top three failure modes in Autosub1's failure history. This report used statistical models for predicting the failure probability as function of mission length. The dataset used for this analysis consisted of the distance travelled for each mission and whether the mission ended due to a fault. Each entry was classed as censored or not censored. A censored observation is one where the mission was terminated not due to failure, whereas a non-censored data captures the distance at which a failure took place. Two types of functions were defined directly from the data [14.12]. The first function, the probability of failure $F(x)$. Where x stands for an instantiation of distance X . In mathematical terms this is captured by $F(x) = P(X < x)$. Where $P(X < x)$ is obtained directly from the probability function that is fitted to the historic data and lower case x is an instantiation of X . The second function of interest is the reliability, otherwise known as survival function, is the probability of survival, $R(x)$. The reliability is the complement of the probability of failure, it is mathematically defined as $R(x) = 1 - F(x) = 1 - P(X < x)$. It stands for the probability of the system surviving failure, the distance x .

Griffiths et al. fitted three well known distributions to the historic data: *Weibull*, *log logistic* and *log normal*. Results showed that the probability of a fault occurring in a 100km mission was 0.352 using the *Weibull*, 0.334 for the *log logistic* and 0.309 for the *log normal*. Software reliability was studied separately using two independent software reliability models: the Poisson model and the Littlewood model. Both models lead to the same reliability estimate for a 100 km mission of 0.93 for the software.

The dataset was later updated to include Autosub2 science missions on MV *Terschelling* 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. A total of 869km were covered in 24 missions [14.12].

Here the authors fitted six different distributions to the data: extreme value, normal, logistic and their logarithmic versions: Weibull, log normal and log logistic. The logarithmic distributions provided a better fit to the data 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. Results were similar to those previously reported.

The first reported work on the effects of upgrades on the reliability growth for AUVs was presented in 2003 [14.17]. The dataset consisted of Autosub missions considered in [14.12] plus missions that were conducted in two engineering trials in June 2002 and September/October 2002. The purpose of these trials was to test new software that gave enhanced autonomy. The results of the analysis showed reliability growth prior to the missions carried out in June 2002. However, after the upgrades there was an increasing failure rate. These were worrying results because the vehicle was due for an under ice campaign in February-March 2003. The team used the reliability model to estimate the number of missions, with an acceptable number of faults, needed to ensure reliability growth. The analysis concluded that 10 missions were required with distances varying from 5km to 144km in length, with only two high impact faultsⁱ allowed on the shorter missions.

A study in reliability growth was carried out for the Dorado AUV designed and operated by Monterey Bay Aquarium Research Institute (MBARI) in preparation for an Arctic campaign in the following summer [14.13]. The dataset consisted of mission data collected during operations in 2003 and 2004, where the vehicle had travelled approximately 1700km in Monterey Bay, California. Several failures emerged, the analysis considered only failures that resulted in an abort - a total of 14. Basic trend analysis using the Laplace test yielded the conclusion that the AUV reliability exhibited an increasing trend. This was verified using two independent reliability growth models due to Duane's and Crow's. Both use graphical approaches to support reliability growth analysis. The estimation of the mean time between failures (MTBF) was conducted for different groups of failures; set according to their criticality. The estimated MTBF for the extremely critical faults was 167.8hr.

Reliability estimation based on revealed operational history helps those involved understand better the failure pattern of these autonomous vehicles, enabling the implementation of failure mitigation mechanisms during the operational lifetime of the AUV. However, for early lifetime, this failure history is not available. For the early lifetime of a vehicle the reliability of a vehicle can be estimated using reliability modelling techniques. These are discussed in the next sections.

14.3.2 Reliability Modelling

Reliability modelling aims to estimate systems reliability based on its intended functionality, and design architecture. The failure model end effect analysis (FMEA) is a well established technique for identifying critical components or phases of the system design and operation [14.18]. However this technique does not accommodate quantitative estimation of the system's reliability and therefore it will not be discussed in this chapter. Event trees and fault trees are two basic methods for system reliability quantification. Event trees use forward logic. They begin with an initiating event (an abnormal incident) and 'propagate' this event through the system under study by considering all possible ways in which it can affect the behaviour of the subsystem. Event trees are useful for accident sequence analysis in which the aim is to estimate the likelihood of a sequence of potential functioning or malfunctioning events. Thus event trees cannot be used for estimating the probability of failure of a system. This figure can be estimated with the support of fault trees. In a fault tree analysis one attempts to develop a deterministic description of the occurrence of an event, the top event, in terms of the probability of occurrence of other (intermediate) events. Intermediate events are also described further until, at the finest level of detail, the basic events are reached. When the

ⁱ Faults that can lead to vehicle loss.

top event is failure of a system then the basic events are usually failures of components. A fault tree can be considered as an expression in a Boolean model, which implements predicate logic. A fault tree contains the same operators, for example conjunctions and disjunctions. However, instead of operating with bits, where 0 corresponds to false and 1 corresponds to true, fault trees operate with probabilities - varying from 0 to 1 [14.19]. Figure 14.2 presents a fault tree devised by Griffiths and Brito for estimating the probability of failure during the deployment of an AUV through an ice hole [14.20]. Graphically a conjunction is represented using 'and' and a disjunction is represented using an 'or' gate. In Figure 14.25, the operators named G_1 , G_2 , G_3 and G_4 are 'or' gates. The G_5 operator is an 'and' gate. A fault tree injects transparency into the fault analysis process. Failure modes are clearly presented and how they propagate in the system is visually displayed.

Figure 14.2 Probability of failure during AUV deployment through an ice hole. Fault tree presented in [14.20].

There is a unique probabilistic model associated with each fault tree, which enable us to calculate the probability of failure for the top event. The probability of failure to launch, P_{10} , is calculated using the following expression.

$$\begin{aligned}
 P_{10} &= G_1 \\
 &= 1 - (1 - G_2) \times (1 - P_3) \times (1 - P_4) \times (1 - G_3) \\
 &= 1 - ((1 - P_1) \times (1 - P_2)) \times (1 - P_3) \times (1 - P_4) \times (1 - (1 - (1 - G_4) \times (1 - G_5)))
 \end{aligned} \tag{1}$$

$$\text{And } G_4 = 1 - (1 - P_5) \times (1 - P_6) \times (1 - P_7) \tag{2}$$

$$G_5 = 1 - (1 - P_8) \times (1 - P_9) \tag{3}$$

The probability of failure for each basic event can be estimated using statistical modelling discussed in the previous sub-section in which a probability function is fitted to observed data. Alternatively, if operational data is non-existent, the probability of failure for the base events can be estimated using expert subjective judgment. In section 14.6 we provide more details on the formal methods that can be used to elicit expert judgment.

14.4 Risk of Collision

The risks of collision for autonomous ocean vehicles are very real. However, as unregulated vehicles, no figures for the number of losses due to collisions can be given; there is no requirement to record losses for these unmanned vehicles. Through anecdote, the incidence of collision for underwater or surface vehicles is sufficiently high that any group operating several vehicles is likely to encounter this risk regularly. Sensitivity of many operators to the details of collisions also means that quantitative, or attributable, information is also rare. The examples given here are based on our knowledge of real incidents.

