1992

AN EXPERT SYSTEMS APPROACH TO COLLISION AVOIDANCE

Grahame Kenneth Blackwell

PHD

AN EXPERT SYSTEMS APPROACH TO COLLISION AVOIDANCE

Blackwell, Grahame Kenneth

Award date:
1992

Awarding institution:
University of Plymouth

AN EXPERT SYSTEMS APPROACH TO COLLISION AVOIDANCE

G. K. BLACKWELL

Ph. D. 1992

Sunrise over Guernsey

(Taken on a cross-channel Decca familiarisation trip aboard the 7 metre yacht 'Dulcinea')

AN EXPERT SYSTEMS APPROACH TO COLLISION AVOIDANCE

Grahame Kenneth Blackwell BTech, PGCE

This thesis is submitted to the Council for National Academic Awards in partial fulfilment of the requirements for the degree of Doctor of Philosophy.

Sponsoring Establishment:

University of Plymouth
Drake Circus
Plymouth
DEVON

Collaborating Establishment:

Plessey Marine Research Unit
Templecombe
SOMERSET

July 1992

DECLARATION

No part of this thesis has been submitted for any award or degree at any other institute.

While registered as a candidate for the degree of Doctor of Philosophy the author has not been a registered candidate for another award of the CNAA or of a university.

ACKNOWLEDGEMENTS

My thanks are due to the following people for their help and support during the preparation of this thesis:

my Director of Studies, Dr. Colin Stockel, for his invaluable guidance, encouragement and constructive criticism;

Dr. Patricia Pearce, my Head of Department, for her practical support in provision of time, resources and opportunities to attend relevant conferences;

Dr. Bruce Colley, Dr. Roland Burns and Dr. Mike Dove, for their advice and involvement at various stages of the work;

Dr. David Norman of Britannia Royal Naval College, Dartmouth, for providing vessels and crew, enabling this research to take to the high seas;

Kelvin Hughes Ltd, for their generous provision of electronic equipment for the test vessel;

my family, for their understanding and encouragement throughout;

and, not least, all of those students who have suffered variously my preoccupied ramblings and my neglect, but have nevertheless continued to show interest and support for my work.

AN EXPERT SYSTEMS APPROACH TO COLLISION AVOIDANCE

Grahame Kenneth Blackwell BTech, PGCE

ABSTRACT

Any object or entity moving in an open environment requires, together with certain navigational skills, the capacity to avoid collisions. Such ability presupposes: manoeuvrability; awareness of the nature of, and ability to detect, potential hazards; and innate intelligence to formulate, evaluate, and implement appropriate avoidance strategies.

The maritime environment, and vessels moving within it, provide a classic illustration of the problems and principles of collision avoidance in two-dimensional space. Events in recent years, leading to loss of life, wastage of resources and pollution of the environment, have highlighted the limited success being achieved in this area. In most cases human error has been identified as the major factor; errors of judgement consequent on either a false sense of security or (more often) information overload - an excess of data from various sources, to be processed in limited time.

This thesis describes the development of a shipboard expert system for marine collision avoidance. This system, to be operated in a microcomputer on the bridge, analyses all pertinent data on an encounter with another vessel and formulates an optimal strategy to deal with the situation. Detailed information on own-ship's characteristics, including a mathematical manoeuvring model, is held in the knowledge base. Situational data relating to own-ship, potential hazards and any other relevant factors are to be input via on-board sensors, including ARPA (Automatic Radar Plotting Aid); keyboard input is an additional option. The user interface makes extensive use of an advanced WIMPS (Windows-Icons-Mouse-Pointer System) environment; scrollable text and graphics windows provide the main display; menus and dialogue boxes provide for input and control functions; all user facilities are mouse-operated.

The system logic incorporates facilities for emergency manoeuvring, and the potential for extension to handle multi-ship encounters; direct input and consideration of electronic chart data on coastal features and shipping lanes has also been anticipated in the design. These issues are considered in some depth in the thesis, and form the basis of two new research projects now in progress.

The system has been validated for two-ship encounters in the open sea, through expert appraisal of simulated manoeuvres and random simulation of large numbers of encounters; sea trials on board a research vessel have shown it to be seaworthy and effective in practical operation.

CONTENTS

CHAPTER 1

INTRODUCTION

"The distance doesn't matter; it is only the first step that is difficult." Marquise du Deffand

1.1 Collision Avoidance As A Way of Life

It is perhaps stating the obvious to observe that the process of entropy - the irreversible levelling of energy states throughout the universe - is directly attributable to collisions of all types and magnitudes. At the macrocosmic level, celestial bodies and/or clouds of dust and gases interact in collisions in a way that reduces the net potential of those systems; the rock pounding at a cliff face under the force of heavy seas both reduces itself to shingle and erodes the land mass (at the same time dissipating wave energy); at the molecular level, Brownian motion illustrates the transfer of energy from fast-moving to slower-moving particles (also evident if one holds a metal poker in the fire too long). In every case, order reduces to disorder, individuality to bland uniformity.

Life, in its myriad forms, is characterised by a high level of order and differential energy levels within its structure. External influences which affect that order, or interfere with those energy levels, will cause the degradation (possibly the termination) of that life force within any organism.

For some lower life forms, protection against such disruption is afforded by a combination of low propulsion velocity (or a static existence), coupled with either a durable outer shell or a suitably elastic body structure. All other organisms incorporate sensory and motor functions which should, in the normal course of events, enable them to avoid damaging or life-threatening collisions with fixed or moving objects (including other organisms). Furthermore, such life-forms have evolved strategies for identifying potential collisions and manoeuvring safely through such encounters. The ability of the humble housefly to avoid a fast-moving newspaper, and the tiny pipistrelle bat manoeuvring at speed through densely-packed tree branches, both illustrate this near-universal concept.

The human species is generally regarded as being the most highly developed in terms of reasoned (as opposed to instinctive) action. We have also adopted techniques and strategies observed in other species; for example, radar, a major feature of this thesis, is but a variant of the navigation instrumentation fitted as standard in the aforementioned bat. In this vital field of collision avoidance, therefore, a study of human behaviour patterns is a useful preparation for any attempt to automate or program the process.

Certain general principles may be gleaned from personal experience by anyone who has walked down an average shopping street. If the following observations, drawn from such experience, seem so obvious as to be not worth stating, it should be remembered that any expert system must be built upon such fundamental truths.

1) Potential hazards may be classified as 'static' (lamp-posts, pillar boxes, etc.) or 'dynamic' (another person, dog, etc.). A static hazard will not ever take avoiding action itself, but neither will it increase the problem by taking untoward action; by contrast, a dynamic hazard may reduce or exacerbate the problem by its own actions.

2) If a static hazard is not on, or very close to, one's set course, it may be ignored - until, possibly, it becomes a consideration because of a change of course. If such a hazard is on or close to one's course (either originally or due to a course change), then that course must be adjusted so that this is no longer the case.

3) A dynamic hazard may be ignored until it comes within reasonable range - it is not normal or reasonable to concern oneself with the potential hazard posed by every other pedestrian within sight! Note that the term 'reasonable range' here implies a time constraint - a skateboarder bearing down on one from 30 metres away would probably merit attention, whereas an old lady 20 metres away would not. Once that hazard comes within such range a plan is formulated (usually subconsciously) as to appropriate avoiding action, if any, and put into action immediately.

4) There are no set priorities or protocols for pedestrian collision avoidance. Simultaneous avoiding actions by two people may cancel each other out; another person may fail to take action as reasonably expected, or even create a potential collision situation unexpectedly at a late stage through negligence or misunderstanding. Individuals tend to adopt 'personal strategies' such as always walking on the left, always taking clear give-way action at an early stage, or (conversely) always acting on the assumption that the other party will give way. In general, however, right of way will be given to a pedestrian with 'restricted ability to manoeuvre' - old lady with walking frame, mother pushing pram, etc.

5) The task of evading an anticipated collision with a dynamic hazard involves: (a) assessment of velocity vectors for self and hazard, and relative displacement vector for hazard; (b) extrapolation of tracks to likely point of impact; (c) consideration of various collision avoidance strategies - turning (which way? how far?), change of pace (faster? slower? how much?), vocal warning (last resort!) - with due regard for other hazards, both static and dynamic; (d) careful attention to restricted traffic lanes (pavements) where appropriate; (e) continual vigilance throughout the 'manoeuvre', in case hazard takes unexpected action. All this without conscious thought, whilst numerous other issues are being 'processed' at a conscious level - shopping lists, bus time-tables, cash-in-hand, plans for evening meal, etc.

6) It appears to be common practice for pedestrians to attempt to preserve a 'personal zone' around them, free of hazards (British reserve?), and to conduct manoeuvres on that basis. This zone would appear to be variable, reducing in size in situations of high-density traffic.

1.2 Control and Guidance of Mechanised Systems

In attempting to order and structure his environment, man has introduced a variety of transport systems which (generally) operate under intelligent control. Since such systems are similarly capable of sustaining, and causing, reduced functionality as a consequence of a collision, that guiding intelligence must take responsibility for any necessary avoidance manoeuvres. Depending on circumstances, the problem may be set in one, two or three dimensions (examples: railways, road traffic, and aircraft respectively), which offer varying degrees of scope for avoiding action. Typical speeds at which such action must be taken will also vary considerably between different types of moving objects; this will affect the nature and severity of the avoidance problem, as will traffic density.

The overriding factor in determining strategy must be the cost function associated with any collision, in terms of both damage to property and risk to human life - offset by costs of avoidance, such as loss of time, fuel wastage etc. The extent of the damage in any collision, and thus the cost, is likely to be largely speed-dependent. Other factors, such as inertia and resilience/ fragility, also enter the cost equation in addition to intrinsic value. A further consideration in certain cases must be second-stage effects of collision - consider the end result of a collision involving loss of a relatively inexpensive section of wingtip from an airliner in flight, or the consequences of an otherwise 'minor' collision which happens to involve a truck loaded with high explosive.

It would not do to carry the parallel too far between pedestrian encounters and those involving mechanised forms of transport. However, most of the concepts highlighted in the previous section are fairly universal: the 'personal space', also termed the 'domain'; the time-based 'area of concern'; vector calculations, extrapolations, and strategies for avoidance (as detailed in chapter 5). The major difference is the cost function - low (in general) for pedestrian collisions, varying from substantial to astronomical in the case of the various forms of transport. For this reason, the collision avoidance protocols for mechanised transport are rigidly formalised and, in most cases, enforced by the rule of law.

Despite such formalisation, collisions still occur, all too often. Almost without exception, the cause is identified as 'Human Error', indicating either misinterpretation of the situation or failure to apply the appropriate rules correctly. In some cases this may be put down to lack of vigilance, in others confusion due to an inability to assimilate all of the relevant data from various sources - often at the same time as having to handle various other vehicle management tasks. Whichever is the case, the situation would clearly be eased if the task of collating the data, identifying the situation, and advising on the appropriate strategy could be delegated to an automated system. This system would be required to show the same expertise, in terms of background knowledge and intelligent decision-making, as the human expert practitioner. Such a system - an Expert System for Collision Avoidance - is the subject of this research.

Rather than taking a 'broad brush' approach to the whole realm of collision avoidance, this thesis addresses one particular area of application, namely the maritime setting. There is a very real need for such automated decision support for mariners, shipping losses being in the order of 1000 tons per week due to collision or grounding [1]. The information handling requirements of sea travel are well suited to support by an Intelligent Knowledge-Based System (IKBS).

Life on board ship may vary between two extremes, in terms of the information being received and analysed in respect of prevailing conditions. On exit from, or entrance to, a port, or whilst negotiating a busy seaway, there is a plethora of information arriving at the bridge from a wide variety of sources - radar, radio, navigation systems, communication from other parts of the ship, observations from those on watch. Such information must be collated and translated into appropriate action. It has been observed [2] that in cases of 'information overload' the human brain will tend to blank out elements of information which it cannot handle - possibly information vital to the safety of the ship and her crew.

At the other extreme, there are periods on a sea passage where no hazardous situation presents itself for a long time; no other vessels or charted features give cause for concern, and the vigilance of port entry or exit may be relaxed somewhat. It is at just such times that an unexpected situation may catch the unwary mariner off his guard, and disaster ensue - witness the collision between the trawler 'Dionne Marie' and the tanker 'Rose Bay' on Cup Final Day (!) 1990.

A number of other considerations support the selection of the maritime field as the focus of this research:

(1) it offers a 2-dimensional environment, with the possibility of complete freedom of movement within those dimensions - but additionally circumstances which may restrict such movement;

(2) there already exists a comprehensive set of guidelines for encounters involving marine vessels, on which to base a rule structure;

(3) a wealth of experience and documentation is available to give further definition to that rule structure;

(4) the research base (Polytechnic South West) has the wherewithal and the contacts to provide facilities for testing, enhancing and validating such an expert system:- advanced simulator, research vessel, maritime industrial liaison, well-qualified mariners;

(5) Polytechnic South West is already a thriving centre of maritime research, notably in the areas of navigation and control. Consequently an experienced and well-informed supervision team is readily available to give direction and support.

Points 1-3 identify an environment offering substantial freedom of movement, but within an organised framework of guidelines and experience - an ideal basis for the development of fundamental principles of inference and strategy. Points 4 and 5 highlight the practical considerations governing the choice of scenario.

1.3 The Maritime Collision Avoidance Task

1.3.1 Outline

Collision avoidance at sea is, almost by definition, a rule-based task. The International Regulations for Preventing Collisions at Sea [3] would appear to provide a clear framework, supported by the force of law, for a simple procedural approach: whatever the situation, apply the appropriate rule, and the problem is solved without incident.

The reality is, of course, rather different, as evidenced by the regular losses of multi-kiloton ships in collision incidents. Such losses are due not only to incompetence or fatigue, but also to differing interpretations of a given situation by the masters of the vessels involved. The 'Collision Regulations' are broad guidelines, rather than rigorous definitions of preconditions and consequent actions. Details as to timing, clearances, suitable course alterations, are left very much to the discretion of the mariner - discretion which is perfected by years of accumulated experience and wisdom. Furthermore, the variety of 'non-standard' situations - particularly those involving a number of vessels - is such that any rule book attempting to cover all such contingencies would be far too cumbersome to provide a useful reference.

In defining a rule structure for this task, it would seem appropriate to first identify 'rules of thumb' for specific types of situation which occur most frequently; supplementary rules may then be added for variations in these situations. This rule structure would be based on the aforementioned regulations, and would incorporate the accumulated wisdom of expert mariners. There would also be a need to tailor the rules to reflect the response characteristics of the system operating that rule structure - speed of response, breadth of information available to the system, confidence limits (including possible effects of misjudgement).

Such a system would require access to fixed characteristics of the vessel (in some cases, specific to current voyage): length, beam, maximum speed, turning curves, safe clearing distance, and a variety of technical data; relevant chart data for the area being navigated could also be considered as similarly 'static' (fixed) information, though of a rather different type. It would also need to be able to receive current values at all times of other, dynamically-varying, parameters: speed, course, rudder setting, etc, plus position, speed and course of any potential hazards in the vicinity. There is, of course, some degree of overlap between these two types of data, static and dynamic: for instance, one's assessment of a safe clearing distance may depend on the nature of the hazard to be cleared.

1.3.2 System Hardware and Software

The task as defined is clearly susceptible to handling by means of a suitable expert systems package, ideally running on a microcomputer system located on the bridge of the ship. However, two major considerations set this task apart from most conventional applications of artificial intelligence (AI) experienced to date.

(1) The dynamic information required must, by its very nature, be input to the system directly via a range of sensors, including such instrumentation as radar; furthermore, the scope and limitations of such data must be given due regard by the expert systems logic.

(2) The decision processes involved require extensive mathematical calculations, primarily of trigonometrical functions, to be carried out in real time.

These considerations would appear to indicate the need for somewhat specialist software tools, operating at near-optimum processor speed but with capability for advanced forms of user interaction. Such requirements tend to preclude standard expert system 'shells' currently available (see chapter 3), and would seem to be mutually exclusive in terms of programming languages, suggesting a two-pronged approach which combines:

(a) a 'calculation' module, to evaluate the trigonometrical and other mathematical functions which form the basis of the decision-making process; and

(b) a 'front-end' module, to operate the rule base and provide a 'user-friendly' form of communication with the mariner.

(a) demands powerful mathematical tools to evaluate a large number of multivariate functions in real time, plus facilities for direct sensory input; (b) requires an AI-style approach to 'fire' (i.e. action) the rule structure, plus a sophisticated and versatile man-machine interface (MMI) - for example, a WIMPS (Window-Icon-Mouse-Pointer System) environment. Three alternative approaches have been considered:

(1) linked modules in two different languages, such as 'C' and Smalltalk (an object-oriented language);

(2) use of an extended language, such as Objective C, offering all the power of 'C', plus MMI and object-oriented add-ons;

(3) an inference engine (the driving logic of the expert system) written in 'C', supported by an advanced WIMPS-based operating environment capable of firing the rules at appropriate times.

In every case 'C' was identified as the ideal language for the mathematical and sensor/control aspects of the task. Options (1) and (2) favour an object-oriented approach to the user interface and actioning of the rules; option (3) relies on the system's operating environment to provide such facilities.

Consideration of an OOPS (Object-Oriented Programming System) approach to the rule base and the user interface was prompted by the observation that many of the entities in the process conform to the OOPS concept of 'objects': ships and rules, as well as windows and icons, may all be usefully classified as 'objects' - indeed, the language Smalltalk grew out of a need for a suitable form in which to define and manipulate a WIMPS environment [4]. However, such entities may also be defined as 'structures' in the standard 'C' language, and (under the new ANSI standard) processed very effectively and efficiently. The two main strengths of OOPS - the concepts of 'class' and 'inheritance' - do not feature largely in the system requirements, and the overheads demanded by such facilities would seem disproportionate to the benefits offered. Consequently, option (3) was selected.

Initial prototyping was undertaken on an Atari ST microcomputer, running under GEM (Graphics Environment Manager). At times, two or three Ataris were linked to provide a 'parallel processing' simulation environment [5]. At a later stage, development was transferred to the Acorn Archimedes, a RISC (Reduced Instruction Set Computer) system. This system was selected for the following reasons:

(a) powerful processing capability at relatively low cost;

(b) a sophisticated multi-tasking WIMPS environment, providing all of the required facilities, accessible through 'C' primitives;

(c) facilities for input/output from/to a variety of devices - sensors, motor controls, communications channels, etc.

(d) compatability with other developments in the ship control field, within the same research group.

1.4 System Development - Rationale

Clearly, the ultimate objective of a system such as the one under consideration here is to negotiate safely (and efficiently) through an encounter situation of any complexity. This implies the capability to deal competently with an encounter involving a number of other vessels of varying types, one or more of which may not be acting in conformity with the regulations, whilst constrained by the need to hold to a traffic separation system and/or negotiate other navigational hazards (e.g. coastal features); the possibility of one's own vessel having to operate under conditions of impaired manoeuvrability should not be discounted, and manoeuvres should be conducted with at least the degree of efficiency that one would expect of an experienced ship's master.

In practice, the task as defined must be ordered in ascending levels of complexity, starting with the simplest possible scenario. It is, of course, essential that at every level the way is left open for any enhancements which may follow (upward compatability). This means:

(a) an initial structure and methodology which make no presuppositions as to what may, or may not, be appropriate or needful in terms of future development; and

(b) since (a) is, at least in the initial stages, an unrealistic ideal, the overall approach must be susceptible to a generative process of 'stepwise refinement', i.e. highly structured in such a way that innovations may be accommodated without major upheaval.

It should be emphasised that these remarks do <u>not</u> refer to enhancements of the rule base in the 'finished' system, based on new experience. The intention is that such enhancements would be incorporated without any modification of the overall system - hopefully via a 'user-friendly' utility to be designed for this purpose (see section 11.4).

The work outlined in this thesis is based upon the following perception of appropriate stages of development:

(i) a system capable of making decisions relating to two vessels in the open sea (own-ship and one other vessel, no additional hazards or constraints), both acting in conformity with the regulations under similar time and distance criteria for decision-making; the required outcome to be a safe resolution of the encounter situation;

(ii) extension of (i) to accommodate the possibility that the hazard vessel may act in a manner likely to necessitate emergency manoeuvring on the part of own-ship; e.g. failing to give way when the regulations indicate that it should do so; altering course in such a way as to increase the risk of collision;

(iii) provision of an 'optimisation' facility, which considers various alternative strategies and selects the one which minimises some appropriate cost function;

(iv) extension of the decision logic to consider a number of potential hazard vessels, and select an optimal (or near-optimal) strategy to negotiate these various hazards; such a strategy should be in the form of a 'flexible response', amenable to revision in the light of new pertinent data (e.g. another potential hazard coming within radar range);

(v) incorporation of a software interface to an electronic chart data-base suitably structured for interrogation by a computer-based decision system; extension of the rule structure to include consideration of such chart data in its decision-making;

(vi) consideration of possible consequences of impaired functionality in any control or sensory subsystem of own-ship (e.g. rudder, engines, radar); adjustment of strategy planning,as appropriate.