14.4.1 Risks of collision on or near the surface

For surface vehicles, and when underwater vehicles operate on or near the surface, the main collision risks are:

- The shore, or water shallower than the minimum operating depth of the vehicle. At least one autonomous vehicle has run ashore because parts of the chart from which waypoints were taken were surveyed in the 19th century, and the position error for the shoreline was substantial. Operating near a coast with very gentle beach gradients without full knowledge of, and accounting for, the tides can give rise to grounding.
- Manned surface craft. Collision with manned surface craft has probably been responsible for many autonomous vehicle losses. In some cases, the manned craft concerned has been the craft deploying the autonomous vehicle. Mistakes and miscommunication during launch and recovery, especially, can lead to collision with the support vessel. Collision with the sides of the vessel may inflict little damage, but collision with propellers has on more than one occasion led to damage to the pressure vessels of underwater vehicles and immediate loss.

Where statistics of vessel traffic are available probabilistic models may be used to assist in estimating the risk of collision. Merckelbach has developed such a model, specifically for undersea gliders [14.4.1]. In principle his approach is applicable to other autonomous vehicles. The model considers the vessel and vehicle speeds, the vessel traffic density, the statistics of vessel length, draft and breadth, the water depth, the operating parameters of the vehicle, e.g. time on surface and within the draft of vessels, and a bow-wave factor. This last factor accounts for the pushing aside that can happen as collision is imminent; in effect the effective breadth of the ship is reduced by this factor.

Mission planning to take account of surface vessels can make use of the Automatic Identification System required on “internationally voyaging ships” of 300 tonnes or more, all passenger ships, and, from 2014, the entire European Union fishing fleet of vessels over 15 m in length. Websites with real-time mapsⁱⁱ enable planning ahead, and also near-real time collision avoidance.

For very intensive shipping areas, e.g. Bosphorus Strait with typically six large ships per hour, autonomous vehicles should avoid the surface.

- Large flotsam and jetsam, such as logs and baulks of timber, lost freight containers. These hazards are less amenable to quantifiable modelling than ship traffic. Local knowledge is likely to be important, e.g. on the known areas where felled trees from boreal forests may be prevalent.
- Floating nets and fisheries-related hazards. Despite being banned by international agreement since 1992 in international waters, drift nets with surface floats that are meant to catch pelagic fish are still in use within the Exclusive Economic Zones of many countries e.g. the

ⁱⁱ Such as at <http://www.shipais.com/index.php?map=uk> accessed 19 March 2014.

USA. Other nations stipulate maximum lengths (e.g. 2.5km for the EU). These nets are a hazard for autonomous ocean vehicles on the surface. Local knowledge may suggest where such fishing happens. Avoidance using telepresence from the vehicle to the command centre of video or radar imagery may assist in spotting the surface floats.

- Offshore structures, e.g. oil and gas installations, wind farms. The positions of these surface hazards are usually well known, and fixed. The onus is on the autonomous vehicle operator to plan missions with sufficient contingency and allowance for navigational error such that a vehicle does not infringe the exclusion zones around these structures.
- Coastal structures such as breakwaters, or moles. While there may be no exclusion zones around these hazards, their positions are fixed, and appropriate prior planning should be adequate to mitigate the risks.

14.4.2 Risks of collision underwater

Apart from the universal case of collision with the seabed, the other risks of collision when underwater depend on the environment of the missions. Near-shore missions may encounter obstacles such as piers and other fixed structures, vessels at anchor, and the submerged parts of vessels underway. In the open ocean, apart from with the seabed, collision when submerged is likely to be a rare event. Under ice, collision is likely with ice at the sea surface, with projecting keels of sea ice, and with icebergs. While grounded icebergs, ice shelves and glaciers and sea ice that is locked to the shore or a shoal may be stationary, other forms of ice should be considered as moving obstacles.

Collision, or obstacle, avoidance for autonomous underwater vehicles is a well-studied topic, combining as it does the challenges of sensing, interpretation and action. The effectiveness of obstacle avoidance is one indicator of the degree of autonomy of a vehicle. Horner et al. [14.22] summarises these challenges, and provides examples of how a forward-look sonar array can be integrated with an autopilot control algorithm that minimises the cross-track error from the desired path. That is, the system provides for adequate, but minimum, deviation from the desired path to avoid the obstacle. The interpretation of the obstacle characteristics in this implementation uses image information from the sonar array, borrows techniques from computer vision, and can be made robust against false targets. This approach is relevant for avoiding the seabed, but is also sufficiently powerful to provide obstacle avoidance in more complex situations, such as near piers and other man-made structures.

There are situations where a simpler approach may give sufficient capability, for example, where the seabed is the only obstacle likely to be encountered. Using a single beam mechanically scanned sonar and a horizon-tracking algorithm McPhail et al. [14.23] showed that effective seabed avoidance could be delivered when terrain following at altitudes down to 10m on the flanks of a seamount, and down to 3m on the flatter, but rocky, summit. Vehicle reactions included simple avoidance in the vertical plane, and turn-around and retry.

Under ice, information is needed on obstacles ahead, below and above the vehicle. One approach, devised for the DEPTHX vehicle [14.24] proposed a suite of 24 narrow beam sonars and 30 imaging sonars. In the subsequent Endurance vehicle, 64 pencil beam sonars provided a 3.5 steradian view of the environment within ice-covered Lake Bonney in Antarctica. A far simpler approach is to use for obstacle avoidance purposes, sonar information from other instruments on the vehicle. Pebody has described how, for an under ice application, a forward-looking sonar may be augmented using range to reflector information from the four beams of an upward-looking and a downward-looking acoustic Doppler velocity log [14.25]. This system was tested using an air curtain behind a ship, to simulate an iceberg, and used under fast ice off Greenland, and under Fimbulisen and Pine Island Glacier in Antarctica to good effect.

As a back-up to collision avoidance, some operators include rubber fenders or other protection on the nose and undersides of their vehicles. Incidents where collision avoidance has failed include repercussions of hardware failures and failures of the control systems to properly appreciate the form of the environment and hence failure to take proper action. Drawing on examples with Autosub AUVs, a sternplane fault led to the vehicle diving far too steeply, such that its bottom avoidance sonar was at the incorrect angle, the sonar could not see the seabed, the vehicle nose collided with the seabed but the vehicle continued to make slow forward progress until the nose became full of sediment. For an example of inappropriate response to an unexpected environment, when terrain following at low altitude up a steep cliff off Sicily, Autosub2 encountered an overhang. The appropriate response would have been to reverse out, but this behaviour was not pre-programmed, and instead the vehicle dropped its abort weight and had to be retrieved using an ROV. These lessons also serve to show that not all collisions are fatal to a vehicle.

14.5 Risk of unavailability

System availability is defined as the probability of the system being available given that the system is needed at a given time. System unavailability is the complement of this figure. The deployment of an AUV consists of several phases. A series of tasks are carried out in each phase of the deployment. These human or machine related tasks are not immune to error. Estimating the availability for the AUV being at a given phase is only possible if we take into account the sequence of phases that precede that phase. This problem can be mathematically modelled using a probability approach denoted as Markov theory.

For a brief description of Markov theory consider a probability problem with a set of outcomes of interest $E_1, E_2, \dots, E_k$. Given that there is a probability, p_j , associated with each event the joint probability for a given sequence is defined by the multiplicative property, thus $P(E_1, E_2, \dots, E_k) = p_1 \cdot p_2 \cdot \dots \cdot p_k$. The Markov chain theory introduces an assumption that simplifies this expression; it considers that the outcome of any trial depends on the outcome of the preceding trial and only on it [14.26]. Therefore, if event E_1 precedes event E_2 and event E_2 precedes event E_3 and so on for the remaining events, then instead of associating a probability to an event E_j , it uses a transition probability p_{jk} for every pair of events (E_j, E_k) . Where p_{jk} is the probability of E_k occurring given that E_{k-1} occurred in the previous trial. A Markov model may have more than one sequence of events; E_j may have two or more posterior states. Therefore j and k are not necessarily adjacent. In addition to p_{jk} , one must also define the probability of E_j occurring at the initial trial, a_{j0} . Therefore for the initial trial the $P(E_j, E_k) = a_{j0} \times p_{jk}$. For the general case, considering a sequence of several transitions, given that event E_{j0} precedes E_{j1} which precedes E_{j2} and so on for the remaining events, the joint probability distribution is computed using the expression in equation 4:

$$P_r(E_{j0}, E_{j1}, \dots, E_{jn}) = a_{j0} p_{j_0 j_1} p_{j_1 j_2} \dots p_{j_{n-2} j_{n-1}} p_{j_{n-1} j_n} \quad (4)$$

It is not unusual to find problems where one state has more than one potential preceding state. When this is the case the mathematical calculation cannot be performed using the simple equation presented above. Instead the calculation is performed using matrix operations. The transition probabilities are arranged in a matrix denoted as transition matrix or stochastic matrix. The transition probability together with the initial state vector completely defines the Markov chain. The availability after n transitions can be calculated by the product between the transition matrix to the power of n and the initial state vector.

Brito and Griffiths [14.27] used Markov chains for modelling and estimating the availability of Autosub3. The Autosub3 deployment sequence was modelled as a Markov chain in which each state corresponds to one phase of the vehicle deployment. A key assumption in the proposed model is that the Autosub is deployed from a stationary vessel. The model can be altered to capture the case where

the AUV is deployed from an ice hole or from the coast. This would consist of adding a state for which the transition probabilities would depend on the failure modes that can emerge in the different type of deployment. Griffiths and Brito in [14.20] give an example of when the approach was adopted for deployments from ice holes.

Here we briefly describe the model proposed in [14.27] for estimating AUV availability based on surface vessel deployments. The Markov chain model for AUV availability consists of eleven phases, these states and the transitions are presented in Figure 14.3. The first state captures the phase when the vehicle is on board of the vessel, switched on, ready for testing (Dp). A series of communication, actuation and navigation checks are carried out during this phase. Having passed the on-board checks the deployment moves to vehicle ready for deployment phase (Dr). Here the vehicle is waiting to be deployed overboard. While overboard there is the risk that the vehicle may run under the vessel causing severe damage, in which case the vehicle may need to be salvaged. This scenario is captured by the transition from phase overboard (O) to salvage (S). Once overboard and at a safe distance, communications checks are carried out before the vehicle starts diving. The vehicle is then set to stop diving when it reaches a predefined depth and holds pattern (Sh). During this period, the vehicle is still within acoustic range, more checks are carried out before committing to the science mission. Once underway (U), there are two possible outcomes: the vehicle can be recovered (R), or it can be lost (L). If historic operational data exists then the probability of AUV loss can be computed using reliability modelling techniques combined with expert judgment - this method is described in the section 14.6. Alternatively if there is no historic data of AUV operation than the probability of loss can be computed using fault tree modelling that takes into account expert judgment with regards to the consequence of each failure mode.

Figure 14.3 Markov state space model capturing the sequence of events undertaken during AUV deployment and operation. A directional arrow from state i to state j means that the process can move from state i to state j .

The AUV can be recovered while it is in the overboard phase, during diving, in holding pattern or underway. The vehicle can also be lost from any of these phases. If a vehicle is lost, the deployment can be maintained in a permanent loss state (L), captured by the transition from $p_{7,7}$ or it can move

to a state of Salvage (S) or Found(F). From these states the AUV can be declared fit for re-deployment, this is captured with the transition from state (S) to (Dp) and (F) to (Dp), or declared scrapped (Sc). A deployment can reach a state of AUV scrapped from overboard (O), dive (Dv), holding pattern (Sh), underway (U) and recovery state (R).

For Autosub3 availability analysis the transition probabilities were calculated based on expert judgments. Three experts with more than sixty years combined experience on AUVs have estimated the probability for each transition. The transition probabilities are a result of the un-weighted linear pool aggregation of the expert judgments. The transition probabilities for the availability of Autosub3 are in Table 14.5.

Table 14.5. Aggregated expert judgments for the transition probabilities.

Transition stimuli	Transition Probability	Transition Stimuli	Transition Probability	Transition Stimuli	Transition Probability
$P_{1,1}$	$1-p_{1,2}$	$P_{4,5}$	0.9565	$P_{7,9}$	0.33
$P_{1,2}$	0.875	$P_{4,7}$	0.0085	$P_{7,10}$	0.33
$P_{2,1}$	0.055	$P_{4,8}$	0.035	$P_{8,1}$	0.998
$P_{2,3}$	0.94	$P_{5,6}$	0.98	$P_{8,7}$	0.002
$P_{2,10}$	0.005	$P_{5,7}$	0.0055	$P_{9,1}$	0.75
$P_{3,3}$	0.0195	$P_{5,8}$	0.0145	$P_{9,11}$	0.25
$P_{3,4}$	0.925	$P_{6,7}$	$1-p_{6,8}$	$P_{10,1}$	0.7
$P_{3,7}$	0.0495	$P_{6,8}$	$p_{survival}$	$P_{10,11}$	0.3
$P_{3,8}$	0.006	$P_{7,7}$	$1-(p_{7,9}+p_{7,10})$	$P_{11,11}$	1

The Markov property states that the sum of the probabilities leaving any given state must equal to unity. If the sum of all transitions leaving a state is constant c , where c is lower than one, then the probability of the process remaining in the same state in the next transition is $1-c$. The Markov condition was applied to calculate transition probabilities p_{11} , p_{67} and p_{77} .

The availability of Autosub3 from ready for test, whilst on deck, to underway was calculated to be 0.75. This estimate is similar to that obtained for unmanned aerial vehicles (UAVs). The United States Office of the Secretary of Defense published the availability figures for five UAVs: Predator RQ-1A (concept demonstrator) (0.40), Predator RQ-1B (early production), (0.93), Pioneer RQ-2A (1990-1991) (0.74), Pioneer RQ-2B (0.78), Hunter RQ-5 (reliability enhanced 1996-2001), (0.98), Average UAV (0.77)ⁱⁱⁱ.

14.6 Risk of loss

The task of estimating AUV risk of loss is affected by many factors. Technical failure rate or the likelihood of a human error are two important factors in the assessment of AUV risk of loss but these factors alone are not sufficient. A failure has different impact depending on the AUV operational environment. For example, a failure leading to unexpected drop of the abort weight causing the vehicle to surface may have a low impact in open water operations, however if the same failure takes place whilst the vehicle is under an ice shelf, sea ice or busy coastal area it may lead to vehicle loss. The human element is also of great importance when it comes to failure mitigation, as highlighted by

ⁱⁱⁱ This study was available online at www.acq.osd.mil/uas/docs/reliabilitystudy.pdf but was subsequently withdrawn.

Stokey [14.11], many technical failures in AUV operations can be mitigated by a more experienced team.