Design of a suitable test environment for such development constitutes an integral part of the research work described here. Various forms of computer-simulated test environment are described in chapter 8, together with the rationale for their use at different stages of development.

Successful resolution of stages (ii)-(vi) depends wholly on the nature of the implementation of stage (i). This thesis describes the formulation, design and construction of such a two-ship system; the system has been so structured as to facilitate anticipated further developments. The underlying philosophy of the expert system shell - inference engine and rule structure - is based on an original concept which ensures optimum safe manoeuvring in any encounter situation (see chapter 6).

Stages (i) - (iii) have been successfully implemented: both standard and emergency manoeuvres are dealt with in a satisfactory manner, with accommodation for increasing levels of sophistication in both areas; the optimisation facility is adaptable in terms of selection and weighting of relevant criteria. In addition, a theoretical basis for stage (iv) is developed in detail in this thesis; that work has since been implemented in a practical extension of the system [6, 7]. Stage (v) depends on major advances in the field of electronic charting; consideration of this aspect of the system, and associated technologies, is the subject of a further research project.

1.5 Organisation of Thesis

The contents of the succeeding chapters of this thesis are organised as described below. The sequence of these chapters largely mirrors the progression of ideas - either structurally or chronologically - leading to the current phase of development. However, chapters 7 and 8 encompass attributes of the system which have evolved continually throughout the work described here; they should therefore be seen as lying in parallel with the other chapters, and their positioning as a matter of expediency rather than natural sequence.

Chapter 2: Fundamental Concepts and Principles

This chapter outlines concepts identified by previous research, through consideration of the practicalities of collision avoidance and observation of good practice - the latter specifically in the maritime setting. It also sets out the guiding principles which have been established for effective avoidance of a collision or near-miss at sea; this includes the regulations published by the International Marine Organisation, but also covers observations regarding interpretations of those regulations and situations of possible uncertainty.

Chapter 3: The Expert System 'Shell'

This covers the design and functioning of the inference engine, the system logic which 'drives' the rules, providing the reasoning capability which characterises an expert system. The structure and format of those rules, and their logical connectivity, is described in detail, without being proscriptive about their actual content. The potential for expansion of this rule base, either to deal with new situations or to provide more sophisticated handling of situations already covered, is outlined.

Chapter 4: The Knowledge Base

The information necessary for intelligent decision-making is drawn from a variety of sources. Some data items are invariant, at least for the duration of the voyage; other data values vary dynamically as the ship proceeds on its passage. This chapter identifies relevant data and their sources; it also considers methods of data acquisition, and consequent considerations with respect to the decision logic.

Chapter 5: Rules for the Two-Ship Situation

The collision avoidance regulations provide an initial basis on which to build a set of rules, or guiding principles, governing safety manoeuvres. By themselves, however, they do not form a complete and coherent framework for such decision-making; they require 'fleshing out' with detail built up from a combination of observed good practice and commonsense. Chapter 5 gives an overview of practical considerations governing two-ship encounters in the open sea, and how such encounters may be safely negotiated in accordance with the regulations.

Conventional collision avoidance manoeuvres are planned on the initial premise that other vessels in the vicinity are under the control of competent and well-intentioned mariners who will act in conformity with the regulations. Since this is not always the case, any strategy must include contingency plans for untoward action by another vessel. The latter part of this chapter considers that requirement in relation to a computer-based decision system.

Chapter 6: Look-Ahead Simulation for (a) Safe Manoeuvring
(b) Optimal Manoeuvring

Any collision avoidance manoeuvre is a considered balance between safety on the one hand and minimum loss of utility on the other. The economic term 'utility' is here taken to mean a weighted combination of speed, proximity to destination, adherence to planned course, and other factors deemed pertinent by those having an interest in ship and/or cargo.

As a minimum requirement, any such manoeuvre should conform to both the regulations and the principles outlined in chapter 2. Since every encounter is unique in terms of relative positions, velocities and manoeuvring capabilities of the vessels involved, the time and sea room needed for safe resolution of each situation varies widely. Chapter 6 deals first with a strategy for identifying, well in advance, a time to initiate a manoeuvre such that it will resolve the encounter safely and successfully, without excessive course diversion. This chapter goes on to develop a technique for identifying the optimum manoeuvring time and rudder setting from a range of possibilities - i.e. that option which completes safely, while at the same time minimising the loss of utility (according to criteria advanced by experienced mariners).

Chapter 7: The User Interface, and User Interaction

The benefits of the Expert System described here depend crucially on relevant information being to hand as and when needed, in a form which may be assimilated quickly and unambiguously. In this chapter, the form of representation of information by the system is described. Also, those mechanisms are outlined whereby the user may, in turn, interact with the system to elicit further information or update the knowledge held by the system.

Chapter 8: Simulated Environments for Development and Testing

In the course of development of this Expert System it has been necessary to provide a 'test-bed' or context within which the system may be tested and assessed. Such development is more conveniently conducted within a research laboratory than at sea; also, a simulated test environment precludes the possibility of spurious errors being indicated through technical problems with sensory inputs (though the simulated environment must itself be thoroughly validated to obviate equivalent problems); last, a simulated environment can provide a wide variety of 'experience' for the system - possibly at an accelerated time-scale - without exceptional cost or risk.

Chapter 8 explores the requirements of such an environment, and various formats which have been employed in bringing the system to its present level. The strengths and weaknesses of these formats are considered, together with factors which must be given due credence in transferring to the 'live data' situation.

Chapter 9: System Validation

This chapter describes the appraisal of the system by expert mariners at its current phase of development. It also outlines the results of 1000 simulated encounters, based on randomised initial parameters, illustrating the efficacy of the look-ahead process in resolving all encounter situations so generated. Finally, parallel running of the system alongside a standard ARPA on board a research vessel is described; manual data entry and manual control of the vessel resolve any technical difficulties, ensuring an objective comparison of system performance with accepted good practice in a 'live' test of the system at sea.

Chapter 10: Multi-Ship Encounters

The logistics involved in handling an encounter with just one hazard vessel are relatively clear-cut. Decisions, and manoeuvres, become substantially more complex as the number of hazard vessels increases. It is essential that the expert system is susceptible to expansion in order to accommodate such complexities.

This chapter considers the additional requirements occasioned by a multi-ship encounter situation, and how those requirements may be met within the framework of this expert system.

Chapter 11: Conclusions

As well as offering conclusions on the work to date, chapter 11 attempts to point the way for future developments, some of which are already under way in other research projects. It extrapolates from the current 'state of the art' to draw inferences as to the feasibility of the ultimate goal - a fully-functioning shipboard expert system capable of giving advice (and, if appropriate, control) in any potential collision or near-miss situation, no matter how complex.

Appendix A: Mathematical Considerations in Avoidance Manoeuvring

An overview of mathematical techniques used in the decision logic.

Appendix B: Research Vessels Used For Sea Trials

Information is given on the picket boats used in this practical aspect of the validation process.

Appendix C: List Of Publications

Papers published by the author are listed chronologically.

Appendix D: Listing Of Current System Rules

The set of rules currently in use on the system are listed, together with an explanation of the structure of these rules.

Appendix E: Extended Bibliography

Additional references to papers pertinent to this thesis are listed, with further information on those of particular interest.

CHAPTER 2

FUNDAMENTAL CONCEPTS AND PRINCIPLES

"the whole of science is nothing more than a refinement of everyday thinking"   Albert Einstein (1879 - 1955) 'Out of my later years'.

2.1 Domains and Arenas

As for the hypothetical pedestrians in the imaginary street, there are certain basic tenets which must be observed if a ship is to avoid coinciding in space and time with another vessel, or other hazard. These may be gleaned from a combination of applied commonsense and observed good practice.

One concept which has been derived from such commonsense considerations, and supported by observation of standard practice among mariners, is the 'safety zone' or 'personal space' around a vessel. This zone, termed the 'domain' by Goodwin [8], is that region around a vessel which the master endeavours to keep clear of any potential hazards. If all manoeuvring is conducted so as to preserve such a 'buffer zone', it follows logically that no collision can take place.

Goodwin hypothesised the existence of such a domain, reasoning that such a region is an essential aspect of any effective avoidance strategy. Observation supported this premise, and further indicated differential levels of concern broadly corresponding to the three sectors around a ship subtended by its navigation lights: port (red) sector, starboard (green) sector and rear/overtaking (white) sector. Figure 2.1(a) shows the relative proportions of a typical 'Goodwin' domain, indicating:

(a) requirement for more sea room to starboard than to port, starboard being the conventional direction for avoidance manoeuvres under most circumstances;

(b) a relatively low degree of concern with hazards to the rear of own-ship, reflecting the relatively low approach speed of an overtaking vessel compared with a vessel approaching head-on or in a crossing encounter.

Clearly, the overall size of the domain applied by a ship's master may vary with size and manoeuvrability of vessel, nature of cargo and prevailing circumstances, and personal preference of that master; however, the relative sizes of the three sectors appear to fairly consistently follow the pattern shown in fig. 2.1(a).

Although the Goodwin model of the domain fits the observed practice, the discontinuities at sector boundaries have been viewed by some as inconsistent and illogical (as well as being difficult to accomodate in manoeuvring algorithms). Davis [9] propounded a variant of the Goodwin domain, which maintained continuity of the circular domain boundary whilst preserving the relative areas of the three sectors as observed by Goodwin. An example of such a domain, with own-ship in an off-centre position, is shown in figure 2.1(b).

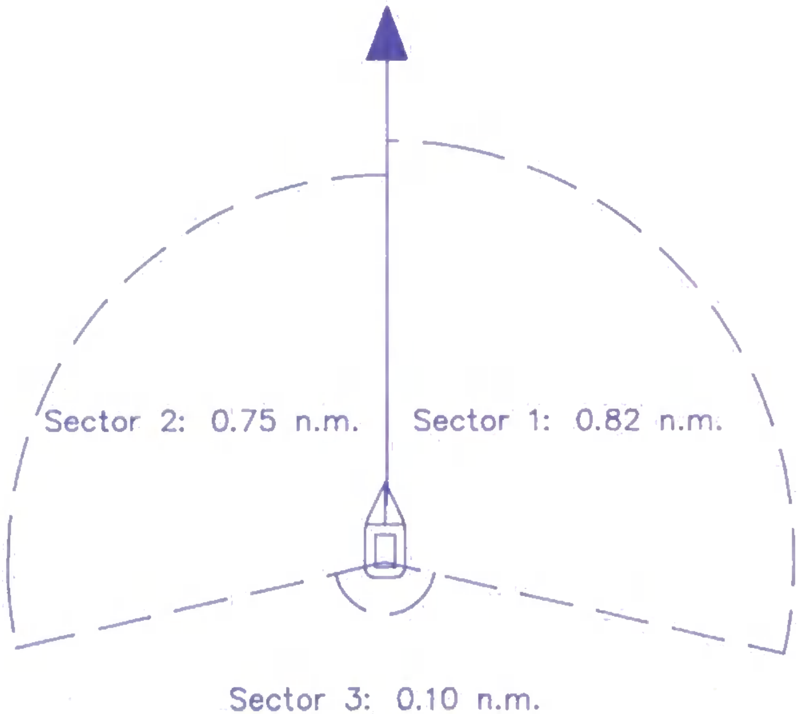

Figure 2.1(a) Example of Goodwin Domain

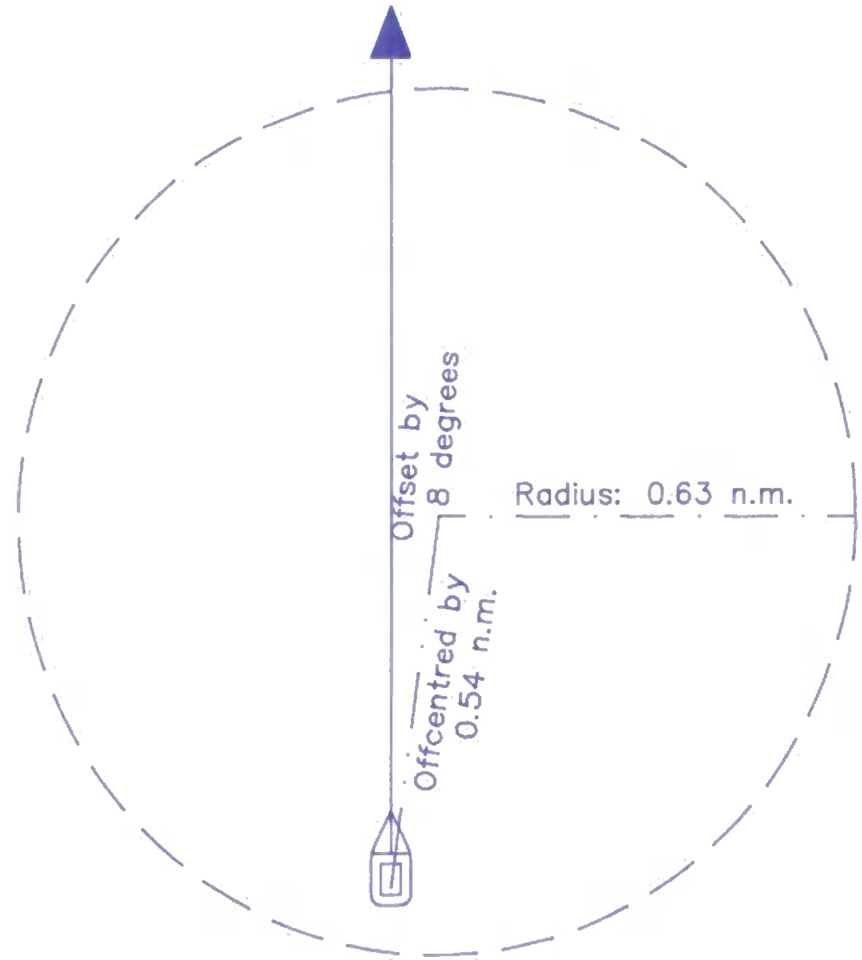

Figure 2.1(b) Example of Davis Domain

Figure 2.1

Such a circle is easier to handle in decision logic than the Goodwin domain, but it still has an awkward tendency to function as an eccentric cam as the vessel turns, giving rise to highly complex mathematical considerations. Furthermore, it is a moot point whether conservation of relative areas is of major significance in comparison to, say, offset of domain boundary (the latter being the more crucial factor in determining whether a particular hazard enters the domain). In the event, it has been decided to sidestep such subjective considerations and mathematical pitfalls, and use a circular domain with own-ship at centre. Such a simplified form has proved adequate for all practical purposes, and arguably provides as solid a basis for decision-making as any more complicated interpretation of observed practice; certainly, choice of a suitable domain radius can ensure safe manoeuvring without leading to over-reaction. Before leaving this subject, it should be noted that a structural concept based on the Davis domain is considered in chapter 10 as a potential discriminator for first-stage resolution of multi-ship encounters.

If a domain of any form is to be implemented, it follows logically that a broader event space must be defined. That space, first identified and termed the 'arena' by Davis [10], represents the area around own-ship within which potential hazards are identified as such, and avoiding action planned and implemented so as to obviate the possibility of domain infringement. The Davis arena is a circle centred on own-ship, of substantially larger radius than the domain - again, size of arena would depend on personal preference of the ship's master and prevailing circumstances.

This construct, though a valuable first step, has limited applicability since it draws no distinction between the nautical equivalents of the skateboarder and the old lady with the walking frame. Curtis et al [11] carried the arena concept into the realm of time-based decision-making by advancing the proposal that the crucial factor is not distance, but time to domain infringement. The 'closing-time' version of the arena was thus formulated, and termed the RDRR (Range-to-Domain/Range-Rate). In this, Curtis effectively drew a time-based 'circle' around own-ship at centre, and indicated a need for action if any hazard fell within that area. For example, a 10-minute RDRR would identify any potential hazard whose relative velocity put it within 10 minutes of infringeing own-ship's domain (assuming constant course and speed for both vessels). This criterion has proved a useful starting point in identifying both the need and the appropriate time for automated decision-making in the context of an expert system. Chapter 6 considers the results obtained when such a discriminator was used in the first prototype expert system.

Konyn [12] further advanced the RDRR concept by considering the combination of 'thinking time' and 'action time', and its variations depending on relative positions of the vessels involved. However, this RDRR+ has still proved too blunt an instrument for the almost infinite variability of juxtapositions and relative velocities of two vessels in contention. It has been rapidly superceded by the perceived need for a decision time tailored to the specific circumstances of a particular encounter. In the final analysis, concern with respect to an observed potential hazard is not related to an area, or a sector, or even a radial arc; it relates simply and solely to the one-dimensional line joining the two vessels - the length of that line, how fast that length is reducing, and what the minimum length of that line will be on current projections. From these considerations stem further concerns as to the likely consequence of a given manoeuvre at a given time. Hence the development of the look-ahead simulation technique described in chapter 6, which has given rise to the PSMT (Predetermined Safe Manoeuvring Time) - since refined to give an Optimal Manoeuvring Time (POMT), based on a variety of criteria.

This 'customised' manoeuvring time now forms the basis of decision timing and strategy in the current version of the expert system. Needless to say, advance identification of such a time for action necessitates a yet broader region, the 'ballpark' - consideration of this region is best left to chapter 6.

2.2 The Collision Avoidance Regulations

The International Regulations for Preventing Collisions at Sea [3] provide a substantive framework on which to hang a rule-based system aimed at improving maritime safety. Furthermore, all other responsible seafarers will be observing these regulations, and anticipating such observation on the part of any vessel they may meet, so consistent obedience to these guidelines is strongly indicated for mutual understanding and co-operation. Not least, these regulations are backed by the force of law in most parts of the world, and divergence from them may be punishable in the courts whether or not this leads to a collision or other untoward event. In short, any expert system for marine collision avoidance must incorporate the IMO regulations as a core element of its rule base.

Having said that, a corollary must be added that these regulations are by no means definitive in their specification of action to be taken in a given circumstance. Various terms used in the text are open to widely-differing subjective interpretations by ship's masters. In the first instance, vessels in contention are required to take avoiding action 'in good time'; such 'good time' may be reckoned as anything from, say, 10 to 20 minutes before closest point of approach (CPA). Since the regulations also require the ship's master to identify the nature of the encounter at that time and initiate action accordingly, and since the perceived nature of an encounter may alter significantly (without any change of course) over a period of 5-10 minutes, there is scope for the masters of the two vessels to view the same encounter in two totally different ways, and apply different decision criteria. This problem is considered in more detail in chapter 5, and illustrated in figure 5.3. Obviously the regulations provide an 'escape clause' where two masters find themselves operating at cross-purposes, but two aims of an expert system would be to be (a) self-consistent, and (b) as tolerant and flexible as possible in the face of a differing interpretation.

Another area of uncertainty is that of what constitutes a 'clear turn'. Since visual (and/or radar) contact is often the only form of communication between two vessels, mariners are required to make their intentions obvious by effecting a 'clear turn' before straightening up.

This is variously interpreted to mean anything from 10-15 degrees up to possibly 45 degrees. Again, such personal vagaries need to be tied down to a greater consistency for a computer-based system - this does not seem an appropriate application for a random number generator! One must, of course, be open to the possible adjustment of this criterion in changing circumstances; for example, a 45-degree turn may be the order of the day in open-sea manoeuvres, whereas 15 degrees may be adequate to make the point in the confines of the Dover Strait.

The regulations are equally general over the required separation distance between vessels. Masters are required to keep a 'safe distance' between their vessel and another, and (by implication) required to give a value judgement on what constitutes a 'safe distance'. Again, such a value judgement must be allowed flexibility for variation in changing circumstances: a substantial 'safe distance' in the open sea may, of necessity, be reduced to a rather lesser (but still safe) separation in a busy seaway. The 'safe distance' alluded to is, of course, represented by the domain radius, which is a system parameter chosen by the ship's master. Since domain radius is a matter of choice, an encounter may occur in which the two vessels manoeuvring are adequately separated for one master's peace of mind, but too close for comfort from the other's point of view. The expert system must be cognisant of possible consequences of such disparity (e.g. hazard vessel taking emergency evasive action when own-ship is manoeuvring satisfactorily, to all appearances), and sufficiently broad in application to avoid causing such reactions if at all possible.

One final, but rather major, area of subjective interpretation of the regulations is on the matter of course changes, as contrasted with collision avoidance manoeuvres. If a vessel undertakes a course change for navigational purposes, rather than a collision avoidance manoeuvre for safety reasons, then the master of that vessel is not bound by the collision avoidance regulations. This apparently clear distinction becomes rather blurred when a course change is, in essence, undertaken to avoid a situation which would necessitate collision avoidance manoeuvring. Figure 2.2 illustrates an example of typical action for a ferry crossing a busy shipping lane: an early course change is taken in preference to either a series of 'slalom' avoidance manoeuvres or one substantial avoidance manoeuvre which would take own-ship further off track. Where such a course change is undertaken well before any possible domain infringement, the master of the vessel is unquestionably entitled to complete freedom of action. However, as a likely domain infringement becomes increasingly imminent, constraints imposed by the collision avoidance regulations must be regarded as increasingly binding.

Different mariners (or even the same mariner under differing circumstances) may regard themselves as bound by the regulations once they are within 20 minutes, 15 minutes, or even 10 minutes of a domain infringement situation. It is considered expedient by some mariners to 'bend' the rules to imply that, even at very close quarters, the collision avoidance regulations need not be invoked unless a direct hit is imminent - a 'near-miss' threat may be avoided by an unrestricted choice of course change. This strategy explains a number of potential

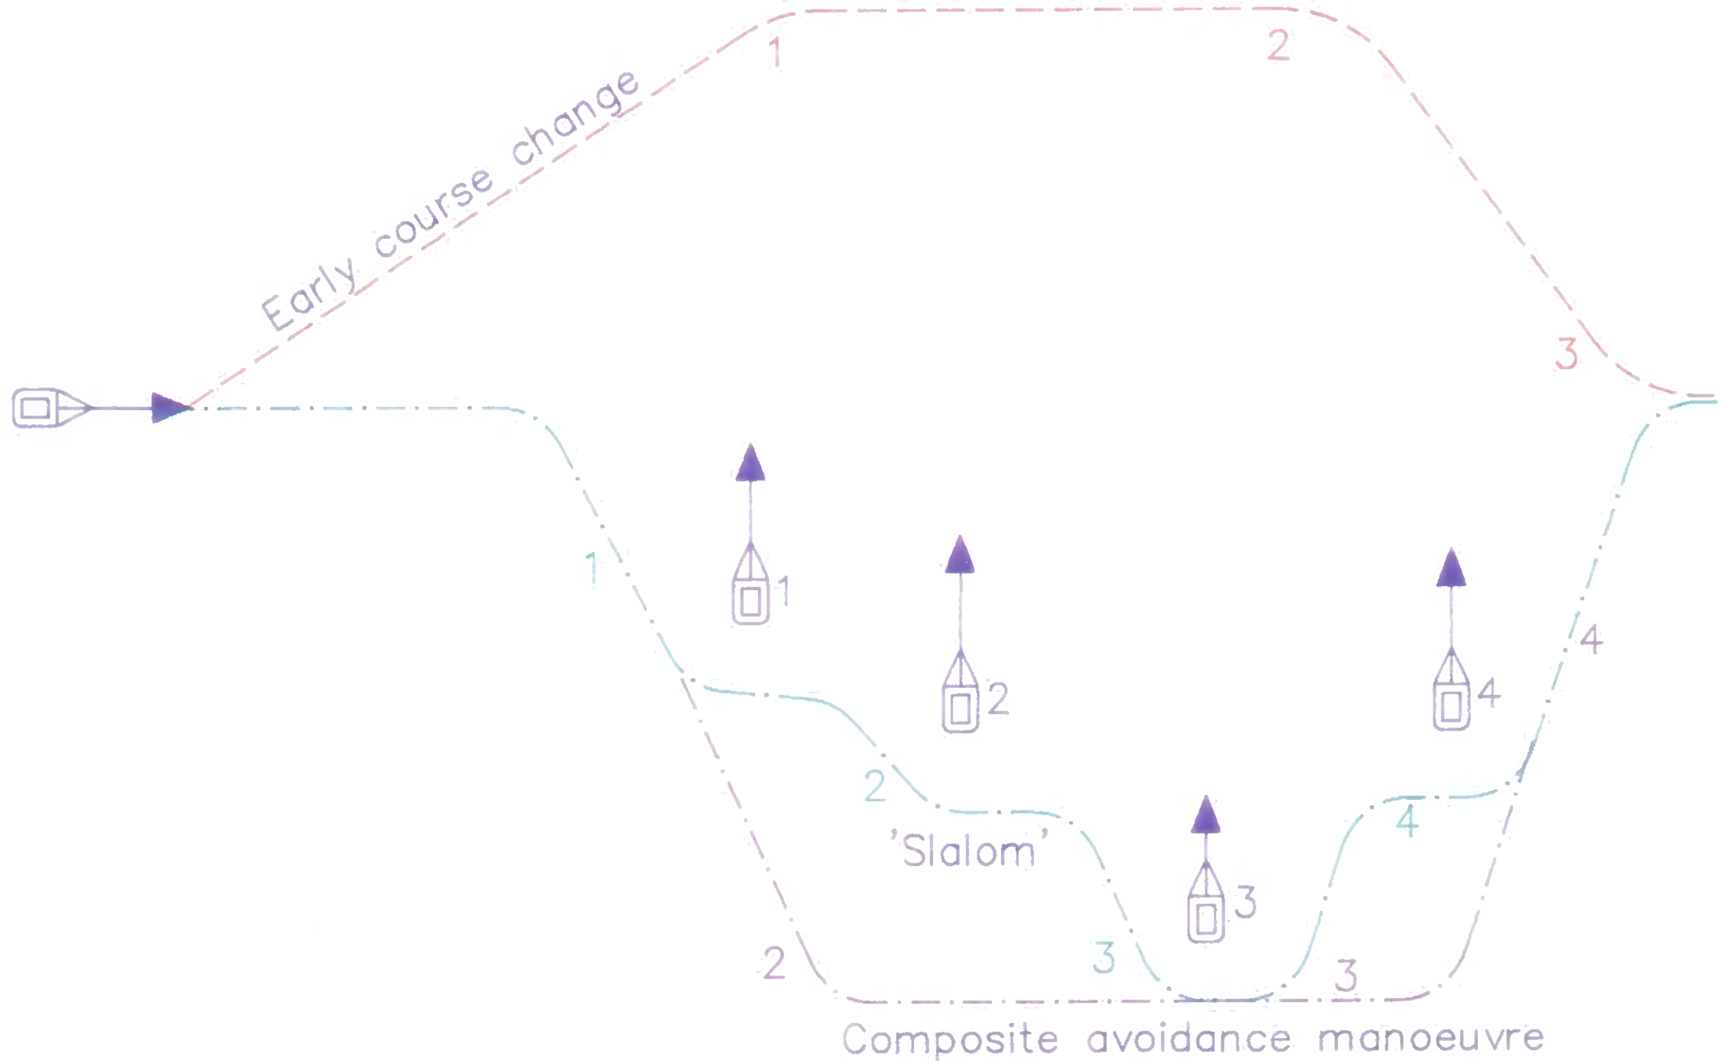

Figure 2.2 Alternative routeings for avoiding a number of other vessels (positions of target vessels are shown for times of corresponding marked positions on own-ship's tracks).

'near-misses' that have been converted into disastrous collisions by one master obeying the regulations whilst the other 'used his initiative' - the classic 'radar-assisted collision'. Figure 2.3 illustrates the mechanics of such an incident with specific reference to one of the great maritime disasters of modern times, the collision between the 'Andrea Doria' and the 'Stockholm' in 1956 [13].

2.3 Ship Manoeuvring Models

In order to ascertain the likely outcome of a particular manoeuvre or strategy, it is necessary to have a working knowledge of the performance characteristics of the vessels involved. In particular, any required outcome identified by the expert system must be recommended to the master of own-ship in terms of actions (rudder setting, engine speed) and their timings calculated to produce that outcome. This presupposes that the expert system holds within its repertoire a hydrodynamic mathematical model of own-ship's behaviour under varying control inputs.

A body moving in three dimensions (such as a ship at sea) may experience any or all of six aspects to its motion - six degrees of freedom: linear, with components along each of three orthogonal axes (longitudinal, transverse and vertical); and the corresponding rotational motions about each of those axes. These are referred to in nautical terminology as: surge, sway and heave; and roll, pitch and yaw, respectively. They are illustrated in relation to a ship's orientation in figure 2.4.

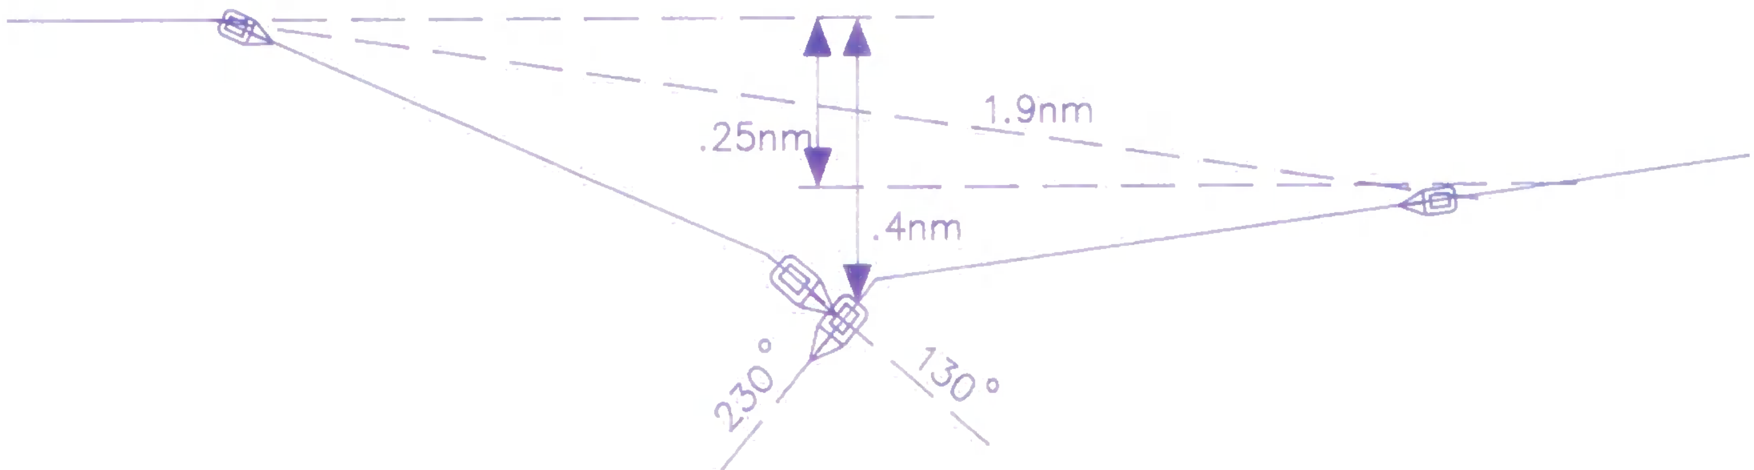

Figure 2.3 Schematic of the collision between the 'Andrea Doria' and the 'Stockholm' (see ref. [13]).

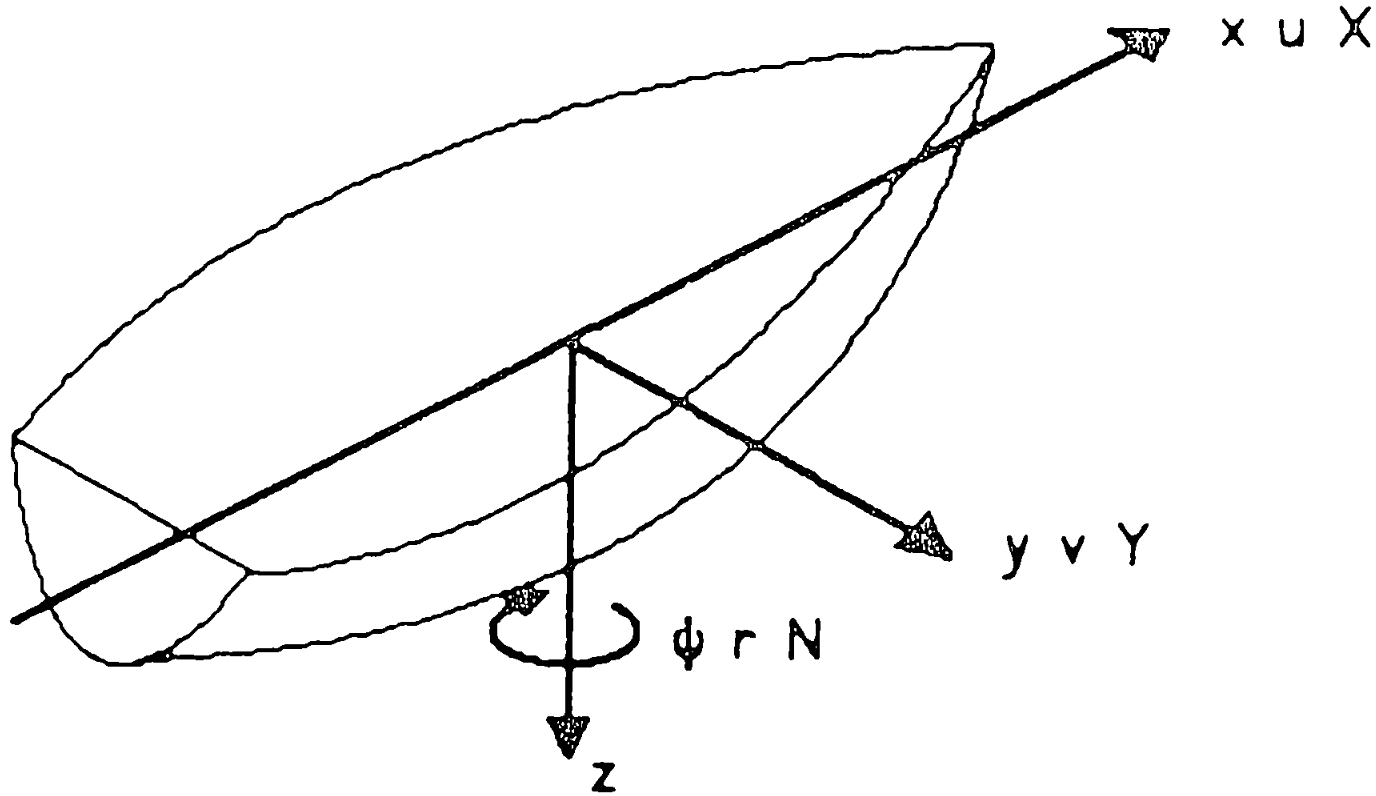

Figure 2.4 Ship axes co-ordinate system, showing displacements, velocities, rotation, forces and moment having effects in the horizontal plane.

Of these, only three relate directly to motion in the horizontal plane: surge, sway and yaw. It would seem apposite, therefore, to model the progress of a ship across a body of water by reference to these three forms of motion - a '3 degrees of freedom' (3DF) model - and many ship manoeuvring models are formulated as a set of three equations representing these three forms of motion. However, recent research [14] has shown possible benefit in a 4DF model, since roll has been shown to have a material effect on motion in the x-y plane.

A linear ship model may be expressed as a function of first-order terms only of the relevant variables. A non-linear model requires higher-order terms, and in return gives a more accurate representation of ship behaviour. Two major difficulties arise in relation to the non-linear model: evaluation of the relevant functions at each step is a more time-consuming matter; and (more significantly) the derivation of the coefficients for the terms in each equation is a substantial undertaking involving major sea trials. By contrast, linear coefficients may be calculated directly from standard technical data on any vessel.

Chudley [15] describes the benefits of a non-linear model (as formulated by Tapp [16]) over a linear model; he also identifies the concomitant difficulties of acquiring the necessary hydrodynamic coefficients for the non-linear model, as contrasted with the lesser task of deriving the coefficients for the linear model. However, his analysis of the 4DF linear model [14] indicates accuracy approaching that of the non-linear model without the accompanying complications - a useful compromise where ease of derivation and speed of processing are major considerations.

Chudley also outlines the benefits of the modular model, in which effects of propeller(s), rudder, wind, tide etc. are identified by separate terms in the ship model (see figure 2.5). This contrasts with the earlier 'holistic' model, in which all such effects were compounded into a single term. Clearly, the modular form facilitates modification of one aspect of the model, or addition/deletion of one of the component effects, depending on perceived significance.

Derivation of a sophisticated manoeuvring model of own-ship is clearly a more practical proposition than formulation of such a model for each hazard vessel encountered. This is more a problem in theory than in practice, since limited ability to second-guess the purposes of the master of a hazard vessel render of limited value any capacity to model that master's hypothetical intentions. The main requirement in modelling hazard motion is straight-line prediction at constant speed - for which no hydrodynamic model is required. It is possible that future enhancements to the rule base may draw from a 'bank' of standard ship models to identify potential manoeuvring capability of an observed hazard; this concept is explored further in sections 4.4 and 11.5.

Within the system as currently implemented, own-ship motion is modelled by straight-line sections, plus circular arcs for turning motion. This is clearly an over-simplification of the real world. However, the algorithms used for manoeuvring calculations may be simply modified to accomodate the manoeuvring models described earlier in this section - or in some cases, acceptable simplifications of those models to enable multiple look-ahead predictions in real-time, as described in chapter 6.

Figure 2.5 Equations for modular manoeuvring model

Ignoring heave, roll and pitch, which do not act in the x-y plane:

Force (surge): $X = m(\dot{u} - vr - r^2 x_G)$

Force (sway): $Y = m(\dot{v} + ur + \dot{r} x_G)$

Moment (yaw): $N = I_z r + m x_G (v + ur)$

By selecting origin of ship co-ordinate system to coincide with centre of mass of the vessel, these reduce to:

$$X = m(\dot{u} - vr - r^2 x_G)$$

$$Y = m(\dot{v} + ur + \dot{r} x_G)$$

$$N = I_z r + m x_G (v + ur)$$

involving only mass, velocity and acceleration in x and y directions, yaw rate and moment of inertia about z-axis.

X, Y and N may each be expressed as a sum of modular components, dependent on effects of hull, propeller, rudder and disturbance factors:

$$X = X_H + X_P + X_R + X_D$$

$$Y = Y_H + Y_P + Y_R + Y_D$$

$$N = N_H + N_P + N_R + N_D$$

Each such component factor may be adjusted independently, as considered approriate, or in certain cases omitted if regarded as negligible.

Specific characteristics of any ship's turn which must be given due credence in the mathematical model are:

(a) a time-lag between putting the rudder over and the ship commencing to turn; during this time lag, the ship may initially 'twitch' fractionally in the opposite direction from the turn, but for all practical purposes motion may be regarded as continuing on the original course;

(b) a turning motion which spirals gradually inward (Dieudonne spiral, [17]), rather than tracing a circular arc.

(a) may be accommodated within the existing framework by calculating the time delay from the mathematical model and simply adjusting the time for putting the rudder over accordingly. In turns of less than, say, 90 degrees, the spiral arc in (b) may be approximated for predictive purposes by a circular arc without significant loss of resolution; since all 'normal' manoeuvres fall into this category, the mathematical model may be applied to provide the appropriate approximation, which may then be used in the existing look-ahead and rule structure.

It is planned that the final system should utilise a modular manoeuvring model, either the non-linear version or the enhanced 4DF linear version. Hydrodynamic coefficients required to model a particular own-ship will be carried as part of the knowledge base, as described in chapter 4.

2.4 The Human Interface

In any system which is concerned with providing information intended to contribute to overall safety, communication of that information must be a major consideration. Where response to that information must be made in real time - possibly a very limited time, under conditions of stress - the medium of communication assumes a yet higher significance. If the relevant facts are not presented in a form susceptible to clear, fast and unambiguous assimilation, then the system itself is valueless. More than that, a system which distracts without informing as required is an added hindrance to safety and efficiency.

Anyone who has either created or (more significantly) used a software package intended for the naive user will know that screens full of text deter rather than aid. On the premise that a picture is worth a thousand words, it is almost axiomatic to state that any event/situation which may be meaningfully represented in graphical form, should be. Textual or numerical supporting information should be as clear and concise as possible; use of colour to highlight aspects of this information can also be most helpful, if judiciously applied. (N.B. These are not highly technical observations; any child who has played a good - or bad - video game will be conversant with these fundamental principles).

Manipulation of the display, using simple controls which do not allow the user to get 'lost' or confused, are a further essential feature. The capacity to elicit further information as needed, and to close off potentially distracting elements when not required, rate highly on the system specification. Facilities for scrolling text and graphics displays, and 'zooming in' on graphics detail (i.e. providing localised enlargement) are likewise of substantial potential benefit. These requirements are indicative of a 'windowing' environment (i.e. sections of screen partitioned off to provide a 'window' onto text or graphics items, generally with facilities for scrolling both vertically and horizontally); icons (representative symbols) prompting for action, menu options displayed and actioned at the press of a 'mouse' (user interface device) button, and a screen 'pointer' indicating intended area of activity or menu option, complete the specification for an advanced WIMPS-based user interface.