These factors are difficult if not impossible to capture deterministically, particularly if the number of previous deployments in the target operational environment is small or if the environment and operational conditions are impossible to predict. Autonomous ocean systems are not unique in this challenge, and many techniques have been developed and widely accepted, for risk quantification in areas such as nuclear, process, oil and gas and aerospace safety. Common to these approaches is the formal use of expert judgment.

In this section we present formal expert judgment elicitation methods focusing on their pragmatic characteristics for estimating AOV risk. When studying component or systems reliability it is not sufficient to have the mean time between failures, designers and operators are often interested in knowing the component or system reliability as a function of time in operation or mission length. This is imperative for the implementation of corrective measures for risk reduction. Similarly, when operating AOVs the decision to go, or not go, for a mission is better informed by a risk profile that captures the probability of loss with mission time or length. The creation of a probability of loss profile is possible by integrating experts' subjective judgments with statistical survival techniques.

14.6.1 Expert judgment elicitation processes

Governments and companies are increasingly relying on expert panels for providing risk assessments for high risk or high profile projects such as the safety assessment of a nuclear power station or the implementation of a major science programme [14.28]. There are many dangers in eliciting expert judgments in an informal way. When providing assessments to events, people often follow one of a number of mental shortcuts, denoted as heuristics [14.29, 14.30]. Research has shown that when used incorrectly, these heuristics can lead to systematic and predictable bias [14.31]. In their work, Tversky and Kahnemann [14.29] have identified three predominant heuristics: representativeness, availability and adjustment and anchoring. Representativeness is described as the tendency to judge the probability that A belongs to B by how representative A is to B. For example if the description of Peter is highly representative of an engineer then people tend to judge the probability of Peter being an engineer as high regardless of the base rate, that is, the proportion of engineers in society. The base rate neglect is one type of bias that can be introduced by following the representativeness heuristics; other types of biases are for example insensitivity to sample size, insensitivity to predictability and misconceptions of chance and regression [14.29]. Availability heuristics is the tendency to judge a particular event by the ease with which the instances and events can be brought to mind. For example one may estimate the risk of heart attack amongst middle aged men by remembering specific personal examples. Here bias can be introduced due to exposure to negative outcomes; they tend to be easier to remember than positive ones. Anchoring and adjustment is a tendency to anchor probability estimates at an initial estimate. Insufficient adjustment results in biases of underestimation or overestimation.

Formal judgment elicitation has been proposed as a solution to reduce biases and facilitate reproducibility of the results [14.32]. A formal expert judgment elicitation is a structured process that consists of a number of phases; a number have been proposed in the literature [14.33]. Below we describe the generic structure of a formal judgment elicitation process tailored to autonomous ocean vehicles:

1. Set out the issues. In this phase the issue that is to be addressed is described in general terms. This is a scope definition stage; it defines the nature and direction of the analysis, the choice of questions to be asked and issues to be considered. For example, the issue of interest is to estimate the probability of vehicle loss.

2. Selecting the experts. An expert is someone with specialist knowledge in the task in hand who is also capable of decomposing a complex problem in subsidiary problems. In vehicle risk assessment an expert should be someone with experience in the type of fault scenarios and in autonomous ocean vehicles operations.
3. Clearly define the issues. The issue in hand is what is the probability of failure F leading to loss (L) in operational environment E ? This is mathematical represented as $P(L|F,E)$.
4. Training the experts and eliciting judgments. This is the most important task in the elicitation process. Different formal judgment elicitation processes follow different approaches. For some, training is quite informal it consists of providing some examples of typical assessments and an explanation of basic principles of probability theory. Other judgment elicitation processes are stricter and encourage the facilitator to elicit expert assessments for a number of seed questions. These are questions for which the facilitator knows the answer but the expert is not familiar with the problem or question [14.32], and where responses can be used to calibrate individual assessments.
5. Analysing and aggregating. This is a critical phase of the assessment. Here the decision maker uses a strategy for combining the expert judgments into a single assessment that represents the group's view. This can be done mathematically or behaviourally. A mathematical aggregation uses analytical functions for combining the expert judgments. The analysis of expert judgments must be conducted prior to the aggregation to ensure that any misunderstandings and bias are removed from the assessments.
6. Complete analysis and write up. The assessments provided by the experts are documented, disagreements between experts are recorded. A report is submitted to the experts for review. Following this review a report is submitted to the decision maker informing the results of the risk assessment exercise.

Mathematical Aggregation

The use of mathematical methods for consensus building has the benefit of facilitating the application of the elicitation process because experts are not asked to agree on the final judgment; in fact experts do not have to be in the same room during the elicitation. Supporters of mathematical methods have argued that behavioural methods, in addition to being time consuming, can lead to systematic biases caused by group polarization [14.34].

The analytical function for combining individual judgments can take into account the level of expertise in the topic and the effectiveness of the calibration. These factors are captured in a weight w_i , where $i = 1, \dots, n$ or in an *a priori* distribution judgment. Methods for mathematically aggregating expert judgments were developed based on a set of widely accepted axioms [14.35]. Perhaps the most appealing method for combining expert judgments is the *linear opinion pool* [14.36]. Here the expert judgments are aggregated by taking the weighted average of the assessments provided by the experts. The *linear opinion pool* complies with the property of *consistent marginalization* [14.37]. A different mathematical aggregation method, the *logarithmic opinion pool* uses a multiplicative averaging of the expert assessments [14.36]. The *logarithmic opinion pool* does not comply with the marginalization principle but unlike the *linear opinion pool* it complies with the *principle of external bayesianity*. This can be explained as follows, suppose that the decision maker or facilitator has reached an aggregated $p(\theta)$ but has subsequently learned new information relevant to θ . One way to update the judgment $p(\theta)$ in light of the new information is by updating the individual expert judgments $p_i(\theta)$ and then aggregate all judgments using an analytical function. A second way is to update the original $p(\theta)$ without updating the individual expert judgments, using Bayes theory. A

mathematical aggregation method that complies with the *principle of bayesianity* will produce the same result regardless of the process used for creating the final judgment.

Other mathematical aggregation methods are more complete from the axiomatic viewpoint. But this comes at a price of being more difficult to implement. *Bayesian aggregation methods* adopt a completely different prerogative to the *linear and logarithmic pools*. Here, the decision maker begins by defining his own prior distribution $f(\theta)$. The expert judgments are then incorporated, using Bayes' rule, to form the decision maker's *posterior* distribution $f(\theta|D)$. Where $D = \{f_1(\theta), \dots, f_n(\theta)\}$ is the set of experts' elicited distributions. In this context, according to Bayes theorem, the $f(\theta|D)$, is proportional to $f(\theta)$ multiplied by the likelihood term $f(D|\theta)$. The problem with Bayesian methods is that the decision maker must specify his own prior belief about the risk $f(\theta)$ and the knowledge and beliefs about the experts $f(D|\theta)$ [14.35].

In 2008, in preparation for the Autosub3 deployment under the Pine Island glacier a risk model for Autosub3 was developed based on the assessments provided by eight independent experts. Experts were asked to assess the probability of AUV loss given that a fault X emerges during the deployment [14.10]. In addition to providing an estimated probability of loss, experts provided a weight, from 1-5, capturing their confidence in the assessment. A weight of 1 represented little confidence in the assessment, whilst a weight of 5 meant that the expert was very confident. The experts considered 63 faults, in four environments: open water, coastal water, under sea ice and under ice shelf. A detailed description of each environment was provided prior to the assessment. The faults were collected during six Autosub3 campaigns from mid 2005 to mid 2008.