With respect to the rule structure governing the response of the system to prevailing circumstances, it would not seem appropriate to allow the aforementioned naive user to update that structure between (or even less during) voyages. Extension of the rule base is envisaged, to accommodate new experience in terms of encounter situations and responses, but this should only be done by one conversant with the functioning of these rules. It is also essential that such updating be done 'off-line', i.e. not on active service, and in circumstances which permit thorough testing of the updated rules before they are used in action. An updating facility is to be added at some future date, to permit such enhancement of the rule base by suitably-qualified personnel in a 'user-friendly' fashion. The nature of such a facility is considered in detail in section 11.4; other aspects of the above user interface form the substance of chapter 7.

CHAPTER 3

THE EXPERT SYSTEM 'SHELL'

"Cogito, ergo sum."

Descartes (1596 - 1650) Le Discours de la methode

3.1 The Inference Engine

Figure 3.1 illustrates the conceptual structure of the expert system. As will be seen, a central feature of this structure is the inference engine, a 'content-free' set of logical connectives which drives the rule base. The inference engine operates on two data structures, the rule structure and the knowledge base, and interacts with the input-output routines which form the user interface.

There are a number of expert system shells available commercially, each consisting essentially of an inference engine interacting with a rule structure accommodating user-defined rules and a knowledge base also built up by the user; appropriate input/output facilities are also essential. A number of such shells have been studied in detail, but none has been considered suitable for the task in hand.
As with programming languages, expert system shells are tailored to the class of application for which they are designed. However, there is a common thread running through most of the expert systems currently in use. Various characteristics distinguish this task from more conventional applications of expert systems.

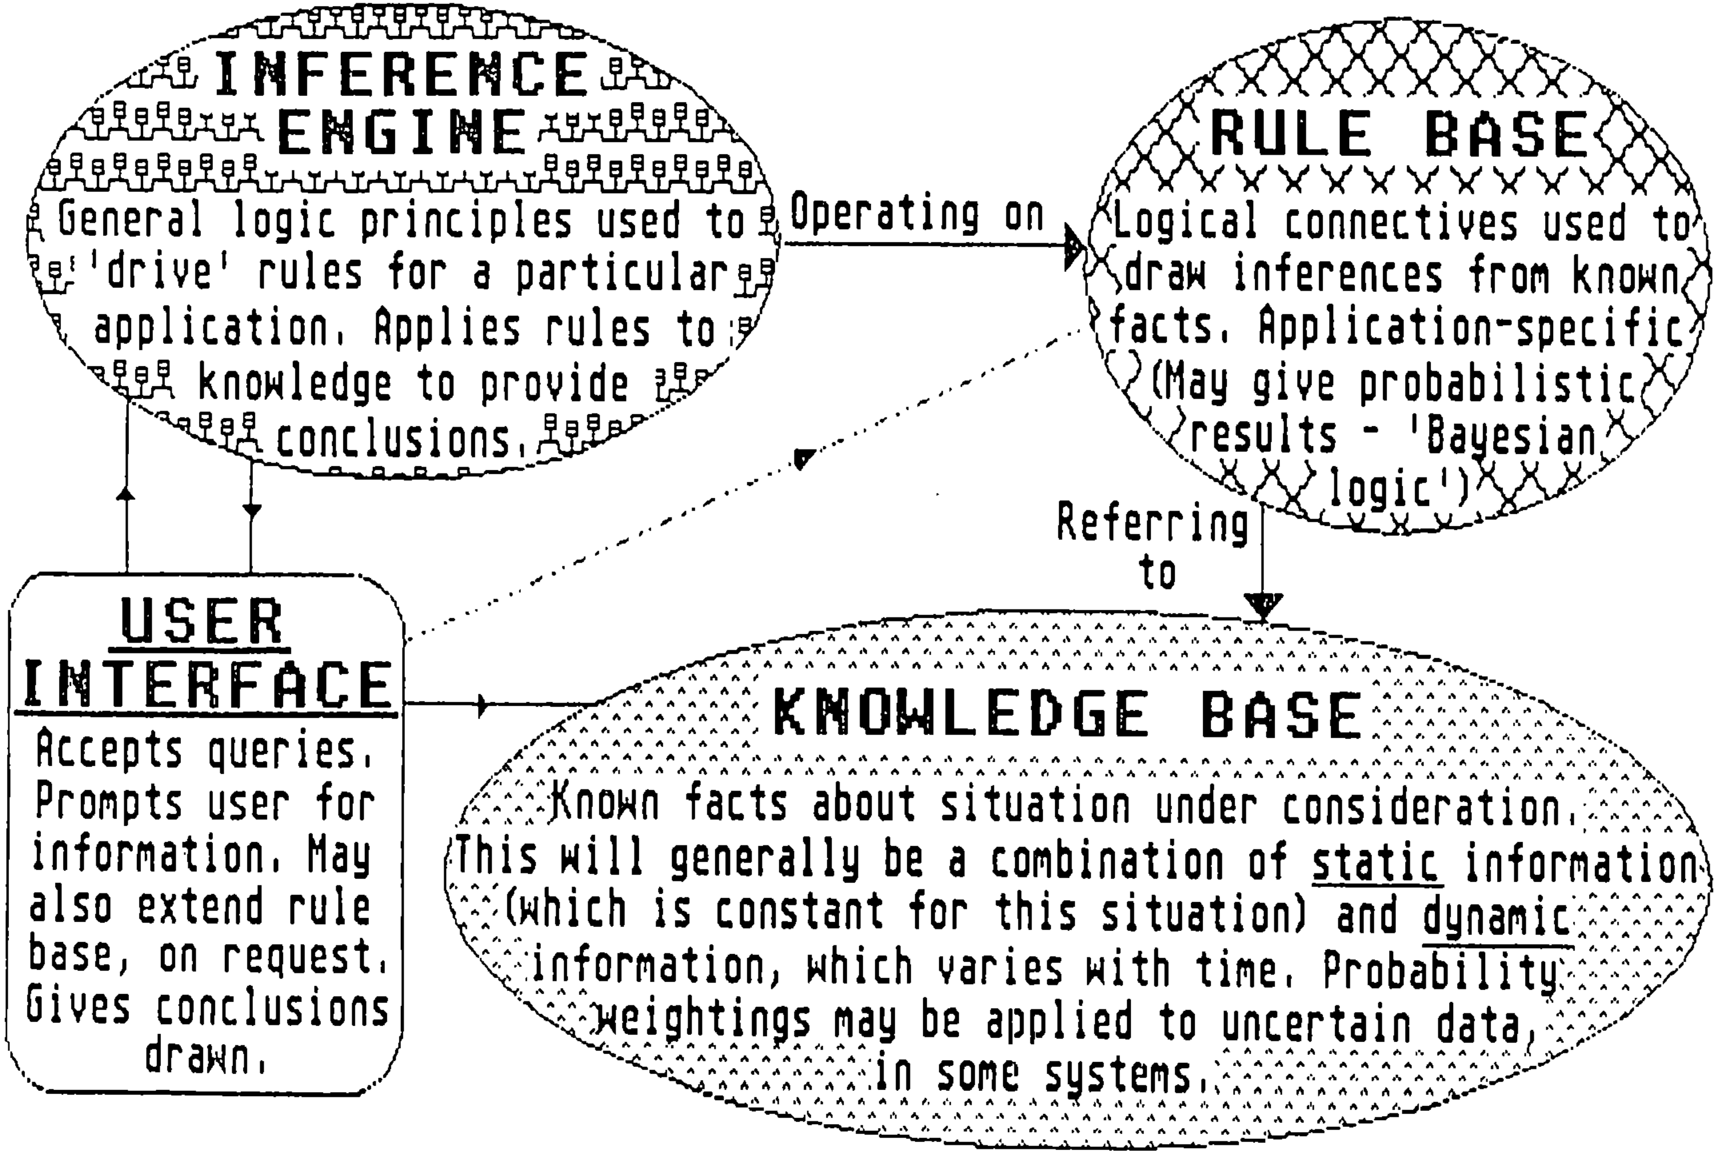

Figure 3.1 Outline structure of expert system shell

(1) Responses from the system are required in real time. In a majority of expert systems studied, there is no serious time constraint on consideration of available data, application of the rules, and presentation of the consequent conclusions and/or advice. Where decisions are being made regarding avoidance of imminent disaster ('imminent' in this context meaning from half an hour down to a few minutes - possibly less in emergencies), the luxury of unlimited thinking time is not an option. There is little merit in a system which gives comprehensive advice on how to avoid a mishap - after that mishap has occurred. It is also significant to note that the 'ground' is moving under one's feet (literally), so that prevailing conditions once the evaluation is made are dependent on the time taken to make that evaluation; such an ever-shifting scenario is not a regular feature of expert system applications, but must be taken into account in the decision-making logic of this system.

(2) The knowledge base is continually being updated via sensory input channels, to reflect the changing nature of the environment to be evaluated by the system. Sensors attached to own-ship's engine(s) and rudder, as well as radar and navigation systems, will provide 'dynamic' reference information to supplement the 'static' information (ship's length, draught, manoeuvring characteristics) held in the knowledge base. This aspect of the system carries with it an implication of knowledge as to the scope and limitations of such sensory input, including error margins; consideration of these factors, and their possible consequences, is integral to the work of the inference engine, acting in tandem with the rule base.

(3) The inferential logic for this system bears very heavily upon evaluation of multivariate trigonometrical expressions. Resolution of potential conflict situations consists primarily of vector calculations, and intersections of vectors with circular regions or arcs. Displacement vectors and velocity vectors (absolute and relative) must be considered in relation to domain boundaries and turning arcs; intersections of domains with turning arcs, and exact times when these occur, are also crucial. In short, evaluation of any rule almost always involves a substantial amount of processor-intensive mathematical calculation (n.b. trigonometrical calculations are heavy on processor time, and look-up tables for trigonometrical functions would be totally inadequate in this context). As yet, floating-point hardware has not been implemented in this system; the speed-up factor (probably around 10) given by such hardware is being held 'in reserve' to meet increasing time constraints as more complex situations are dealt with.

Given these three major departures from 'standard' expert systems technology, a 'standard' approach has been found wanting on various fronts. Specifically, the 'off-the-shelf' expert systems shells currently available appear over-simplistic in their reasoning, unable to handle the type and depth of mathematics needed here, and not in any way geared to a time-constrained, time-dependent decision process. The question of sensory input is probably not insuperable, but it is not a standard feature of any of these packages; further investigation into its feasibility seems pointless in the circumstances.

On the broader front of programming languages conventionally used for expert systems applications, the above constraints apply equally. Declarative languages such as Prolog and Lisp are not geared towards the type of application described here (though some incorporate 'add-on' mathematics libraries which go some way to meeting that particular problem). Likewise, object-oriented languages are implemented in ways which militate against the most efficient functioning of such a package.

It is apparent that an air of mystique has arisen around the whole field of artificial intelligence (AI), fostered by the 'real AI' school. Unfortunately, this school of thought appears to hold the view that an 'AI application' must be implemented in an 'AI language' and using a specific 'AI approach'. This philosophy has been found to be seriously limiting in relation to non-standard applications of AI such as the expert system under consideration. The notable mathematician Penrose explodes this AI myth [18], and leaves the way free for non-conformists to exploit fully <u>any</u> language(s) and techniques which might prove useful. More recently, some highly successful AI applications have been coded, partly or wholly, in 'C' [19], a procedural language which combines the best of high-level structuring with the potential to utilise the full extent of the host computer's capabilities.

Section 1.3.2 considered the software and hardware requirements of this system. The possibility of a dual-language package was evaluated and rejected. As indicated, the inference engine for the collision avoidance system is written completely in 'C', and comprises a number of largely independent modules as described in detail below; these modules incorporate the user interface facilities and support the data structures which, taken all together, form the expert system shell.

3.2 The Windowing Module

Any WIMPS application requires a very substantial amount of setup code, plus maintenance routines for any windows (actual or potential) and features of those windows. On the Atari computer, each window (display area) must be defined in detail, and every window function (sliders for horizontal/vertical scrolling, 'close' icon to delete window, etc.) coded explicitly for each screen window. Any window operation (mouse button clicked on slider, etc.) must then be 'trapped' by the program code and appropriate action taken.

On the Archimedes, a Window Template Editor is provided as a software tool, to enable the programmer to create and/or modify any window to be used in a particular application. Features such as sliders, icons, pop-up menus may be added, and the exact mode of operation of each of these features defined. These window definitions are held in a file, to be called as required from within the application program. This file may be regarded as a user library of windows, replacing the code that would otherwise be necessary to define each of these windows in the main program.

User interaction on the Archimedes is 'event-driven'; any action by the user - mouse moved, mouse button pressed, keyboard action - is termed an 'event'. Events are queued as they occur, triggering CPU interrupts on a regular cycle (or on a priority basis). Each window must be supported by a set of event-handling routines, to which control is passed if an event happens within that window. Those routines will then identify the type of the event, and handle the situation accordingly - display menu, close window, scroll display or whatever. RISC-OS (the operating system for the Archimedes) has the capacity to handle certain events in a standard manner, without the need for specialist user routines; alternatively, the programmer may opt to deal with such events in a way tailored to the requirements of the program. The former would normally be the case, for example, in scrolling a standard text window; mouse operation on the sliders would automatically invoke the relevant RISC-OS 'standard scroll' routine. By contrast, scrolling of a window incorporating a scaled graphics display will require specialist user routines to relocate the origin of that display, apply the appropriate scaling factor, and ensure that any future updates of that display are directed to the correct screen area. This option, between automatic and user-defined windowing functions, may be selected from within the window template editor.

The windowing module holds the initialisation code for the window displays used by the system. This code includes:

(a) declaration of event-handling routines (including error-trapping procedures) for the various windows, and also for 'non-window-dependent' events such as messages from other applications and time-initiated events;

(b) initialisation routines for the system, actioned on:

   (i) clicking on the application icon in the relevant directory, to establish the application as active and place the application icon on the 'icon bar' (bottom of screen display) for subsequent action;

   (ii) clicking on the application icon on the icon bar, to start the 'run' phase of the process - windows opened, input dialog box(es) activated, time-sequence and display begun; this routine also broadcasts a message triggering any other related applications currently on 'stand-by' within the system (typically, simulation tasks, as at present, or datalogging applications such as radar, satellite navigation input, etc).

Naturally, this module contains numerous cross-references to other program modules, since much of the event-handling (apart from straight window manipulation) relates to rule evaluation and/or navigational considerations - see below. The logical interaction of this module with others in the system is illustrated in figure 3.2.

3.3 The Navigation Module

At regular intervals (every 20 seconds, real time), a time-initiated interrupt event triggers a sequence of operations. These include appraisal of own-ship's current position, speed and course, and corresponding data for any hazards under consideration in the vicinity. At the present stage of development, hazards in the vicinity of own-ship

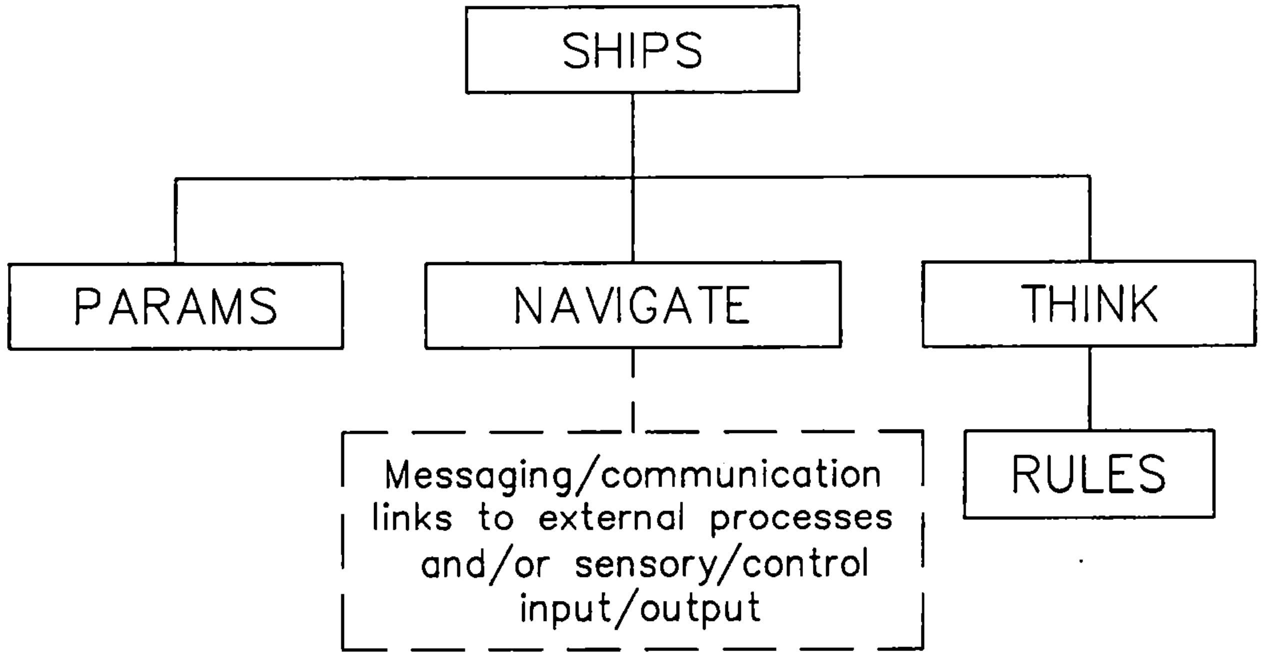

Figure 3.2 Module dependency chart for expert system, showing primary logical connections between program modules. Interaction between these modules is also effected by event-handling interrupt routines.

are simulated by separate tasks running concurrently under the multitasking RISC-OS (see chapter 8 for details). Upon each timer activation, a broadcast message triggers a response from each of those 'hazard' tasks, providing the pertinent navigational data. A similar module simulates the action of own-ship (under the guidance of the expert system), and provides the expert system with navigational data on own-ship in response to the aforementioned broadcast message. A recent enhancement of this process is detailed in an attached paper [6]. The final 'live' version of the system will, of course, receive such data from navigational and sensory equipment such as radar, GPS (Global Positioning System - a satellite navigation facility), sensors on engines and rudder, etc.

This navigational information is used to update the graphics window display of ships' tracks, and to calculate relative headings and ranges of hazard vessels for (apart from other considerations) display in the 'status' window. Such data is further used for evaluation of the current situation in the 'rules' module - see below.

The graphics window requires substantial maintenance by this module, for a number of reasons:

(1) it permits two different scaling factors to be used in translating 'actual' to 'graphic' position, giving rise to a need for a scale-dependent position transformation algorithm;

(2) the ability to scroll vertically and horizontally within the window exacerbates the problems associated with the above algorithm (given that OS-supervised 'automatic' scrolling is not an option under the circumstances);

(3) 'exclusive-or' plotting is used to toggle a 1-mile square grid superimposed on the display (scaled in accordance with current scale of the display);

(4) a colour translation table is used to select appropriate and clear colours for tracks of ships; this table also ensures that the toggling of the grid mentioned above will under no circumstances cause discontinuities in those tracks.

Details as to the nature of the screen display are given in chapter 7.

3.4 The Rules File and Rule Evaluation Module

This combination effectively forms the heart of the inferential logic for this expert system. The former defines, and the latter processes, the conditions and considerations which govern recommendations (or possibly control functions) for safe manoeuvring through a variety of encounter situations. Neither of these elements can reasonably be considered in isolation, since the two are mutually dependent in terms of structure and processing requirement.

3.4.A The Rules File

The rule structure is hierarchical in nature, being based upon the form of a binary tree. Each 'rule' constitutes a node in that tree, with one (or more) link(s) into it from some other node(s) - and/or possibly that selfsame node; two links ('left' and 'right') connect each node onward to yet another node in each case (or again, in either case, possibly the same node). In brief: each node must have left and right links, to two nodes within the structure; each node must be the destination of <u>at least</u> one of those links from some other node (with the possible exception of the root node, active at the start of a voyage). The rule structure is also a 'state table', each node representing one possible stage in a specific type of encounter between own-ship and a hazard vessel. At any point in the voyage, a 'state indicator' points to the currently-active node:- all-clear, overtaking from port stage B, etc. An example of a section of a typical rule structure is shown in figure 3.3.