Figure 14.4a shows the relative frequency of the average of the weights used by the experts. The figure shows that on average, for 70% of the assessments, experts assigned a weight of 3 or higher, for open water, coastal waters and ice shelf environments. Experts were least confident with the sea ice environment.

Figure 14.4b shows the un-weighted linear pool probability of loss, for the five most critical failures in the four different environments. This makes it clear that, in these experts' opinion, faults that would have a low probability of leading to loss in open or coastal water would have a high probability of loss under an ice shelf, but not through simple scaling.

Figure 14.4 a) relative frequency of the average weights used by experts Pareto plot of the faults collected for the expert judgment elicitation. b) Un-weighted linear pool, probability of loss, for the five most critical failures.

Expert judgment has also been applied to estimating the risk of loss of commercial AUVs, including two Remus 100^{iv} vehicles in coastal areas and under sea ice. These vehicles were operated by the Center for Coastal Marine Sciences at California Polytechnic State University (CalPoly). They were used for a range of missions to better characterize and improve understanding of coastal waters, Moline et al. [14.38]. The operational data set contains faults and incidents recorded from 186 missions between July 2001 and February 2009.

For this elicitation experts were asked to specify an unimodal distribution for the probability of loss given a fault. A narrow distribution would show that the expert is confident about the assessment, a widely spread distribution would show that the expert was very uncertain about the assessment. Expert specified this distribution using five quantities:

- The lower bound, L. The minimum value that $P(\text{loss})$ can take.
- The upper bound, U. The maximum possible value of $P(\text{loss})$.
- The median, M. The value for which there is 50% chance of $P(\text{loss})$ being above or below it.
- The lower quartile, LQ. The value for which there is 25% chance that $P(\text{loss})$ is between L and LQ and 25% chance that it is between LQ and the median, M.
- The upper quartile, UQ. The value for which the expert is 25% confidence that $P(\text{loss})$ is between the median and UQ and 25% confidence that it is between UQ and the upper bound.

^{iv} Shallow waters AUV manufactured by Hydroid Inc.

The experts' judgments were aggregated using the un-weighted linear pool, but in two separate groups, the optimists and the pessimists; results are presented in [14.39]. A key conclusion was these vehicles were being knowingly operated with very low probability of loss vulnerabilities, which, on a statistical basis could become important given the large number of missions undertaken.

The limitation with mathematical aggregation approaches is that experts do not have the opportunity to discuss different views and thus to remove bias from their assessment. The aggregation of expert judgments in two different groups helped solve this potential bias issue. However by doing so, the decision maker needs to decide whether he or she is going to be an optimist or pessimist. One way to mitigate the risk of bias is by having experts in the same room, allowing them to discuss their views before agreeing on a judgment. Such an elicitation method is denoted behavioural elicitation method.

Behavioural Aggregation

The decision of whether to aggregate expert judgments mathematically or behaviourally ultimately rests with the decision maker. Research has shown that a group consensus will never outperform the assessment of the best expert in the group. The difficulty is in identifying the best expert [14.40]. Behavioural aggregation brings the experts together as a group, provides a structured process so they can share their knowledge and allow persuasive arguments to change their views, and help them to generate a consensus distribution [14.41]. This approach is often more transparent and more immune to criticism than mathematical aggregation methods.

A behavioural expert judgment elicitation was conducted to build a risk model for two ISE Explorers AUVs, operated by Defence Research Development Canada (DRDC) [14.20]. As these vehicles were to be used on data gathering of national importance in the high Arctic [14.42], a transparent and justifiable process was needed for how the risk of vehicle loss was evaluated, hence a behavioural aggregation approach was taken. Analysis of the outcomes of the discussions central to behavioural aggregation showed that there were seven classes of fault assessments:

- a. Assessments where the panel reached unanimity that the fault would inevitably lead to loss under ice, which included five faults where experts set all parameters of the distribution to 1.
- b. Assessments where the panel reached unanimity that the fault would have no impact at all on survivability, such as failure of a component or sub-system that would not be present for the Arctic.
- c. Assessments for faults where the experts considered that the phase of the mission may affect the consequence. This class of faults, typified by a failure in the vehicle control computer, resulted in vigorous discussion by the experts on the probability of loss, as some experts considered the outcome to be strongly dependent on the phase of the mission during which the fault occurred.
- d. Assessments of faults where individual experts shared particular insights affecting the aggregated outcome.
- e. Assessments where there was an agreement that the fault leads to a wide range of probability of loss.
- f. Assessments of faults that provided insights into instances of where a fault implied a consequential vulnerability.
- g. Assessments where the panel reached an agreement to use heuristic shortcuts. The group agreed collectively that they would spend little time on those faults that had a very low, but non-zero, consequence for the risk of loss. For these, they agreed on a standard distribution with a lower limit of 0, a median of 6.2×10^{-8} and an upper limit of 10^{-6} .

A novel aspect of this study was the quantification of the risk mitigation activities. For each fault experts were asked to assess of the impact of risk mitigation plan, setting P_M to 0 if they believed that the mitigation strategy would not mitigate the fault and to 1 if the mitigation plan would completely mitigate the fault, with intermediate values reflecting intermediate belief in mitigation effectiveness. When these judgements were plotted in an histogram three distinct distributions were observed.

Figure 14.5 Histogram of the assignments of probability of successful mitigation for the 51 faults, showing three distinct distributions.

One distribution, with a mode at zero, covered faults for which the experts agreed that the cause of the fault was unknown or unproven and thus experts were unconvinced that the proposed mitigation strategy would prove effective. The second distribution, with a sharp mode at 0.5, represented those faults the experts considered that, although the proposed solution was appropriate, the mitigation strategy had not been sufficiently tested or proven in field trials, or where a recurrence of a similar fault could not be ruled out. The third distribution had a mode at over 0.9, indicating a high to very high level of confidence by the experts that the causes of faults were well understood, the solutions known and tested.

The resulting probability of loss given the mitigation is calculated using the following expression $P(L|F,E,M) = P(L|F,E) \times (1 - P_M)$. Combining the assessments on $P(L|F,E)$ with P_M identifies those faults where $P(L|F,E)$ is high but P_M is low. These form an important subset of faults for the engineers to address. Most critical was a vehicle control computer configuration problem, where $P(L|F,E)$ was 1 and P_M was 0. All other faults where P_M was less than 0.1, $P(L|F,E)$ was less than 0.01, consequently, the need for effort into improving the understanding of the mitigation required was far less important. Of the 14 faults where $0.4 < P_M < 0.6$ eight were assessed with $P(L|F,E) > 0.5$. This was the most important set of faults for further investigation and improvement in P_M .

14.6.2 Survival Prediction

The risk assessments discussed in the previous section can be used for directing engineers towards where to put more effort in fault mitigation. However, the AUV risk of loss will vary with mission time or length. The decision to deploy an AUV in an extreme environment is better informed by a risk profile that captures the variation of risk with mission distance or time. The creation of such risk profile is possible by integrating the expert judgments of probability of loss with statistical survival models.

Survival modelling is very popular in medical statistics, where the survival function $S(t)$ is defined as the ratio of the number of individuals with survival times $\geq t$ to the total numbers of individuals in the study. These methods have been adopted for modelling reliability of systems and components [14.43].

An adequate and representative dataset is central to the application of statistical survival techniques. Some entries will consist of recorded time of failures; however there may be missions with no failures. When this is the case an entry is denoted as right-censored data.

There are several mathematical models for representing censored data; these are divided into parametric and non-parametric models. Parametric models assume that the failure history follows a particular shape whereas non-parametric models make no assumption with regards to the shape of distribution. In [14.12] the authors used the Weibull parametric model and the Kaplan Meier non-parametric estimator to model the probability of Autosub2 loss. For this analysis faults were discriminated between low impact faults and high impact faults. Only high impact faults were considered in the analysis as these were considered to lead to loss in the target operational environment. In their analysis the Kaplan Meier estimator was used in its usual form [14.12] [14.44].

$$\hat{S}(r) = \prod_{r_i < r} \frac{n_i - d_i}{n_i} \quad (5)$$

The survivor function $S(r)$ is defined as function of range r , where n_i is the number (of missions) at risk immediately prior to range r_i and d_i is the number of high-impact faults (or ‘deaths’) at range r_i . In the previous section we have shown that the uncertainty over whether a fault can lead to loss should, and can, be captured through expert subjective judgments. Thus instead of having a number of ‘deaths’ at time i , the statistical estimator must operate with the probability of loss given a fault in a given environment ($P(L|F_i, E)$). Brito et al. [14.10] derived an extended version of the Kaplan Meier estimator to account for this uncertainty:

$$\hat{S}(r) = \prod_{r_i < r} \left(1 - \frac{1}{n_i} \cdot P(L|F_i, E) \right) \quad (6)$$

The expressions for the variance for this estimator were also derived.

$$\hat{V}(r) = \frac{1}{(\log \hat{S}(r))^2} \sum_{r_i < r} \frac{p_i}{n_i(n_i - p_i)} + \hat{V}_{EEJ_r} + 2 \cdot \frac{1}{(\log \hat{S}(r))^2} \sum_{r_i < r} \frac{p_i}{n_i(n_i - p_i)} \cdot \hat{V}_{EEJ_r} \quad (7)$$

Where p_i is equal to $P(L|F_i, E)$, p_i is used above only to simplify the expression. $\hat{V}_{EEJ_r}$ is the variance in the expert judgements.

This estimator was used for calculating the risk of losing Autosub3 under the Pine Island glacier in January-February 2009. Before the campaign, risk estimates based on prior history and the extended Kaplan Meier approach helped devise a mitigation strategy based on a monitoring distance - reducing the risk to an acceptable level. In addition to a test mission prior to the start of the campaign the

vehicle, each run beneath the glacier included an initial 5 hr in open water where progress could be monitored [14.10]. The logic being that faults emerging during this time could be discovered while the vehicle was in open water, and recovery effected. The appropriate monitoring period being determined from the initial reduction in probability of survival in the Kaplan Meier plot (Figure 14.6). More recently the same approach was used for estimating the probability of losing an ISE Explorer vehicle in the Arctic. The risk profile for the ISE Explorer under sea ice is presented in Figure 14.6. Post campaign a risk assessment exercise was carried out to calculate the observed risk. This profile is presented in full black line in Figure 14.6. Results show that following the mitigation the maximum error between the predicted risk and the observed risk was 9% [14.14].

Figure 14.6 Kaplan Meier probability of survival against distance for the data set as considered by the experts with the pre-campaign assessment of individual fault mitigation included (in dashed line). In full black line is the observed risk during an ISE Explorer Arctic campaign.

14.7 Legal risks

The international rules for the law of the sea, as embedded in conventions such as the UN Convention on the Law of the Sea (UNCLOS) and the Convention on the International Regulations for Preventing Collisions at Sea (COLREGS) are clear with regards to the obligations of “ships” or “vessels”. These legal instruments were written and agreed before the use of autonomous ocean vehicles became widespread in the academic, military and commercial arenas. It is not surprising therefore that it is not clear at all whether, in law, autonomous ocean vehicles are indeed “ships” or “vessels”. The situation is further muddled by the possibility for a vehicle to be a “ship” for the purposes of one law, while not being a “ship” for another law. Consequently, the question, “Is an autonomous ocean vehicle a ship?” needs to be asked for each international Convention or national law. Only after the question has been answered each time can the legal issues, and the legal risks, within the scope of that Convention or law be fully understood. Unfortunately, as made clear in the exhaustive analysis in Brown and Gaskell [14.46] there is rarely a definitive answer to this question.

Given this situation, the vehicle owner or operator may face legal risks following incidents if their interpretation of laws or Conventions is found to be at variance with the opinions of courts. For example, if a court took a view that autonomous ocean vehicles were vessels for the purposes of the Standards of Training, Certification and Watchkeeping (STCW) Convention 1995 their autonomous use on the surface would probably be proscribed [14.46]. It is not surprising therefore that action is

being taken to prevent wholesale blocks on the use of unmanned systems in the ocean. These actions may be legal, procedural or technical.

First, examining the legal context, James Kraska, a lawyer with the US Judge Advocate General's office, considers that the way forward is to recognise autonomous ocean vehicles as "vessels" [14.47]. His opinion is that the "liberal legal architecture" of the key international Conventions would allow for inclusion of autonomous systems. If he were proved correct then they would indeed gain benefit from being considered "vessels". For example, the need to register an autonomous ocean vehicle as a vessel would provide for state protection. Rights in international law are given to "international persons", such as states, and not to individual artefacts, such as an autonomous ocean vehicle. But a vehicle could benefit from those rights if it was clearly linked to a state; hence the usual practice of registry under a particular state and flag. Henderson reinforced this point, "it is in the interest of the United States to establish the sovereignty of its UUVs to protect them from foreign seizure" [14.48].

Bork et al. [14.49] agree with Kraska that there is a need for action, and that a proper legal regime would give protection to autonomous (and free drifting) ocean vehicles, but they are not as sanguine that the existing framework of international maritime law can provide what is needed. They call for a new legal regime, but "it is not probable that one will soon come about".

One aspect of international law that is being followed by many is that of seeking permits for use of autonomous ocean vehicles in the exclusive economic zones of foreign nations, in accordance with UNCLOS Part XIII. Showalter gives an accessible summary of these issues [14.50]. Notable precedents are being set particularly relevant to the command and control of these vehicles. For example, the UK Foreign and Commonwealth Office has accepted that it is where the vehicle is piloted from that determines which nation applies for Diplomatic Clearance from the foreign nation, not the nationality of the vessel making the deployment. Of course, had the vehicle the benefit of registration and a flag state in its own right, this question would be moot, but at least this acceptance is coherent with the outcome if indeed the vehicle did have a flag state.