Each rule incorporates a boolean decision function - most of these are based on trigonometrical considerations regarding own-ship and hazard vessel - which forms the basis for left or right branching to the next rule in any particular situation. In some cases a subsequent rule is evaluated/actioned immediately, in others such action is deferred until the following time-step (shown as solid and dotted links respectively in figure 3.3). As may be seen from the diagram, most encounter states incorporate a test for whether the current condition still pertains (E.g. still head-on stage 2?) - if so, the rule structure cycles on that

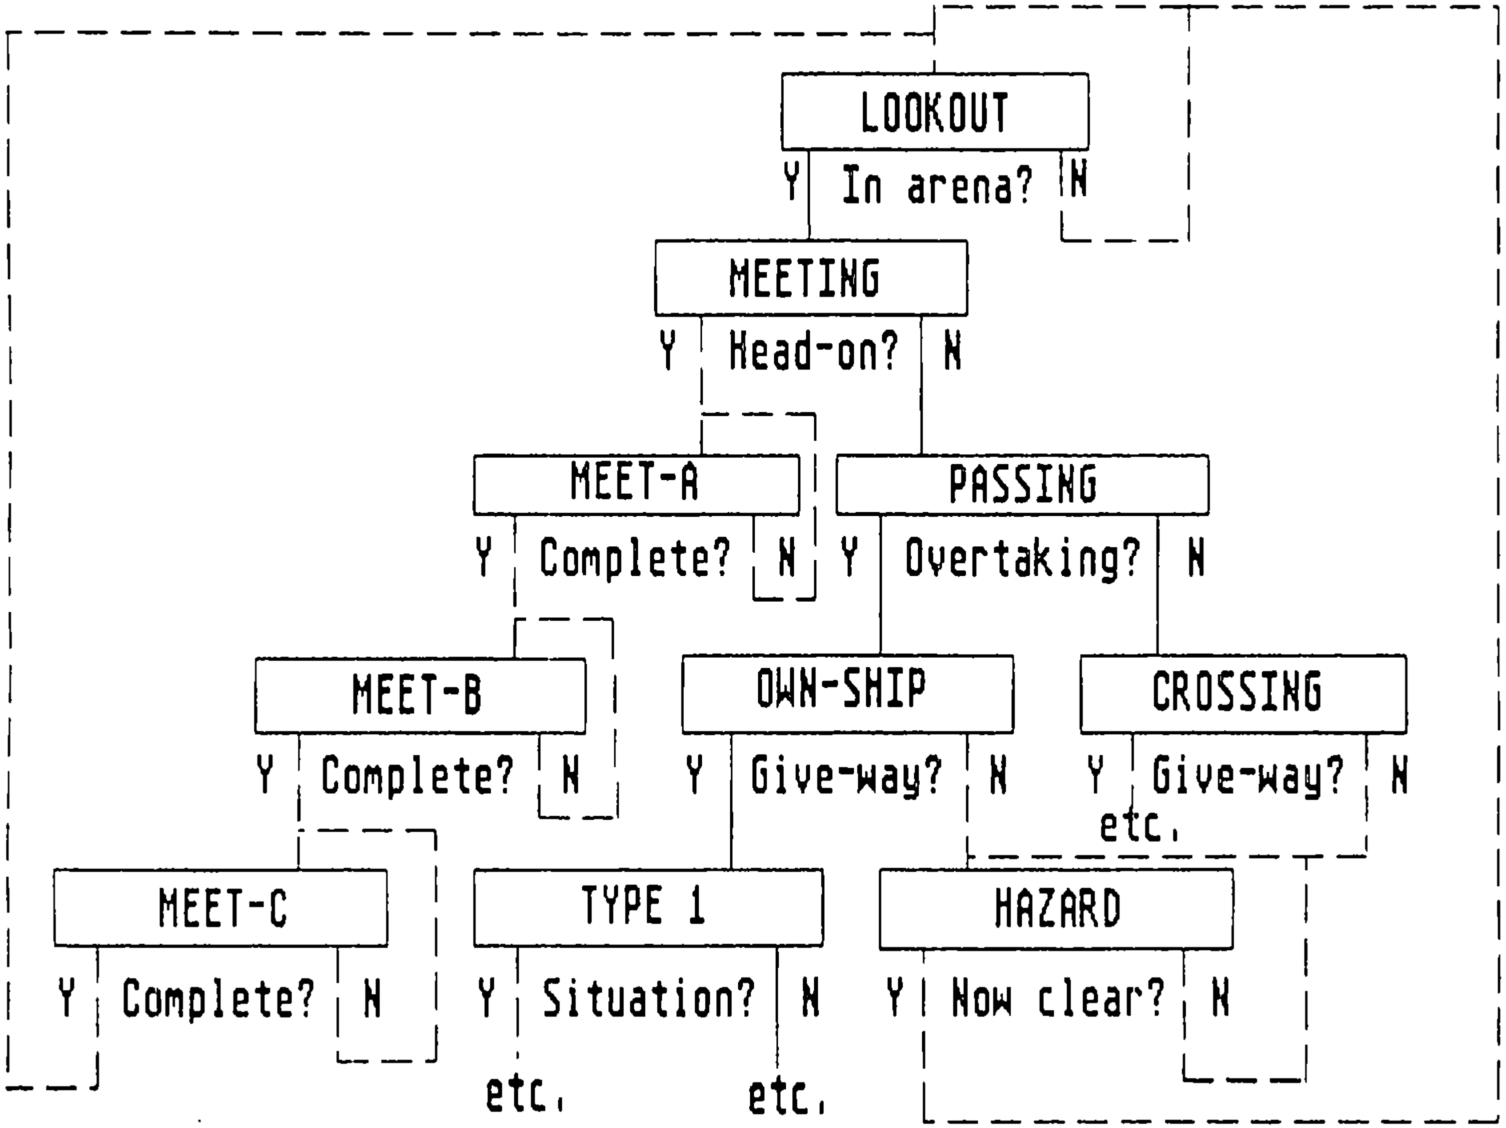

Figure 3.3 Schematic of rule structure.
Dotted lines show deferred action, solid lines show immediate action.
The look-ahead optimisation routine is actioned immediately on the 'Y' branch from the 'lookout' node.

node, with a deferred evaluation of the same condition at the next time-step, otherwise control moves to the next stage of the encounter; action on that subsequent state may itself be deferred for one time interval. The structure also permits transfer from one route through the 'tree' to an alternative route, or 'sharing' of a common end-sequence by two or more encounter types. Control may even loop back, to an earlier point in the same or another encounter sequence. The only specific requirement for each node is that it incorporates a true-false function and associated links to two nodes in the rule structure (either of which may be that same node itself).

Each rule node is a 'C' structure, holding various data pertaining to the current state plus indicators relating to each of the two possible following states:

(a) a function pointer, indicating the boolean function appropriate to this rule;

(b) a brief text string indicating current status, for display in the appropriate screen window - e.g. 'READY', 'MEET-A', 'PASS2-C', etc.

(c) a sub-structure for each of the two alternative paths, holding in each case:

   (i) a structure pointer, pointing to the next node;

   (ii) a flag indicating deferred/immediate action;

   (iii) an indicator as to whether rudder needs adjustment;

   (iv) review text for this situation.

This form of rule structure has been found more than adequate for all purposes to date. Inherent in the design concepts are considerations of likely requirements in more complex situations, notably multi-ship encounters. For example, a switch from 'crossing encounter with Hazard A' to 'head-on encounter with Hazard B' should be totally feasible with incorporation of the appropriate rule(s). Inclusion of more sophisticated emergency manoeuvring should likewise be straightforwardly accomplished by inserting the relevant rule nodes at their appointed places in the decision structure.

The complete set of rules is held as an array of such nodes, cross-referencing between nodes being made by means of array subscripts. A new rule (or more realistically, sequence of rules) may be incorporated by:

(a) adding the node(s) to the end of the existing array;

(b) resetting the link(s) from the relevant existing node(s) and setting the link(s) in the terminal node(s) for the new rule (or sequence); and

(c) providing the required boolean function(s) for the new node(s) - if an existing function will suffice, it may be used repeatedly.

This brief guide to the mechanics of extending the rule base does not, however, adequately convey the technical complexities involved. Any new rule, or rule sequence, must be carefully planned both in respect of its effects and the consequences to the existing rule framework. The core of each node, a boolean function (generally highly mathematical in nature), is not susceptible to formulation, or even clear understanding, by the lay mariner. Moreover, although a sequence of nodes may be likened in some respects to the physical handling of a vessel through the stages of a manoeuvre, the parallel is somewhat obscured by the arcane form of representation. The task of extending the rule base must therefore be seen as a specialised branch of knowledge engineering, the competent mariner co-operating with the computing specialist to devise the appropriate amendments. Naturally, such a task would only be undertaken off-line, never on an actual voyage; thorough testing of the extended rule base should also precede use on active service.

It seems probable that later developments will include a utility to amend/extend the rule base in a 'user-friendly' way. This would include dialog boxes to set up new nodes and re-establish links, a stylised approach to representing the required boolean logic, and a test for consistency and integrity in the modified structure. Such a facility would also be expected to generate a two-dimensional map of the new rule base. This possibility is considered further in section 11.4.

3.4.B The Rule Evaluation Module

This module may be considered to be 'content-free' insofar as the rules themselves are concerned. It makes no presuppositions as to the demands of the rule base, apart from the clear need, whatever the rules, for calculation of a number of standard parameters at each time-step. These include: time to go (on present tracks) before domain infringement by hazard ship (-1 if no domain infringement anticipated); range and bearing of hazard vessel; relative heading, and relative velocity components (W-E and S-N), of hazard vessel; and various other similar factors. At each timer-initiated 'event', the communications input lines are interrogated for the necessary primary data (on own-ship and hazard vessel) to calculate these parameters, before proceeding to the analysis phase.

The analysis, or 'thinking' phase consists of processing the current rule node: evaluation of boolean decision function, using above parameters; testing relevant flags, and setting any further advice/control parameters as indicated; updating status line in window display, and (if appropriate) review text for optional display; advancing rule pointer to next rule as indicated. If the next rule is to be actioned immediately, this process is repeated, until a 'defer' flag for next-rule activation signals the end of the current 'thinking' phase. In simulation or control modes, this module also transmits directives to the control system for own-ship.

A central feature of this module is the pre-evaluation, or 'look-ahead' routine. On identification of a potential hazard, this routine performs an extrapolation in time of various possible avoidance manoeuvres. This involves traversing a number of sequences within the rule tree, each with a range of parameters; a mathematical model of own-ship, held within the system, indicates the likely outcome of each manoeuvre. From this extrapolation, certain control parameters for the best strategy (notably time to initiate the manoeuvre) are calculated, to be used when the actual manoeuvre takes place. Naturally, this strategy will be revised if circumstances change.

Full details on the functioning of this look-ahead evaluation routine are given in chapter 6.

CHAPTER 4

THE KNOWLEDGE BASE

Logic (n): The art of thinking and reasoning in strict accordance with the limitations and incapacities of the human understanding.

Ambrose Bierce (1842-1914), The Devil's Dictionary.

In any IKBS, the operational information is drawn from a variety of sources, which between them portray the 'world picture' as it relates to the decision-making process. In the maritime setting under consideration here, this 'world picture' comprises all the various facets of the environment affecting own-ship, plus all pertinent information on own-ship itself. The following list, though possibly not exhaustive, gives an overview of the range of information required for such a task.

(1) Technical data on own-ship: length; beam; draught; engine performance, and ship speeds under varying conditions (including deceleration times and stopping distances) - parameters for shallow-draught manoeuvring should be held in addition to deep-water parameters; turning curves associated with different rudder settings, and manoeuvring characteristics in general.

(2) Technical data relating to this particular voyage: safety zone to be maintained around vessel (possibly variable with different cargoes); optimisation criteria for voyage - e.g. is speed a higher priority than fuel conservation? (again, variable with cargo, or

scheduling considerations); any special considerations regarding manoeuvring restrictions, if applicable.

(3) Data relating to vessel routeing for the voyage: charted waypoints; navigational channels to be negotiated; potential hazards (charted or notified, e.g. bad weather zone) to be avoided; navigational details for ports of departure and arrival.

(4) Performance characteristics for any hazard vessels encountered (more properly, data for likely hazards, to be fitted as best possible to actual hazards once identified).

(5) Current data (at any specific point in time) on own-ship: position; heading; speed; turning rate (if applicable); rudder setting; engine speed.

(6) Current data on any other vessels in the vicinity, as may be derived from radar and associated aids: relative position; course and speed; any observations as to whether course change is being effected.

(7) Current data on fixed hazards in the area; coastal features; navigational buoys; restricted navigation channels or traffic separation schemes.

It should be noted that these sets of information fall into two clear categories:

(a) static information, which remains constant - at least for the duration of the voyage;

(b) dynamic information, which is continually changing in nature or value throughout the duration of the voyage.

Items (1)-(4) may be classified as static information, whilst items (5) and (6) clearly qualify as dynamic information; item (7) constitutes the current (real-time) points of concern extracted from some of the data held under item (3), and should therefore also be considered as dynamic (i.e. the nature of pertinent data changes with time).

As with any expert system, acquisition of the relevant 'knowledge', and representation in a form susceptible to real-time processing, are not trivial tasks. The following analysis considers the practicalities of acquiring the information detailed in the above categories.

4.1. Own-Ship Technical Data.

Length, beam and draught are immediately available from standard documentation. Speeds and manoeuvring characteristics under varying conditions come under the general heading of 'ship modelling' - a much less clear-cut area.

As explained in section 2.3, the complexity of the manoeuvring model depends primarily upon two factors: the number of 'degrees of freedom' incorporated; and whether a linear or non-linear model is required.

At the current phase of implementation, a very simple model is in use, representing motion by straight lines and arcs of circles. However, it is anticipated that future developments will include use of a 4 degrees-of-freedom linear model - possible within the current structure of the inference engine - giving accurate modelling of manoeuvres within real-time and without the need for extensive sea trials for every vessel to be fitted with the system.

4.2. Technical Data Relating to Voyage.

Complementary to the fixed parameters for own-ship, there is a further set of parameters, dictated by circumstance and personal preference of the ship's master, which apply for the duration of a voyage (or even just for one leg of a voyage).

One obvious case in point is the size of the domain, the safety zone around the vessel which is to be kept clear of any hazards. The radius of this zone will depend on (at least) four factors:

(a) manoeuvrability of the vessel;

(b) nature of the cargo, or task being undertaken;

(c) traffic density in the vicinity; and

(d) temperament of ship's master.

The system allows for input of domain radius, via a dialog box, at the beginning of a (simulation) run. Since a complete voyage could include both open-sea travel and negotiation of a crowded seaway (such as the Dover Strait), variation of the domain radius during a voyage is anticipated; manoeuvring in such a seaway must of necessity be rather more constrained than in the open sea.

The requirements of the domain necessitate specification of a broader region of interest, encompassing any hazard for which current action is required in order to avoid domain infringement. As indicated in chapter 2, this region has been resolved to a time-based criterion, identifying the ideal time for avoidance decisions and (if appropriate) action. This POMT (Predetermined Optimum Manoeuvring Time) is identified dynamically by the system; however, certain parameters used in that decision process are of necessity provided as input at the start of a run (ultimately, at the start of a voyage).

These include: the minimum RDRR, a value below which the POMT will not be allowed to fall so as not to unduly concern other mariners (again, this could vary with conditions); the maximum RDRR, used to ascertain to what extent avoidance may be considered to be the responsibility of the other vessel - i.e. if on working backward from min-RDRR to max-RDRR, at all points conditions indicate responsibility lies with the other vessel, then this may be assumed to be the case (unless and until later circumstances show that the master of the other vessel is not taking action); the emergency RDRR, indicating the latest time at which own-ship may take avoiding action if the other vessel does not act in accordance with the regulations; the ballpark, that broader time-band at which impending encounters should be evaluated to find the appropriate

POMT. Current developments also include the super-domain, a circle of specified radius with own-ship off-centre which is used to classify all vessels in the vicinity which may become involved in an encounter, and should therefore be entered in the knowledge base and monitored on a regular basis (see chapter 10 for details).

Use of optimisation techniques also raises questions as to costs and benefits of different strategies. The course which favours safety, without regard for efficiency or speed, may seem very laudible but prove highly impractical. The safest course of action for any vessel is not to move it from its harbour berth; less ludicrous, but still somewhat excessive, would be to adopt as standard practice the last-ditch measure of describing a complete circle, as a delaying tactic to avoid every impending encounter. Between these over-cautious approaches at one extreme, and a reckless dash at the other, is an infinitely-variable range of options balancing safety against expediency. Such options may be expressed as weighting factors in a cost function, quantifying the relative prominence of different considerations in the decision process. The choice of those factors would relate to circumstances on a particular voyage, with a particular master, and would therefore be included in the performance parameters for that voyage.

At the present phase of development, a nominal set of weightings have been written into the look-ahead routines (see chapter 6). It is anticipated that later developments will include initial specification of appropriate weightings via an input dialog box.

With regard to manoeuvring restrictions, no specific accommodation has been made as yet in the expert system. It is recognised that ultimately restrictions due to disabled equipment (rudder, propeller etc.) or specialised function (dredging, hydrographic surveying) must be catered for within the system. Such constraints will again be expressed as parameters relating to a specific voyage.

4.3. Vessel Routeing Data

At present, the expert system does not take account of any data on navigational features. However, the inference engine and rule base are so structured as to permit incorporation of such data in the future. Two parallel technologies are anticipated as being of direct relevance: electronic chart digital information systems (ECDIS); and automatic weather routeing. Current work on weather routeing [20] centres on a dedicated system which identifies passage way-points so as to optimise voyage management specifically with regard to avoiding bad weather areas. Since the system under consideration here is concerned with the immediate problem of avoiding hazards between such way-points, the task of weather routeing falls within the general sphere of 'passage planning', which is not the responsibility of this system. Apart from the possibility of interfacing the two systems at some later date, no action is anticipated on weather routeing in this expert system.

The present state-of-the-art in the field of ECDIS [21, 22], whilst providing some impressive screen graphics, fails to come up to the mark in respect of the requisite digital data for processing by electronic decision systems. It is not possible, for example, to interrogate any (known) present ECDIS to elicit answers to such questions as: how far to port is the 20-metre contour? if I turn 30 degrees to starboard, will I encounter any navigational hazard within 4 cables (0.4 nautical miles)? what course should I steer to line up with the deep-water channel into the harbour?

To date, the International Hydrographic Organisation have issued guidelines as to standards required in implementations of ECDIS [23]. These include such factors as:

(a) use of colours on VDU screens;
(b) minimum sizes of symbols for navigational features where charts are scaled down for display;
(c) content of default display, on start-up;
(d) facilities for clear display of technical data on chart features - e.g. light and radiobeacon data (where appropriate) for a navigational beacon;
(e) optional blanking-out of non-essential data on screen at any given time;
(f) refresh speed for display;
(g) accommodation of chart updates.

In short, defined standards to date are concerned with the form of display of chart data. The corresponding requirement for a digital database for intelligent electronic decision-making has not, as yet, received the same consideration.

The nature and format of such a database is a matter requiring careful planning. An electronic database should be capable of providing data in real time on a variety of features. One obvious area of concern is the depth below the vessel, and variation of depth in different directions radiating from the vessel; this implies either a close-mesh grid of depths for the whole area of the chart (very expensive on memory) or some form of digital representation of depth contours (extremely difficult to analyse in relation to one fixed point, and directions from that point). Another area of interest is that of navigational aids - buoys, lights etc. These may simply be held each as a record of reference information (in which case all such records must be checked every time, in case any is relevant), or somehow cross-referenced to the close-mesh grid suggested above. Following the latter idea to its logical conclusion, each grid square becomes a data node, holding various information about that small area; such an approach could prove totally unworkable, both in the time taken to convert paper charts to such a form, and in the memory requirements of a set of such charts.

The 'least-cost' option, of extracting digital data directly from the video disk (or screen display) graphical representation, may be dismissed on grounds of extreme inefficiency, time-wise. What is needed is a representation which permits fast retrieval of data pertinent to a specific chart position, for processing in real time.

Formulation of a definitive standard for such data is both essential and extremely difficult. If ECDIS databases are created on an ad hoc basis by various organisations, then devising a standard interface for such databases to decision-making systems could prove impossible. At the same time, any defined standard for digital chart data must take account of all present and likely future requirements on such data - at least on an 'upward compatability' basis.

Given this impasse, the only way forward with the collision avoidance expert system is to make certain assumptions - and leave the way open for modification if those assumptions prove false (hopefully, the final solution will not be too far removed from any intelligent informed conjecture!). One line of advance on this research [24] is currently identifying and categorising the informational needs of this expert system and similar systems with respect to digital chart data. The next step will be to digitise an area of chart in the form so defined, and design a software interface enabling the expert system to use such data in its decision-making.

A viable operational expert system will require such chart data for the whole of a passage - i.e. the full area for which conventional navigational charts would normally be carried. At the 'passage planning' phase referred to earlier, it is anticipated that way-points for the voyage would be 'marked' on the database in some way; one defined objective of the expert system would then be to correct back towards the next waypoint each time a course alteration is made for avoidance purposes.

To sum up: the informational requirements for digital chart data are fairly clear; as yet, no defined standard has yet been formulated; development from the current stage of the expert system must therefore anticipate (and hopefully be directly involved in) such advances in this parallel technology.