Given strict interpretations of international maritime conventions, where autonomous vehicles would be considered as vessels, they would be hard to adhere to, consequently, many organisations and users are preferring self-regulation through collective codes of practice. An early attempt was the Code of Practice for Autonomous Underwater Vehicles produced by the AUV Legal Working Group of the Society for Underwater Technology in 2000 [14.51]. With representation from military, industry and academic AUV users, the Code of Practice was a practical document that dealt with immediate issues of water-space management, environmental protection, operational considerations and training as well as vehicle marking and other matters. It was updated in 2008 with additional sections, among others, on deployment in congested waters, rules on salvage, protection of data, and damage to third parties, and broadened its coverage to autonomous ocean vehicles [14.52]. A succinct code of practice specifically for commercial use of AUVs concentrating on safe operations, mission planning and launch and recovery has been produced by the International Marine Contractors Association [14.53]. Research supported by the European Defence Agency (EDA), through the Safety and Regulations for European Unmanned Maritime Systems (SARUMS) programme, has led to an extensive draft best practice guide on handling, operations, design and regulations connected with unmanned maritime systems. Directed towards military use of these systems, the starting position in this paper is that these autonomous systems are vessels, and that "the main content of the COLREGS apply". An earlier EDA study on maritime unmanned surface vehicles summarised "what additional rules and guidelines that should be developed", which included adopting a classification scheme, sponsoring research into systems to enable the vehicles to comply with the COLREGS, and promoting cooperation with US agencies to examine whether these vehicles could become a specific class for the purposes of COLREGS.

There is active research into providing autonomous marine vehicles with both a rules-based approach to preventing collisions and with a common sense capability. Benjamin et al. examined a behaviour-based control framework, and demonstrated the system in operation on an autonomous surface platform [14.54]. An alternative, path planning, approach has been investigated and demonstrated by Naeem et al. focusing on manoeuvring and preventing head-on collisions [14.55].

There is a role for insurance in mitigating or limiting the financial consequences of legal risks, especially risks to third parties. The practical issues have been set out by Edwards [14.56], and Griffiths et al. provide case studies from scientific and industry users that have made insurance claims connected with autonomous ocean vehicles [14.57].

References

- 14.1 J.G. Bellingham, K. Rajan: Robotics in Remote and Hostile Environments, *Science* **318** (5853), 1098-1102 (2007)
- 14.2 J.A. Dowdeswell, J. Evans, R. Mugford, G. Griffiths, S.D. McPhail, N. Millard, P. Stevenson, M.A. Brandon, C. Banks, K.J. Heywood, M.R. Price, P.A. Dodd, A. Jenkins, K.W. Nicholls, D. Hayes, E.P. Abrahamsen, P.A. Tyler, B.J. Bett, D.O.B. Jones, P. Wadhams, J.P. Wilkinson, K. Stansfield, S. Ackley: Autonomous underwater vehicles (AUVs) and investigations of the ice-ocean interface: deploying the Autosub AUV in Antarctic and Arctic waters. *J Glaciol*, **54** (187), 661-672 (2008)
- 14.3 A. Jenkins, P.D. Stanley, S. Jacobs, S.D. McPhail, J.R. Perrett, A.T. Webb, D. White: Observations beneath Pine Island Glacier in West Antarctica and implications for its retreat, *Nat Geosci* **3**, 468-472 (2010)
- 14.4 T.B. Curtin, D.M. Crimmins, J. Curcio, M. Benjamin, C. Roper: Autonomous Underwater Vehicles: Trends and Transformations, *Mar Technol Soc J* **39**(3), 65-75 (2005)
- 14.5 D. Yoerger, A. Bradley, M. Jakuba, C.R. German, T. Shank, M. Tivey: Autonomous and remotely operated vehicle technology for hydrothermal vent discovery, exploration, and sampling. *Oceanography* **20**(1), 152-161 (2007)
- 14.6 M.V. Jakuba, C.N. Roman, H. Singh, C. Murphy, C. Kunz, C. Willis, T. Sato, and R.A. Sohn, R. A: Long-baseline acoustic navigation for under-ice autonomous underwater vehicle operations. *J. Field Robotics*, **25**(11-12), 861-879. (2008)
- 14.7 J.E. Strutt: *Report of the inquiry into the loss of Autosub2 under the Fimbulisen*, National Oceanography Centre Southampton Research and Consultancy Report 12 (National Oceanography Centre, Southampton 2006)
- 14.8 L. Lippsett: <http://www.whoi.edu/oceanus/viewArticle.do?id=70906>. Accessed May 31st 2012.
- 14.9 G. Griffiths, A. Trembanis: Towards a risk management process for Autonomous Underwater Vehicles, in: G. Griffiths, K. Collins: *Proc. Masterclass on AUV Technology for Polar Science* (Society for Underwater Technology, London 2007) pp. 103-118
- 14.10 M.P. Brito, G. Griffiths, P. Challenor: Risk Analysis for Autonomous Underwater Vehicle Operations in Extreme Environments, *Risk Anal.* **30**, 1771-1788 (2010)
- 14.11 R. Stokey, T. Austin, C. von Alt, M. Purcell, R. Goldsborough, N. Forrester, B. Allen: AUV bloopers or why Murphy must have been an optimist, *Proc. 11th International Symposium on Unmanned Untethered Submersible Technology*, New Hampshire, 1999 (Autonomous Undersea Sys. Inst., Durham 1999) pp. 32-40
- 14.12 G. Griffiths, N.W. Millard, S.D. McPhail, P. Stevenson, P.G. Challenor: On the Reliability of the Autosub Autonomous Underwater Vehicle, *Underwater Technol.* **25**, 175-184 (2003)

- 14.13 K. Podder, M. Sibenac, H. Thomas, W. Kirkwood, J.G. Bellingham: Reliability growth of Autonomous Underwater Vehicle – Dorado, Proc. Oceans 2004, Kobe, 2004 (IEEE, Piscataway 2004) pp. 856–862
- 14.14 M.P. Brito, G. Griffiths, J. Ferguson, D. Hopkin, R. Mills, R. Peterson, E.A. McNeil: A behavioural probabilistic risk assessment framework for managing autonomous underwater vehicle deployments, J. Atmos. Ocean. Tech. doi: <http://dx.doi.org/10.1175/JTECH-D-12-00005.1> (in press)
- 14.15 P.D.T O'Connor: *Practical Reliability Engineering* (Wiley, Chichester 1995)
- 14.16 G. Griffiths, N. Millard, S.D. McPhail, P. Stevenson, P. Challenor: On the reliability of the Autosub Autonomous Underwater Vehicle, Southampton Oceanography Centre Research and Consultancy Report 48, 1-45 (2000)
- 14.17 G. Griffiths, N.W. Millard, S.D. McPhail, J. Riggs: Effect of upgrades on the reliability of the autosub AUV, Proc. 13th International Symposium on Unmanned Untethered Submersible Technology, New Hampshire, 2003 (Autonomous Undersea Sys. Inst., Durham 2003)
- 14.18 J.D. Andrews, T.R. Moss: Reliability and risk assessment. Longman, Harlow, Essex, 1993.
- 14.19 M.V. Frank: Reentry safety: probability of fuel release, Proc. European Safety and Reliability Conference, Rotterdam, 1999 (Taylor & Francis, Rotterdam 1999)
- 14.20 G. Griffiths, M.P. Brito: Risk Management for Autonomous Underwater Vehicles Operating Under Ice. In Proc. Arctic Technology Conf., Houston, Texas 2011 (Offshore Technology Conference, Texas 2011).
- 14.21 L. Merckelbach: On the probability of underwater glider loss due to collision with a ship. J. Mar. Sci. Tech., doi: 10.1007/s00773-012-0189-7 (2012)
- 14.22 D.P. Horner, A.J. Healey, S.P. Kragelund: AUV experiments in obstacle avoidance. Proc. of MTS/IEEE OCEANS, Monterey 2005 (MTS/IEEE, Monterey 2005), pp. 1464-1470
- 14.23 S.D. McPhail, M. Furlong, M. Pebody: Low-altitude terrain following and collision avoidance in a flight-class autonomous underwater vehicle, Proc. Inst. Mech. Eng. M J. Eng. Marit. Environ. **224** (4), 279-292 (2010)
- 14.24 W. Stone, N. Fairfield, G.A. Kantor: Autonomous Underwater Vehicle Navigation and Proximity Operations for Deep Phreatic Thermal Explorer (DEPTHX), G. Griffiths, K. Collins (Eds): Proc. of an Int. Masterclass in AUV Technology for Polar Science, National Oceanography Centre, Southampton, UK, 28-30 March 2006 (Society for Underwater Technology, London 2007) pp. 127-142
- 14.25 M. Pebody: Autonomous underwater vehicle collision avoidance for under-ice exploration, Proc. Inst. Mech. Eng. M J. Eng. Marit. Environ. **222** (2), 53-66 (2008)
- 14.26 W. Feller: *An Introduction to Probability Theory and its Applications*, vol. 1. (John Wiley and Sons, New York 1950)
- 14.27 M. Brito, G. Griffiths: A Markov Chain state transition approach to establishing critical phases for AUV reliability, IEEE J. Oceanic Eng. **36** (1), 139-149 (2011)
- 14.28 R.L. Keeney, D. Winterfeldt: Eliciting probabilities from experts in complex technical problems, IEEE Trans Eng Manag. **38** (3), 191-201 (1991)
- 14.29 A. Tversky, D. Kahnemann: Judgment under uncertainty: Heuristics and biases, Science **185** (4157), 1124-1131 (1974)
- 14.30 D. Kahnemann, A. Tversky: Subjective probability: A judgment of representativeness, Cogn Psychol **3**(3), 430-454 (1972)
- 14.31 M.L. Finucane, A. Alhakami, P. Slovic, S.M. Johnson: The effect of heuristic in judgments of risk and benefits, J. Behav. Decis. Mak., **13**(1), 1-17 (2000)
- 14.32 M. Kynn: The 'heuristics and biases' bias in expert elicitation, J. Royal Stat. Soc. STA, **171** (1), 239-264 (2008)