4.4. Hazard Vessel Performance Characteristics

As indicated in chapter 2, the requirement for a detailed model for own-ship is not mirrored by such a requirement for observed hazard vessels - nor is such a detailed model a practical proposition. However, there is some merit in having at least an approximate picture of the likely manoeuvring capabilities of a potential hazard. It would then be possible, for example, to define more precisely the earliest and latest times at which such a hazard might take avoiding action, if they are deemed to be the give-way vessel. This would aid in decision-making as to whether emergency avoiding action is advisable.

Until identification transponders are fitted on ships (as on aircraft), it is clearly impracticable to expect automatic hazard identification by the expert system. Some degree of identification may be possible from speed and size of radar signal, but such a system would necessarily carry a high margin of error. It would seem preferable to operate on a standard 'default' setting, with the possibility of operator override if a vessel type is visually identified.

Since the intentions of the master of a hazard vessel cannot be known (though they may sometimes be deduced), a sophisticated model of a hazard ship yields little or no benefit over a more simplistic model. Any enhancement of the default setting would therefore be at a fairly simple level - adjustment of turning circle and/or deceleration rate (though these may vary substantially between, say, a cross-channel ferry and a VLCC (Very Large Crude Carrier)). It would suffice to have a table of reference data for, say, ten standard ship types, and match an observed hazard to the closest of those types.

At the present stage of implementation, hazard manoeuvring is considered only on the basis of either constant-velocity motion or (if deemed the give-way vessel) the appropriate avoidance manoeuvre in good time; the latter consideration is not at all specific, except insofar as identifying the latest possible time for emergency action if such a manoeuvre does not appear to be forthcoming. The simulation of the hazard vessel in the development environment is a different issue, covered in chapter 8.

It is intended that, as a later development, data on a comprehensive range of ship types should be added (see section 11.5). The operator would then have the option of identifying a specific hazard as one of those types, via menu selection. Such data could be utilised by the rule base to refine decisions on:

(a) responsibility for manoeuvring; and
(b) timing and action, should emergency manoeuvring prove necessary.

4.5. Current Data on Own-Ship

One overriding concern in the collision avoidance task is the immediacy of its application - if it can't be done now, there's no point in doing it at all. More than that, the system must actually be anticipatory, rather than just considering past or present circumstances. This implies a constant process of data capture and analysis with very strict real-time constraints; that analysis must give an appraisal of the future situation - in some cases, several alternative future situations, depending on action taken in the interim period.

To achieve such a goal, it is first necessary to constantly have a current picture of the state of own-ship. It is further necessary to be able to extrapolate from that data what the state of own-ship may be at various times in the future, given different possible control inputs to the total ship system.

One major consideration is that of the frame of reference: should all data be assessed relative to own-ship, or is there a fixed frame of reference from which all measurements should be taken? In the open sea this decision is less clear-cut than in coastal areas, where certain very firmly fixed features demand consideration. Even far from land, however, waypoints are based on charted positions, and fixed traffic separation zones may apply. Given the position-fixing systems now in operation (Decca, Loran-C, GPS), absolute measurements would appear to be the convenient and expedient approach. Relative measurements from radar can be simply converted to absolute terms, giving true position and motion (this is also considered beneficial as regards screen display format - see section 7.1).

Simulated data used in the present version of the expert system comprises own-ship position (with respect to a fixed reference grid), speed and heading, plus current rudder setting. It is anticipated that a live implementation would receive such data from sensors on engines, rudder and gyro compass, plus direct input from GPS or another position-fixing system. Such live data would, of course, be subject to certain error bounds; these would need to be determined and allowed for - possibly by a percentage increase in domain size, and other similar safety margins. The topic of 'error analysis and compensation on live data input' is a substantial area for investigation at a later date.

4.6. Current Data on Hazard Vessel

The only form of communication that can be guaranteed between two vessels at sea is that of mutual observation (even that may be difficult in bad conditions). Radio contact is a bonus - but that can itself lead to hazardous situations, if such communication is relied upon in place of the standard maritime protocols [25]. As previously stated, ships do not as yet carry transponders for identification.

It follows that the only data which the expert system may use, in planning strategy with respect to an observed hazard, comprises observations which may be made through available instrumentation -

notably radar. It is not uncommon practice for mariners to identify a familiar vessel, and even surmise the likely intentions of its master - such refinements, in terms of user input, are being considered for the future.

An ARPA (Automatic Radar Plotting Aid) provides data on position and velocity for a number of 'target' vessels, either picked out by the operator or acquired automatically as they enter pre-set 'guard bands'. Such data enables the expert system to calculate the future position of such a vessel, on a constant-velocity basis. This information is used to determine:

(a) whether the vessel constitutes a hazard on current heading;

(b) where the responsibility lies for giving-way, should the need arise;

(c) what is the best time and course of action, if own-ship is the give-way vessel.

Velocity data may also, in the future, assist in identifying the type of the hazard vessel and filtering it to one of the standard types referred to in section 4.4 (see also section 11.5).

The present expert system uses simulated data of the type expected from an ARPA; as detailed in appendix B, an ARPA has been fitted to a research vessel, and developments in direct data acquisition will commence shortly. Clearly the same caveats which apply to data integrity for own-ship are also applicable here, and appropriate error margins/strategies will have to be considered.

4.7. Current Data on Fixed Hazards

Section 4.3 considered in detail the requirement for navigational chart data for the complete passage. This section is concerned with the need, at any given time, for data on any features needing consideration in the immediate vicinity. I.e. real-time retrieval of relevant chart data, for inclusion in the decision process. This lends a rather new perspective, since own-ship will, to all intents and purposes, be at some point defined within that database, and will need access to data on other features at, or in close proximity to, that same point. This does not correspond to the requirements of any other database structure known to the author. Developments are awaited with interest - and hopefully active participation.

CHAPTER 5

RULES FOR THE TWO-SHIP SITUATION

'Alas! (thought I, and my heart beat loud)
How fast she nears and nears!
Are those her sails that glance in the sun,
Like restless gossameres?'

Samuel Taylor Coleridge (1772-1834) The Rime of the Ancient Mariner.

When devising a complex system for an advisory or control function, it would seem sensible to start with the simplest possible situation, then build on that. In so doing, care must be taken that such an approach does not lead down a blind alley, precluding expansion by the very nature of the solution. At all times in the development of the rationale the way must be left open for growth, in whatever direction is considered necessary at any later stage.

The present expert system is designed on this principle, with two levels of expansion envisaged for the future. At the simplest level, handling of various encounter types may be enhanced by extension of the rule base as required - greater flexibility in identifying and countering a wide range of threatening situations, improved judgement and dexterity in dealing with emergencies, more sophisticated response to changing circumstances. At a higher level, the underlying principles of the system may be modified, without major upheaval, to incorporate new features - interfacing of electronic chart database, provision of a more discriminating cost function for decision optimisation (see chapter 6), variation of hazard performance model to fit identified vessel type.

The simplest possible encounter situation, and hence the basis for this expert system, is that of two vessels in the open sea. This reduces avoidance considerations to the simple expedient of manoeuvring so as to avoid the other vessel (or leaving such manoeuvring to that other vessel, if appropriate), without concern over knock-on effects (literally!) engendered by other vessels or fixed hazards in the vicinity. The intention is that such effects may be dealt with in later enhancements of the inference engine and/or rule base - indeed, recent work building on this research has shown the efficacy of this approach [6, 7].

5.1 Rules of Engagement

In any two-ship encounter, the nature of that encounter is determined by the relative positions of the two vessels at the time of evaluating the situation. That time is itself indeterminate, since the collision avoidance regulations state that avoidance action should be taken 'in good time' - the possible consequences of such a subjective requirement will be considered shortly. The juxtaposition of the two vessels in contention is viewed in relation to the navigation lights of those vessels, and the sectors subtended by those navigation lights.

Figure 5.1 illustrates the three sectors subtended by the navigation lights of a ship: green (starboard), red (port) and white (rear). If a vessel is approaching another within the white sector of the latter, it

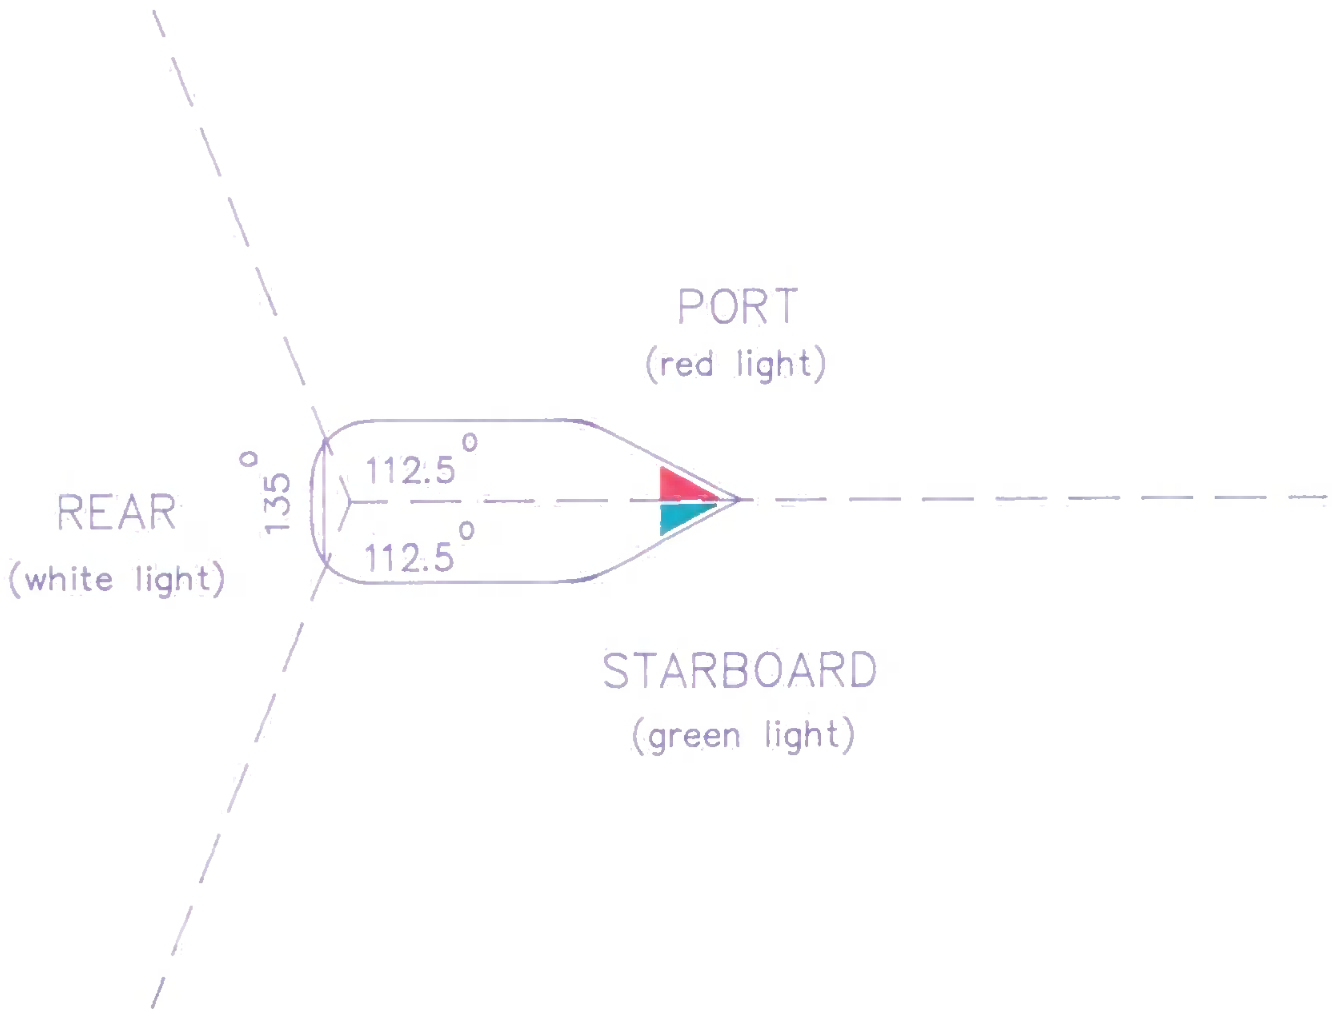

Figure 5.1 The three sectors around a vessel, as designated for navigational purposes and subtended by the ship's navigation lights.

is deemed to be overtaking; if within the red sector, crossing from port side; if within the green sector, crossing from starboard side. The only exception to these definitions is where the two vessels are approaching on near-reciprocal courses (i.e. within a few degrees of head-on); this is deemed a head-on encounter.

The collision avoidance regulations are very specific in certain respects, less so in others. The phrases 'in good time' (for avoidance decision and action), 'a clear turn' (as indicative of intent), and 'maintain a safe distance' are all phrases which, although eminently reasonable, leave a degree of subjective interpretation with the mariner (and possibly slightly different interpretations with the masters of two vessels in contention).

Figure 5.2 illustrates the five basic encounter situations, as identified in the regulations. The requirement on each vessel is as follows (taking hazard vessel as ship B in each case):

(1) Head-on

Both vessels are required to alter to starboard, maintain a safe distance, then alter-back when safely past;

(2) Crossing from port side (ship A)

The vessel observing the other's red light (i.e. ship A) is required to alter to starboard, pass a safe distance behind the other vessel (ship B), then correct back onto course when safely past; the vessel observing the other's green light (ship B) is not required to alter course unless the other vessel is seen not to be taking appropriate action.

(3) Crossing from starboard side (ship A)

The roles of ship A and ship B are reversed with regard to responsibility for avoidance, as indicated in (2);

(4) Overtaking from port quarter (ship A)

The overtaking vessel is required to alter to starboard, pass behind at a safe distance, then correct back onto course when safe to do so;

(5) Overtaking from starboard quarter (ship A)

The overtaking vessel is required to take avoiding action as follows: if it is safe to do so, it may alter to starboard, parallel-up and overtake at a safe distance on the starboard side of the vessel being overtaken, then correct back onto course and cross at a safe distance ahead (observed by Colley [27] to be an accepted practice); otherwise, alter to port and cross at a safe distance behind, correct back onto course when safe to do so.

It is also stated in the regulations that, once an encounter has been identified as being of a particular type (as defined above), then the procedure for that encounter type should be followed through. A change in apparent circumstances does not warrant a change in policy, unless occasioned by untoward action on the part of the other vessel (or, exceptionally, mechanical failure of some sort).

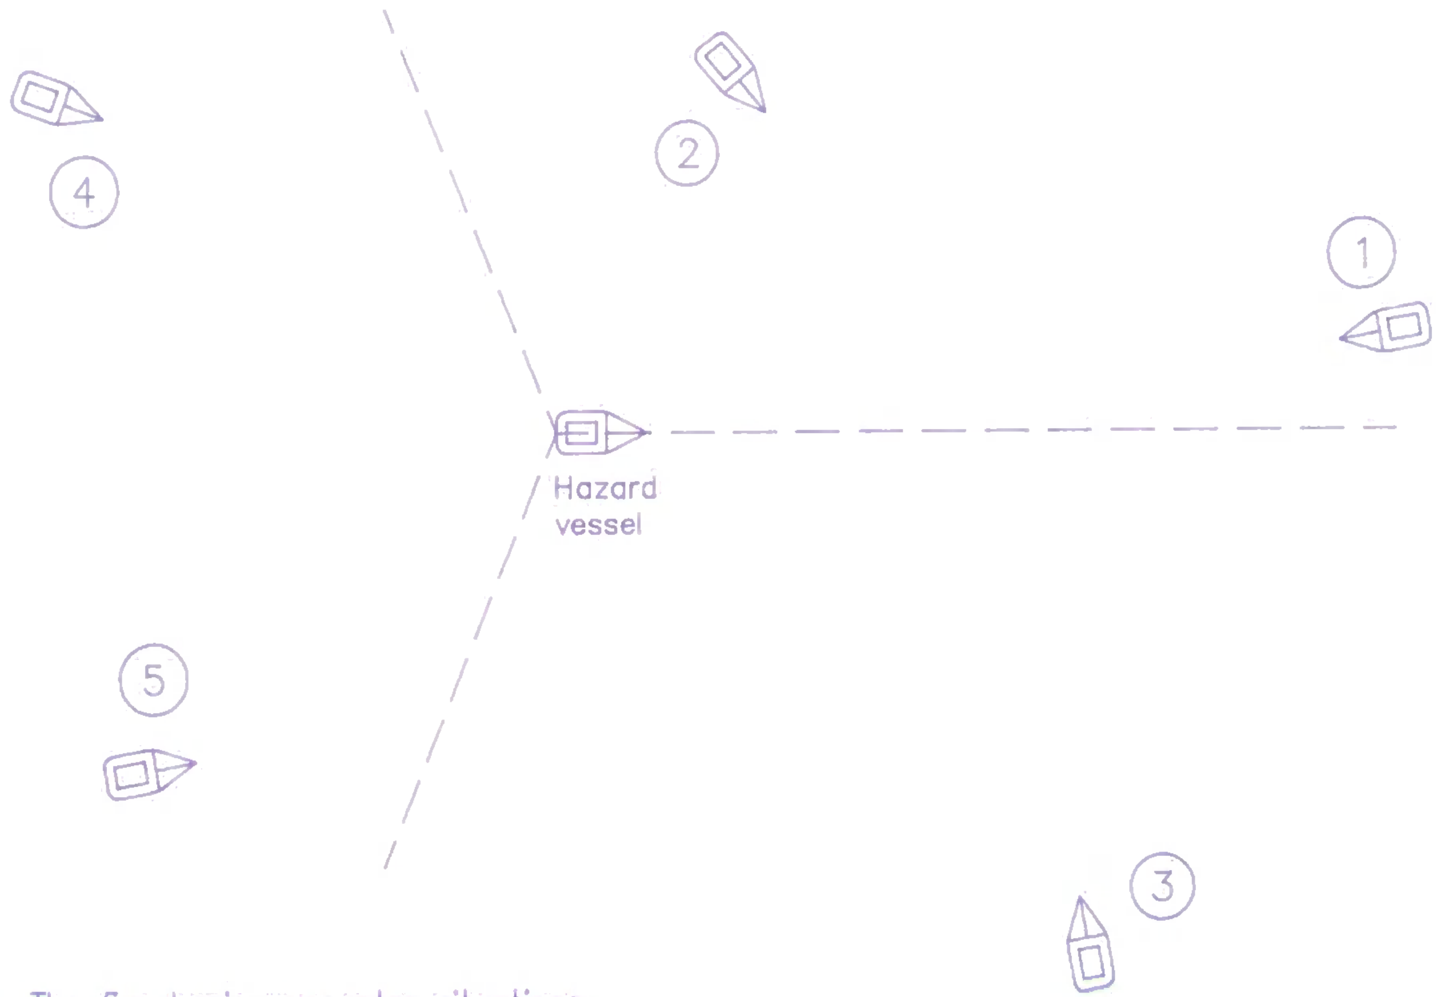

Figure 5.2 The five basic encounter situations:

(1) head–on
(2) crossing from hazard's port side
(3) crossing from hazard's starboard side
(4) overtaking from hazard's port quarter
(5) overtaking from hazard's starboard quarter

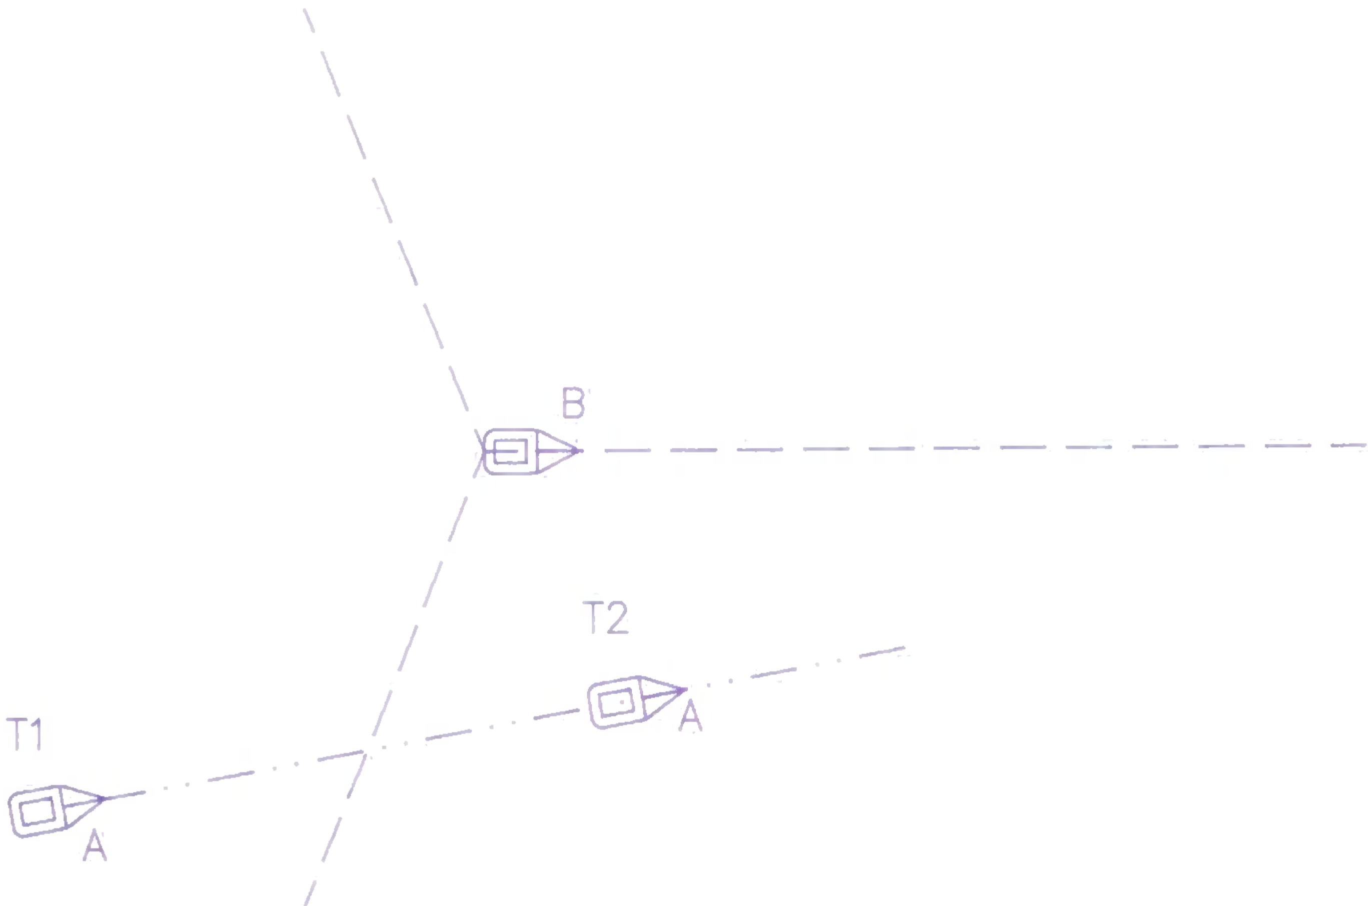

Figure 5.3 Two differing perceptions of a single encounter situation.
At time T1, A would be seen as the give–way vessel in an overtaking encounter.
At time T2, B would be seen as the give–way vessel in a crossing encounter.

This point is illustrated by figure 5.3. Vessel A is approaching vessel B from behind, on the starboard quarter. At time T1, A is in B's white (overtaking) sector; shortly afterwards, at time T2, A is in B's green (starboard) sector. The implication of this is that: at time T1 the situation would be evaluated as an overtaking encounter, with vessel A having responsibility for avoiding (by turning to port, and behind B); whereas at time T2 it evaluates as a crossing encounter, vessel B now being responsible for giving-way by altering to starboard and behind A.

There are two points to be drawn from this illustration:

(1) if the encounter is assessed at time T1, then the assessment as an overtaking encounter, and the consequent responsibility for action, should be followed through irrespective of the fact that the situation may shift to a crossing encounter at some stage in the process;

(2) more significantly, if master B assesses the encounter at time T1, he will see it as an overtaking encounter with A the avoiding vessel, and will adopt a stand-on stance; if master A then assesses the situation at the later time T2, he will regard it as a crossing encounter with B to avoid, and will also stand-on; such a (valid) difference in outlook could well lead to emergency action by either or both vessels at some time in the future.
Just such a difference of opinion was the root cause of a collision between the 'Nowy Sacz' and the 'Olympian' in 1972 [26]. The fact that the court of appeal reversed both the judgement and the

apportionment of blame (as given by the lower court) is indicative of the disparity of views possible in such circumstances.

The regulations also allow for the possibility of a course change at some earlier stage, so as to obviate the need for any avoidance manoeuvre (see section 2.2). Whilst this derestriction is interpreted in a variety of ways by mariners, a definitive system such as this must have a clear, self-consistent approach - this does not preclude flexibility - which makes due allowance for whatever interpretation is put upon it by the master of a potential hazard vessel.

5.2. The Expert System Rules

The rules as defined in the expert system, for collision avoidance in the two-ship situation, follow the requirements of the collision avoidance regulations, in a manner which seeks to remove any uncertainty or subjectivity from their interpretation. The basic 'standard' avoidance manoeuvre is as follows:

(1) alter course (to port or starboard, as appropriate) by at least a standard 'minimum turn' - this is defined as a system parameter, set at 15 degrees on some occasions, 45 on others (covering the observed range used by mariners); the turn may need to continue beyond this value to ensure domain clearance, or in certain circumstances to meet specific requirements of a particular rule;

(2) proceed in a straight line on the new course until

(a) returning to original course would no longer threaten domain infringement, and

(b) any specific requirements for this type of encounter have been met;

(3) alter back to original course; proceed parallel to original track.

In the open sea, correction back onto original track to pick up a way-point is not a major consideration, since this may generally be accomplished over a protracted distance. This would be the responsibility of the (computerised) navigation system (i.e. an autopilot) rather than the collision avoidance system.

The expert system functions by continually cycling through a system of rules defined in a form based loosely on a binary tree structure - see section 3.3. At a particular time-step, the vessel's current state will be identified by one of those rules (indicated by a system pointer), with either a change in state to a new rule or continued/repeated application of the same rule at the next time-step.

Initially the system cycles on a 'look-out' rule, which notes any hazard on course for domain infringement. Once such a hazard enters the 'ballpark' (a time-based boundary, currently set at 20 minutes to domain infringement), control transfers to a new rule which prepares for the encounter. Figure 5.4 illustrates the concept of the 'ballpark' in relation to three potential hazards travelling at different speeds.

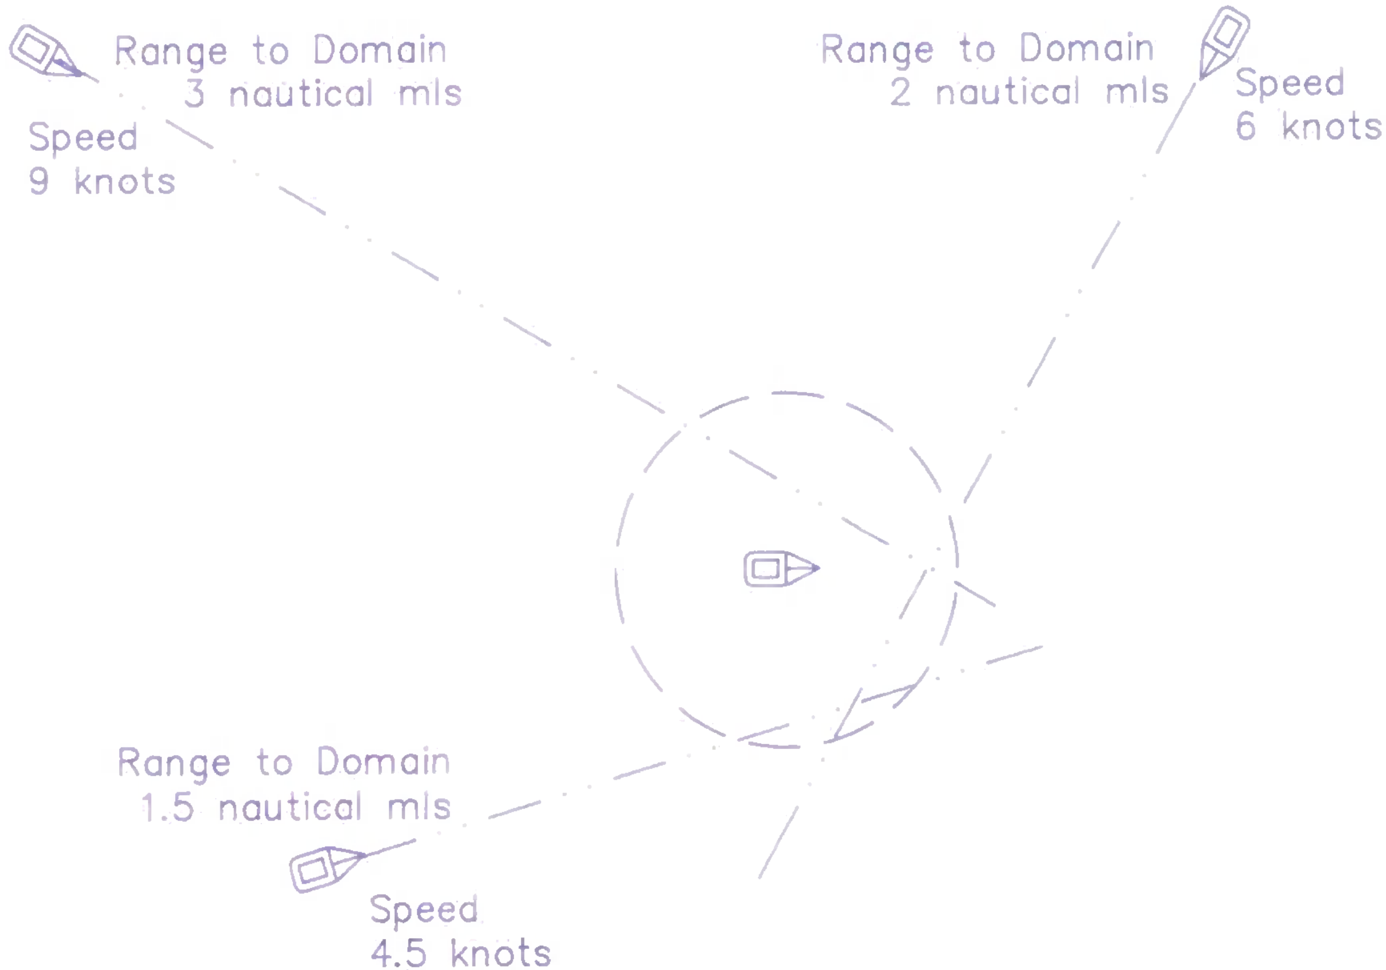

Figure 5.4 Illustration of the 'ballpark' concept as a time-based threshold. The three hazard vessels shown are at differing distances from own-ship's domain, but all exactly 20 minutes from projected domain infringement – a 20-minute ballpark.

On entering the ballpark, a hazard triggers the look-ahead mechanism (described in detail in chapter 6), which identifies the nature of the impending encounter and the optimal manoeuvre with which to counter the threat. The timing for this manoeuvre (POMT - see section 2.1) is noted, and the 'ready' rule is invoked.

In the 'ready' state, own-ship is allowed to proceed without any change in course through successive time-steps, until the POMT is reached. At this point, the first stage of the appropriate avoidance manoeuvre (which might be 'stand-on') is selected, and the system proceeds to work through the rules for the particular encounter type under consideration. The sequence for each type of encounter is described in detail below.

(1) Head-on

Two vessels are approaching each other on reciprocal bearings, plus or minus a given tolerance (typically, 10 degrees). A vessel in such an encounter must consider itself give-way, altering course to starboard (right) by the required turn as defined earlier. It will continue on the new course until the other vessel has passed abeam to port (left) - i.e. its relative bearing from own-ship is between 270 and 180 degrees. It will then alter-back onto original course.

(2) Crossing

Two vessels are approaching such that one can see a green light, and the other a red light, on each other's vessel. The ship viewing red - i.e. such that it is on a relative bearing greater than 247.5 degrees, viewed from the other ship - should give-way, altering to

starboard for the required turn as defined earlier. It should continue on this new course until the other vessel has passed 15 degrees to port of own-ship's original heading (i.e. if own-ship were instantaneously back on course, the other vessel would be on a relative bearing of less than 345 degrees, having crossed the bows of own-ship). It will then alter-back onto original course.
The ship viewing green - i.e. on a relative bearing of less than 112.5 degrees seen from the other ship, adopts a stand-on stance.

(3) Overtaking

An overtaking vessel (relative bearing between 112.5 and 247.5 degrees, seen from the other ship) will give-way as follows:

(a) if on a near-parallel course (to within a fixed tolerance), go behind to starboard by required turn as defined earlier; alter back onto course once doing so will not infringe the domain;

(b) if on port quarter of stand-on vessel (relative bearing from other ship between 180 and 247.5 degrees), action as for (a);

(c) if on starboard quarter (relative bearing from other ship between 112.5 and 180 degrees) <u>and</u> current course would take own-ship ahead of hazard <u>and</u> it is possible to parallel-up (by turning to starboard) without infringeing domain, do so; when it is possible to return to course (passing ahead) without infringeing domain, alter back to port, onto original course;

(d) otherwise, turn to port by required turn as defined, pass behind hazard vessel; alter back onto course (to starboard) once such action will not cause domain infringement.

The additional clearance requirements included in some of the above rules are taken from earlier research by Colley [27]. A simulation of ship encounters, and consequent manoeuvres, used criteria based on observation of common practice by experienced mariners. It seems possible that a system of rules which does not involve such (apparently fairly arbitrary) figures could perform equally well, though not necessarily in a style familiar to mariners.

This raises the broader question of the extent to which the expert system should actually mimic conventional human behaviour, as opposed to simply performing the task competently and safely, making best use of its electronic capabilities. This is a significant issue, since it hinges on the relative strengths and weaknesses of the computer-based system as contrasted with the human officer of the watch. For example, the change in aspect of a vessel starting to turn would be readily observable very quickly by eye, but a sequence of several radar readings would be required to confirm this electronically; by contrast, the human observer would not readily extrapolate the time and location of the CPA (Closest Point of Approach) of a target vessel, whereas his electronic counterpart could provide such information in a fraction of a second.

It is inappropriate, however, for the computer-based expert system to implement the optimal electronic solution without regard for human conventions. Since a vessel communicates its master's intentions primarily by its actions, unconventional behaviour could lead to uncertainty on the part of other mariners. It is therefore important to strike a balance between the most efficient performance achievable by such an expert system and a style of operation which will be familiar and reassuring to other mariners in the vicinity. This is not simply an affirmation of the need to conform to the official regulations - quirky behaviour or 'brinkmanship' within the constraints of those regulations could still be most disconcerting to experienced mariners. Conventional behaviour is a valid aim in itself.

As a consequence it is considered essential that the strategies implemented by this system should be acceptable to experienced mariners, and recognisable as conforming with standard practice. A substantial element in the validation of the rules and their mode of implementation consists of just such a test of acceptability - see chapter 9.

5.3. Emergency Manoeuvring

The rules defined in section 5.2 have been formulated on the premise that both vessels in the encounter behave in conformity with the anti-collision regulations. This is not always the case; the master of a hazard vessel may fail to take the action required of him, or even act in such a way as to aggravate the situation.

Such circumstances are not necessarily attributable to negligence or incompetence. As indicated in section 5.1, two well-intentioned masters may, in good faith, allow their vessels to reach a point where normal safety manoeuvring will not resolve the potential conflict. Nonetheless, whether due to inattention, incompetence, inclement conditions or incompatible assessments of the situation, such conflicts must be resolved. There is therefore a need within any suite of collision avoidance strategies for a contingency plan (or preferably, plans) for emergency manoeuvring.

The criteria for decisions on emergency manoeuvring are by no means clear-cut. In particular, it is very difficult to judge, before it is too late, whether the master of a hazard vessel is in fact neglecting his responsibility to avoid, or whether his interpretation of the term 'in good time' is simply rather less generous than one's own. Clearly, if one is forced (or at least, persuaded) to implement an emergency avoidance manoeuvre, it should be done:

(a) in time to be effective, without pre-empting a genuine safe manoeuvre by the other vessel in such a way as to aggravate the situation;

(b) so as to defuse, rather than compound, the problem should the master of the other vessel also attempt to take action belatedly.

In transferring from a human agent to a computer system for such decision-making, this takes such reasoning from the intuitive to the definitive. It is not practicable to act on the basis of what 'looks right' or 'seems appropriate'; the rationale behind any decision to take emergency action, and what that action should be, must be spelt out in definitive terms (albeit possibly involving an element of 'fuzzy logic').

The principal aim of this computerised avoidance system is to preserve the integrity of the domain; a secondary concern is the selection of a strategy which achieves this with optimum efficiency (see chapter 6). All other considerations stem from these two objectives. It follows that, where required to do so under the regulations, own-ship will manoeuvre so as to clear the domain but not by an excessive margin.

It is reasonable to expect that any other vessel, whether under human or computer-assisted guidance, would perform likewise when required to give-way. Following through such logic, it would in many cases be premature to assume negligence on the part of the (master of the) other vessel until domain infringement is virtually inevitable. This presupposes, of course, that the other vessel's manoeuvring capabilities, and her master's perception of domain size, are comparable to one's own.

One possible response to such argument is the though that "Surely the master of the other vessel could be expected to give a good wide clearance, if acting under non-emergency conditions? If this was not forthcoming, then emergency manoeuvring could be undertaken by own-ship with still ample, if less, safe clearance". This is, of course, exactly what one would expect in such circumstances.

There is no contradiction here with the principles embodied in the computerised system under consideration. Indeed, the reasoning outlined above serves to both define and highlight those principles. In essence, there is a requirement for a substantial domain which is to be preserved under normal manoeuvring and a reduced domain for emergency manoeuvring. The proviso that the emergency domain should not be violated under any circumstances (barring unavoidable aggressive manoeuvring on the part of a hazard) will determine the threshold at which own-ship (as the stand-on vessel) decides that emergency avoidance action is necessary.

The criteria for emergency manoeuvring, then, are much the same as for normal manoeuvres. Once the hazard ship has been identified as the give-way vessel, with own-ship stand-on, the look-ahead routine (see chapter 6) may be invoked further to assess the situation re. emergency manoeuvring. This involves identification of the latest time at which own-ship may manoeuvre if the emergency domain is not to be breached - if the hazard vessel continues on its present course.

In some cases this may require no course alteration, since hazard vessel's track may clip own-ship's full domain but not the emergency domain. In other cases, one of a range of options may be selected to avoid a near-miss or collision. These options would form a branch, or set of branches, of the rule structure.

As yet, the full range of possible emergency manoeuvres has not been considered in detail. A rule has been incorporated to identify the need for emergency manoeuvring - when own-ship is stand-on and the hazard vessel shows no intention of manoeuvring. The latest possible time for manoeuvring is calculated, and at this time own-ship performs a complete 360 degree turn - one full circle.

This somewhat extreme manoeuvre is virtually guaranteed to be effective in a two-ship encounter in the open sea. It could prove less so in restricted waters, or where there are a number of other vessels in the vicinity. However, it at least serves to demonstrate the efficacy of the rule base in accommodating rules for emergency situations; further, more sophisticated rules to this end may be added without difficulty in the future.

One further point for consideration here is the use of a mathematical model for hazard vessel manoeuvring. Such a model may indicate the ability of the hazard to manoeuvre safely at a late stage, when own-ship manoeuvring would not even be able to preserve the emergency domain. The question then arises as to whether one trusts to the other vessel until it is too late for own-ship to avoid, or whether one 'plays safe', possibly pre-empting a clear, safe manoeuvre by the hazard vessel. This is a subject for further consideration at a later date.

CHAPTER 6

LOOK-AHEAD SIMULATION

"I have seen the future and it works."

Lincoln Steffens (1866-1936) US Journalist: Biography.

6.1. Problems of Avoidance Strategy and Timing

When considering the current state of shipping around his own vessel, the competent master will first identify those vessels in the vicinity which could at some future time pose a threat to own-ship. Such information may, in the first instance, be provided via collision warning facilities on the ARPA. Some extrapolation of tracks and speeds - i.e. velocity vectors - may be applied to estimate likely time(s) and position(s) of projected collision or near-miss situation(s). At a more fundamental level of observation, any vessel seen to be approaching on a constant relative bearing - i.e. not changing its bearing relative to own-ship - will be considered as a likely collision threat.

When a vessel is seen to pose a threat, the next step is to formulate a policy for avoidance. This will incorporate both a time for action and a strategy to be implemented at that time. In general, the detailed strategy will not be worked out until close to the time for avoidance action, since the situation will become clearer, and may even change, as that time approaches. However, some form of avoidance plan must be decided on in principle, in order to identify the time at which action must be taken.

In devising a strategy for avoidance, a number of criteria must be taken into consideration. First and foremost, any strategy must conform with the International Regulations for Preventing Collisions at Sea [3]. Additionally, it must make allowance for any other fixed or moving hazards in the vicinity - rocks, buoys, other ships - as well as navigational requirements such as traffic separation zones. Last, but still of significance to shipping operators (and therefore ship's masters), manoeuvring should not lead to excessive deviation from track, loss of speed, loss of time or any other 'amenity', beyond what is prudent.

Any avoidance manoeuvre, then, is a balance between safety and expediency. If the correct balance is to be achieved, it is important to ensure the former without undue sacrifice of the latter; hence the need to plan ahead.