- 14.33 R.M. Cooke, L.H.J. Goossens: Expert judgement elicitation for risk assessments on critical infrastructures, *J. Risk Res.* **7**(6), 643-656 (2004)
- 14.34 D.J. Isenberg: Group polarization: A critical review and meta-analysis, *J. Pers. Soc. Psychol.* **50**(6), 1141-1151 (1986)
- 14.35 P.A. Morris: An axiomatic approach to expert resolution, *Manage. Sci.* **29**(1), 24-32 (1983)
- 14.36 R.T. Clemen, R.L Winkler: Combining probability distributions from experts in risk analysis, *Risk Anal.* **19**(2), 187-203 (1999)
- 14.37 K.J. McConway: Marginalisation and linear opinion pools, *J. Am. Stat. Assoc.* **76**, 410-414 (1981)
- 14.38 M.A. Moline, S.M. Blackwell, C. Von Alt, B. Allen, T. Austin, J.F. Case, N. Forrester, R. Goldsborough, M. Purcell, R. Stokely: Remote environmental monitoring units: An autonomous vehicle for characterizing coastal environments, *J. Atmos. Ocean Tech.* **22**, 1797-1808 (2005)
- 14.39 G. Griffiths, M. Brito, I. Robbins, M. Moline: Reliability of two REMUS-100 AUVs based on fault log analysis and elicited expert judgment. *Proc. Int. Symposium on Unmanned Untethered Submersible Technology*, Durham, New Hampshire 2008 (Autonomous Undersea Systems Institute, New Hampshire 2008).
- 14.40 A. O'Hagan, C.E. Buck, A. Daneshkhah, J.R. Eiser, P.H. Garthwaite, D.J. Jenkinson, J.E. Oakley, T. Rakow: *Uncertain judgments: Eliciting experts' probabilities* (John Wiley, Chichester, UK, 2003)
- 14.41 L.D. Phillips, M.C. Phillips: Facilitated work groups: Theory and Practice, *J. Oper. Res. Soc.* **44**, 533-549 (1993)
- 14.42 T. Crees, C. Kaminski, J. Ferguson, J.M. Laframboise, A. Forrest, J. Williams, E. MacNeil, D. Hopkin, R. Pederson: UNCLOS under ice survey - An historic AUV deployment in the Canadian high arctic, *Proc. OCEANS 2010, Seattle, 2010* (IEEE/MTS, Seattle 2010) pp. 1-8
- 14.43 R.L. Prentice, J.D. Kalbfleisch: *The Statistical Analysis of Failure Time Data* (Wiley, New York 2002)
- 14.44 E.L. Kaplan, P. Meier: Nonparametric estimation from incomplete observations, *J. Am. Stat. Assoc.*, **53**(282):457-481 (1958)
- 14.45 A. Jenkins, P. Dutrieux, S. Jacobs, S. McPhail, J. Perrett, A. Webb, D. White: Autonomous underwater vehicle exploration of the ocean cavity beneath an Antarctic ice shelf, *Oceanography* **25**(3), 202-203 (2012)
- 14.46 E.D. Brown, N.J.J. Gaskell: *Questions and answers on the law, state practice and procedure relating to autonomous underwater vehicles* (Society for Underwater Technology, London 2001)
- 14.47 J. Kraska: The law of unmanned naval systems in war and peace, *J. Ocean Tech.* **5**, 43-68 (2010)
- 14.48 A.H. Henderson: Murky Waters: the legal status of unmanned undersea vehicles, *Naval L. Rev.* **55** (2006)
- 14.49 K. Bork, J. Karstensen, M. Visbeck, A. Zimmermann: The legal regulation of floats and gliders: in quest of a new regime? *Ocean Dev. Int. Law* **39**, 298-328 (2003)
- 14.50 S. Showalter: The legal status of autonomous underwater vehicles, *Mar. Technol. Soc. J.* **38**, 80-83.
- 14.51 SUT: The operation of autonomous underwater vehicles – recommended code of practice – Volume 1 (Society for Underwater Technology, London 2001)
- 14.52 SUT: Recommended Code of Practice for the Operation of Autonomous Marine Vehicles, 2nd edition (Society for Underwater Technology, London 2009)
- 14.53 IMCA: Considerations for the safe operation of autonomous underwater vehicles (AUVs) (International Marine Contractors Association, London 2009)

- 14.54 M.R. Benjamin, J.J. Leonard, J.A. Curcio, P.M. Newman: A method for protocol-based collision avoidance between autonomous marine surface craft, *J. Field Robot.* **23**, 333–346 (2006)
- 14.55 W. Naeem, G.W. Irwin, A. Yang: COLREGs-based collision avoidance strategies for unmanned surface vehicles, *Mechatronics* <http://dx.doi.org/10.1016/j.mechatronics.2011.09.012> (2011)
- 14.56 S.R.P. Edwards: AUVs. How do you insure them? *Int. Ocean Sys. Des.* **4**, 10-12 (2000)
- 14.57 G. Griffiths, N. Bose, J. Ferguson, R. Blidberg: Insurance for autonomous underwater vehicles, *Underwater Technol.* **27**, 43–48 (2007)