Just as the human agent has to forward plan for collision avoidance, so a computerised system designed for this purpose must likewise evaluate possible future scenarios on the basis of both safety and economy. In this respect it is vastly better equipped (from the analytical point of view) than its human counterpart, since a number of possible 'endgames' may be evaluated in a very short time. This advantage must, of course, be balanced against intuitive evaluation of situations by the human - the consequence of years of wisdom and experience, either accumulated at sea or handed down by others.

Since safety is the primary consideration, it is first necessary to devise an approach which will guarantee this (insofar as that is possible). Once certain predefined safety conditions have been met, various alternative strategies which satisfy those conditions may be compared for cost-effectiveness in terms of factors considered relevant.

6.2. Look-Ahead Simulation for Safe Manoeuvring

As outlined in chapter 5, any collision avoidance manoeuvre (including stand-on or emergency action) is conducted by traversing a specific sequence of rules within the rule structure. This presupposes that, at the time it is reached, a rule is able to provide the stimulus for avoidance action in time for it to be effective. Since each rule involved in a particular manoeuvre is one in a chain of events, it is necessary to time the initiation of the first event in the sequence such that the entire chain completes satisfactorily.

The following example, illustrated in figure 6.1, may serve to clarify this point.

> Ship A is engaged in an overtaking manoeuvre, approaching ship B on the port quarter. A must first turn to starboard sufficiently to clear B's domain on the next phase - which is to travel on a new straight course behind B. (Two points: (1) A's course alteration must also satisfy the 'minimum observable turn' criterion; and (2) it makes no difference whether the domain is

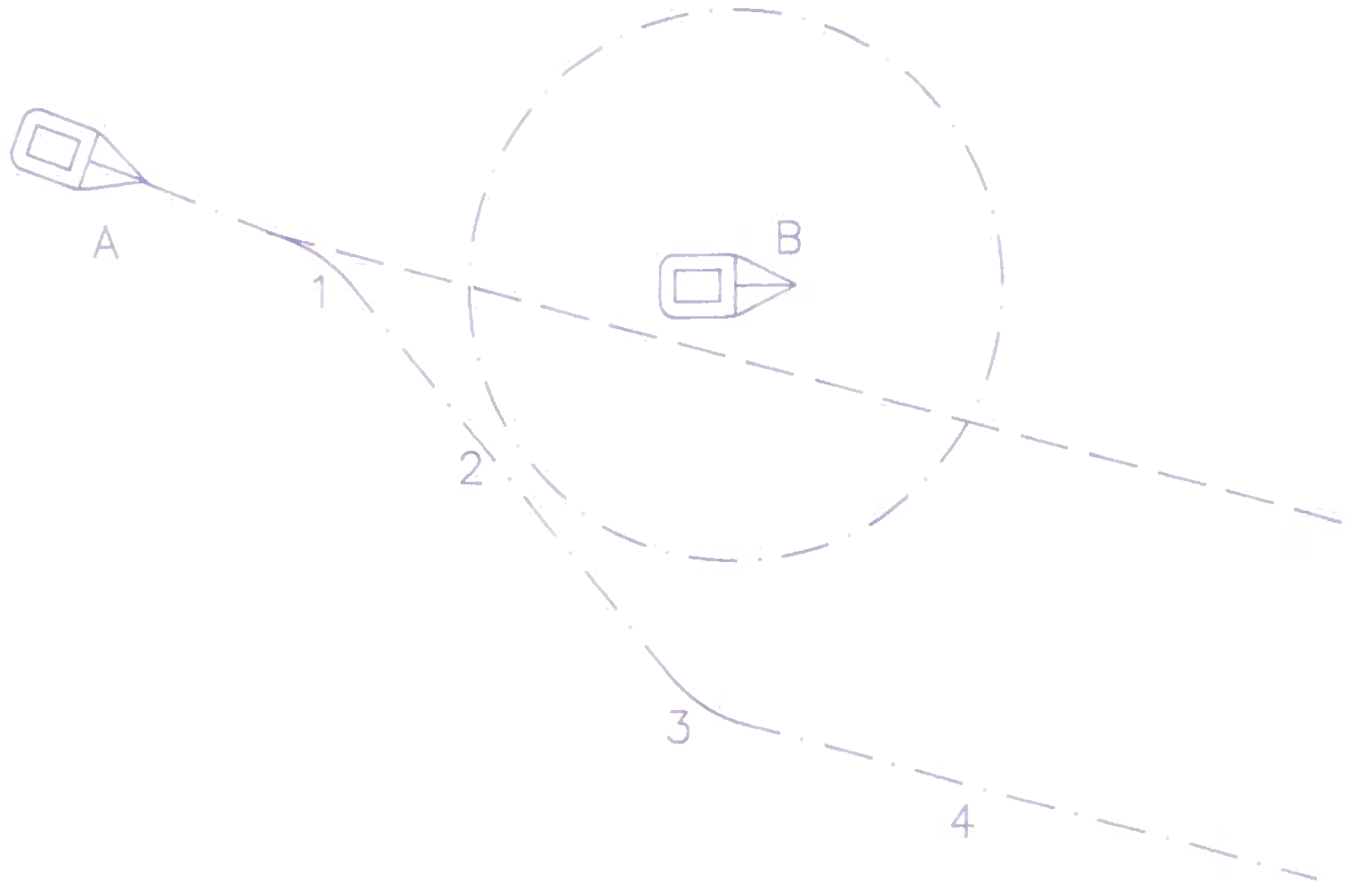

Figure 6.1 The four stages of an overtaking manoeuvre (vessel A overtaking vessel B)

(1) Initial alteration to avoid domain infringement.
(2) Continuation on altered course to clear domain.
(3) Alteration back onto course.
(4) Back on course, continuation parallel to original track.

considered to surround vessel A or B for separation purposes; B has been chosen in this case for convenience).

The point at which the manoeuvre may fail is on the second stage, as ship A passes behind the domain around ship B; if the timing is left too late, then A's course will actually intersect the domain, giving an unacceptable CPA. However, the success of that second stage is itself dictated by the timing of the first stage,in two senses:

(a) for a fixed angle of turn, a delayed commencement of that turn leads to a corresponding delay in completion, leading to a later enactment of the straight-line second phase and consequently greater risk of domain infringement;

(b) as ship A approaches ship B, so the turn required for effective avoidance becomes more extreme; such an increased turn takes a proportionally greater time, again reducing the likelihood of completing that turn in time for phase two to complete successfully.

Curtis [11] put forward a time-based criterion for when action should be taken, the RDRR (range-to-domain/range-rate) or anticipated time to domain infringement on preset courses and speeds. As detailed in chapter 2, a specific chosen RDRR - say 15 minutes - would dictate that avoidance decision-making and action should be initiated 15 minutes before anticipated domain infringement. Such an approach takes account of the closing speed of the two vessels, rather than simply being based on some fairly arbitrary separation distance.

However, the same RDRR value does not suit all encounter situations. Even one specific vessel, with specific turning characteristics, approaching a specific encounter type at a specific speed, may need to vary its time for action according to the relative positions of the vessels involved. Action 15 minutes before projected domain infringement may give a wide margin in one case, just achieve domain clearance in another, and result in a dangerously close CPA in a third.

The three crossing encounters shown in figure 6.2 illustrate this point. In each case, the velocity of own-ship relative to hazard vessel is initially perpendicular ($\pm$ 10 degrees) to hazard's track, as shown. As defined in the regulations, own-ship is the give-way vessel, being in hazard's red (port) light sector, and is required to alter to starboard and pass behind the hazard vessel. In each of the three cases own-ship applies an identical rudder setting 15 minutes before domain infringement (domain being shown around hazard vessel).

In case (a), own-ship is originally set to just clip the domain around the hazard vessel as it passes behind; the manoeuvre 15 minutes prior to infringement gives a wide clearance,increasing CPA to almost double the domain radius.

In case (b), own-ship is initially on a collision course with hazard; altering to starboard at the 15-minute RDRR point gives effective (but not excessive) domain clearance before returning to original course.

(a) Own-ship set to infringe domain to rear of hazard. 15 minute RDRR gives more than adequate clearance.

(b) Own-ship on collision course with hazard vessel. RDRR of 15 minutes gives adequate domain clearance.

(c) Own-ship set to cross ahead of hazard vessel, turns to starboard at 15 minute RDRR. This does not clear domain, but instead substantially reduces clearance between vessels.

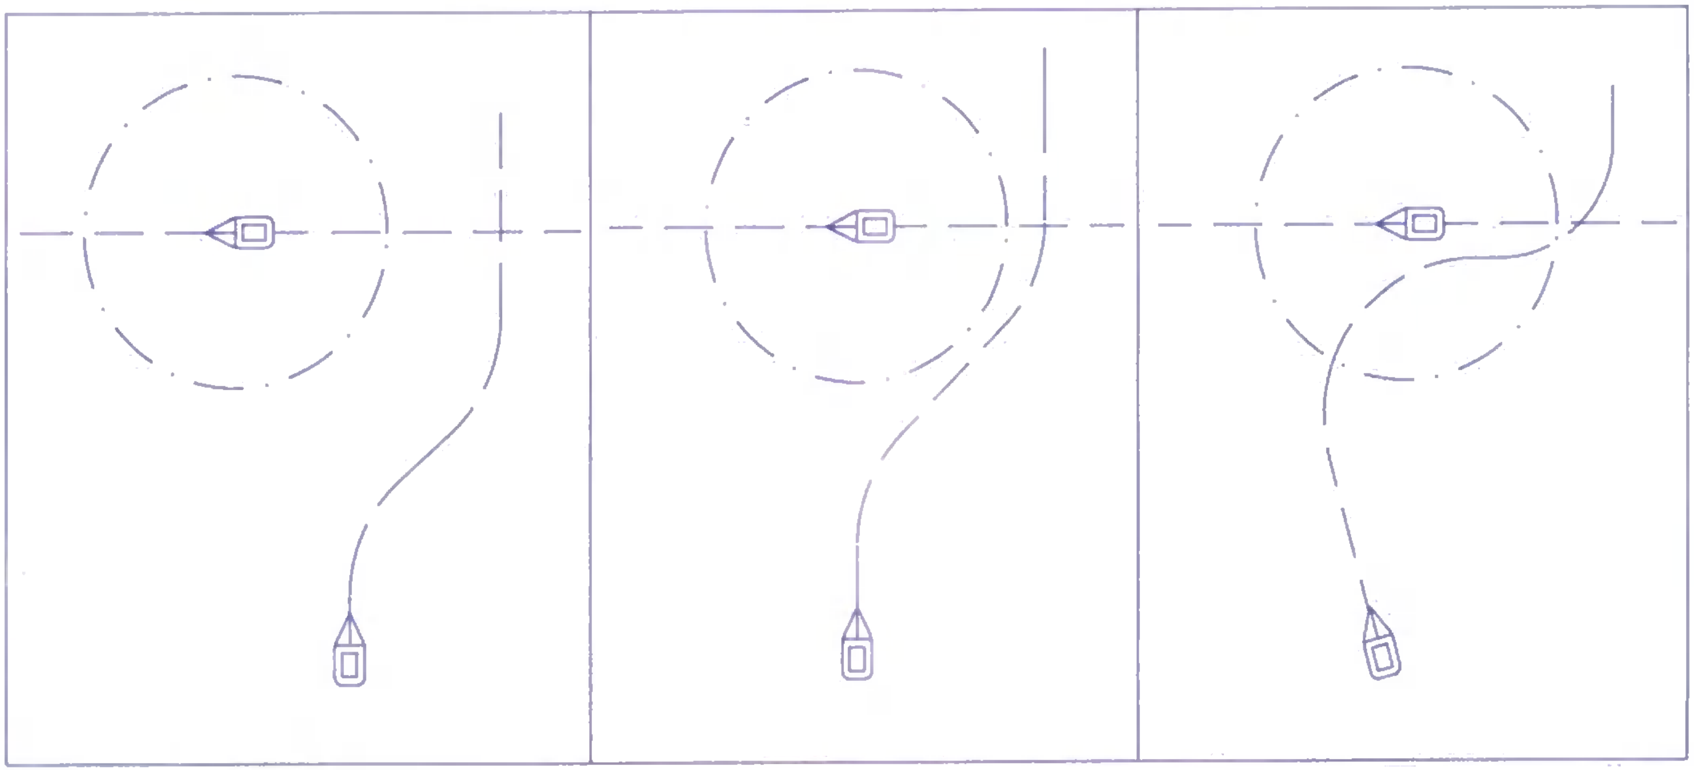

Figure 6.2

Three crossing encounters, each involving own-ship manoeuvring at 10-minute RDRR. The consequences of this action vary widely, depending on position and course of own-ship relative to hazard vessel.

In case (c), the original projected CPA is virtually the same as for case (a) - the difference being that own-ship is set to pass ahead of hazard in (c), as opposed to passing behind hazard in (a). However, the same avoidance manoeuvre leads to a dramatically different outcome, namely a substantial reduction in the clearance between own-ship and hazard; paradoxically, this supposed safety manoeuvre leads to a more dangerous situation, because the manoeuvre was initiated too late to be effective. Figure 6.3 shows how the same manoeuvre, initiated at an earlier time (20 minutes before domain infringement) would be effective. The alternative, an alteration to port to clear hazard's domain and pass ahead (see figure 6.4) is in direct contravention of the regulations and could precipitate a collision (depending on vessel speeds, and speed lost on the turn).

This example considers only three different instances of the same type of encounter. Where different encounter types are considered, further factors come into play. For example, if two vessels are travelling at speeds of 10 and 14 knots respectively, then the closing speed in a head-on encounter is 24 knots, compared with just 4 knots in a near-parallel overtaking encounter. In the former case a 15-minute RDRR would come into effect at a separation of 6 nautical miles (+ domain radius), whereas the latter would be down to a separation of just 1 nautical mile (+ domain radius) before a 15-minute RDRR triggered any decision/action. Though equal in terms of closing time, the dramatic difference in physical separations gives pause for thought. Not least, in the latter case, a minor change in direction of the vessel being overtaken could disproportionately affect the closing time.

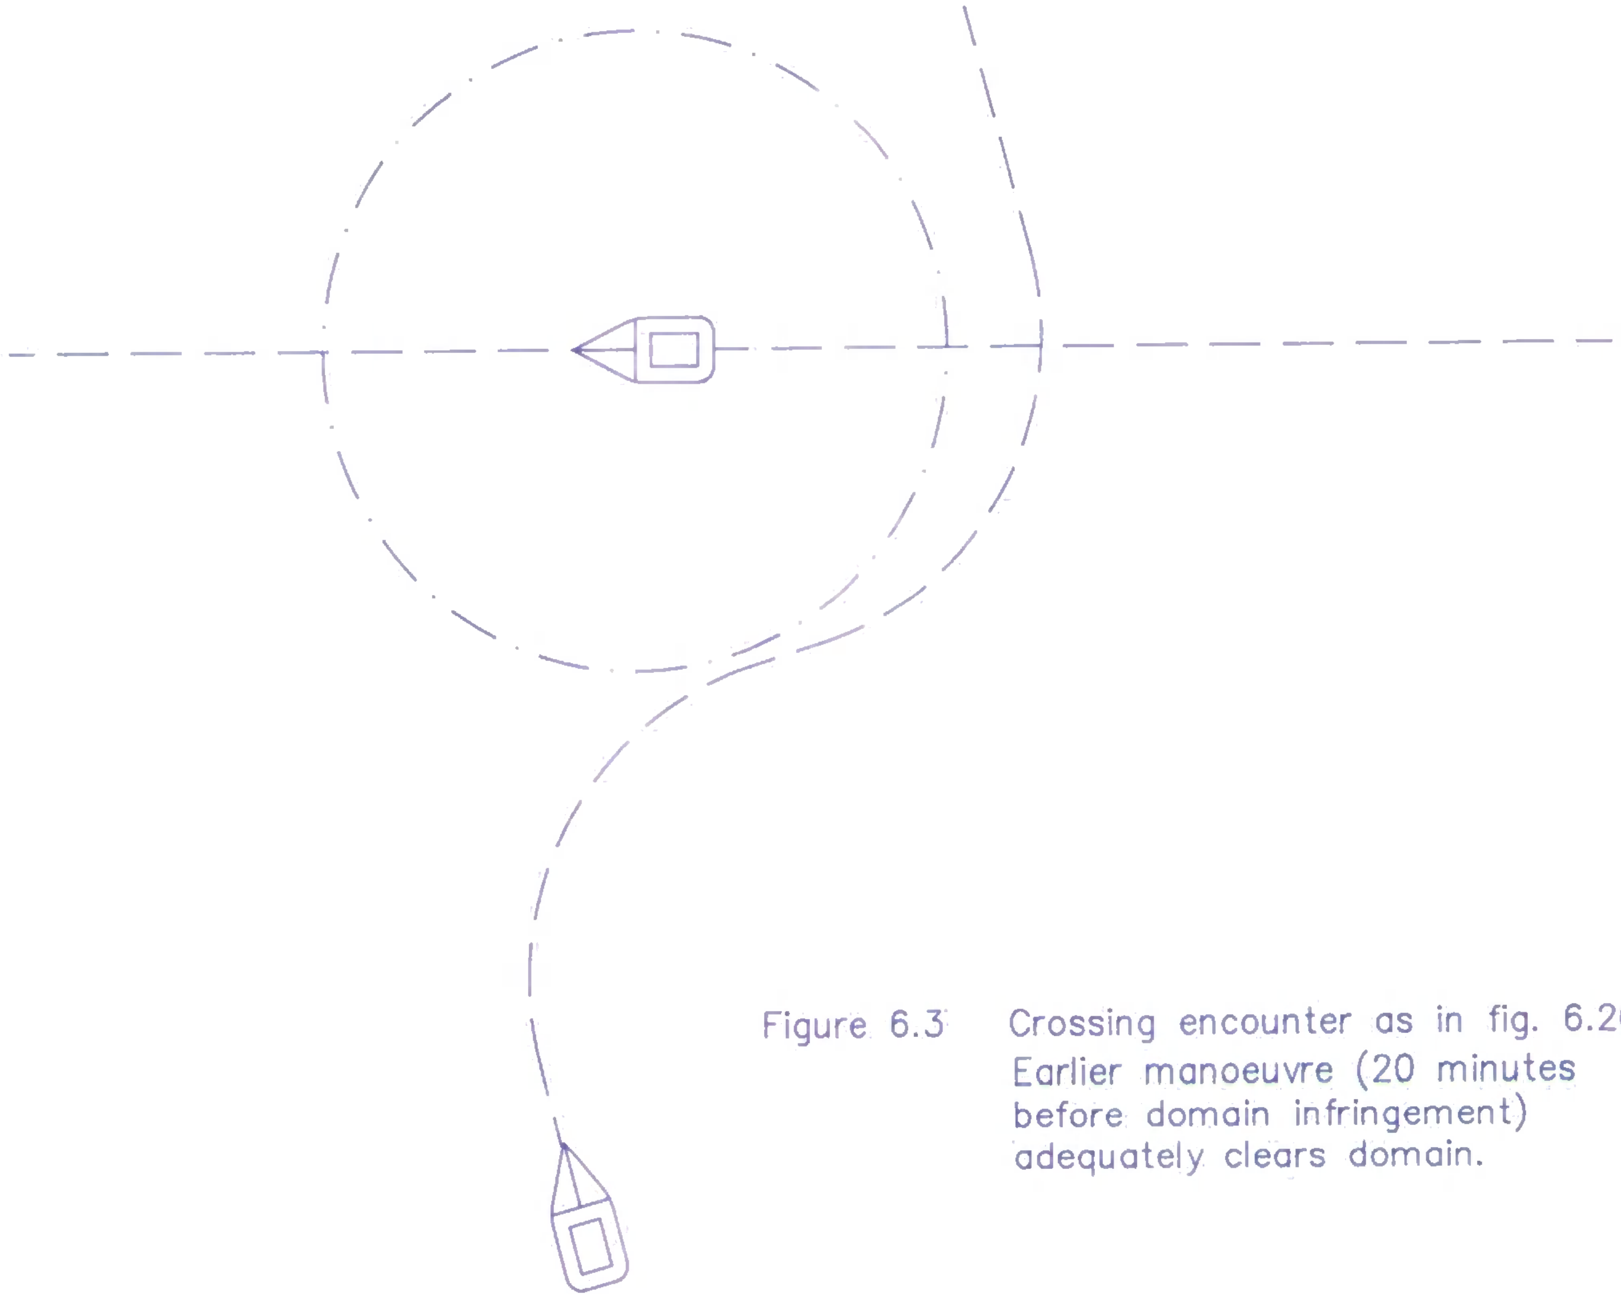

Figure 6.3 Crossing encounter as in fig. 6.2(c). Earlier manoeuvre (20 minutes before domain infringement) adequately clears domain.

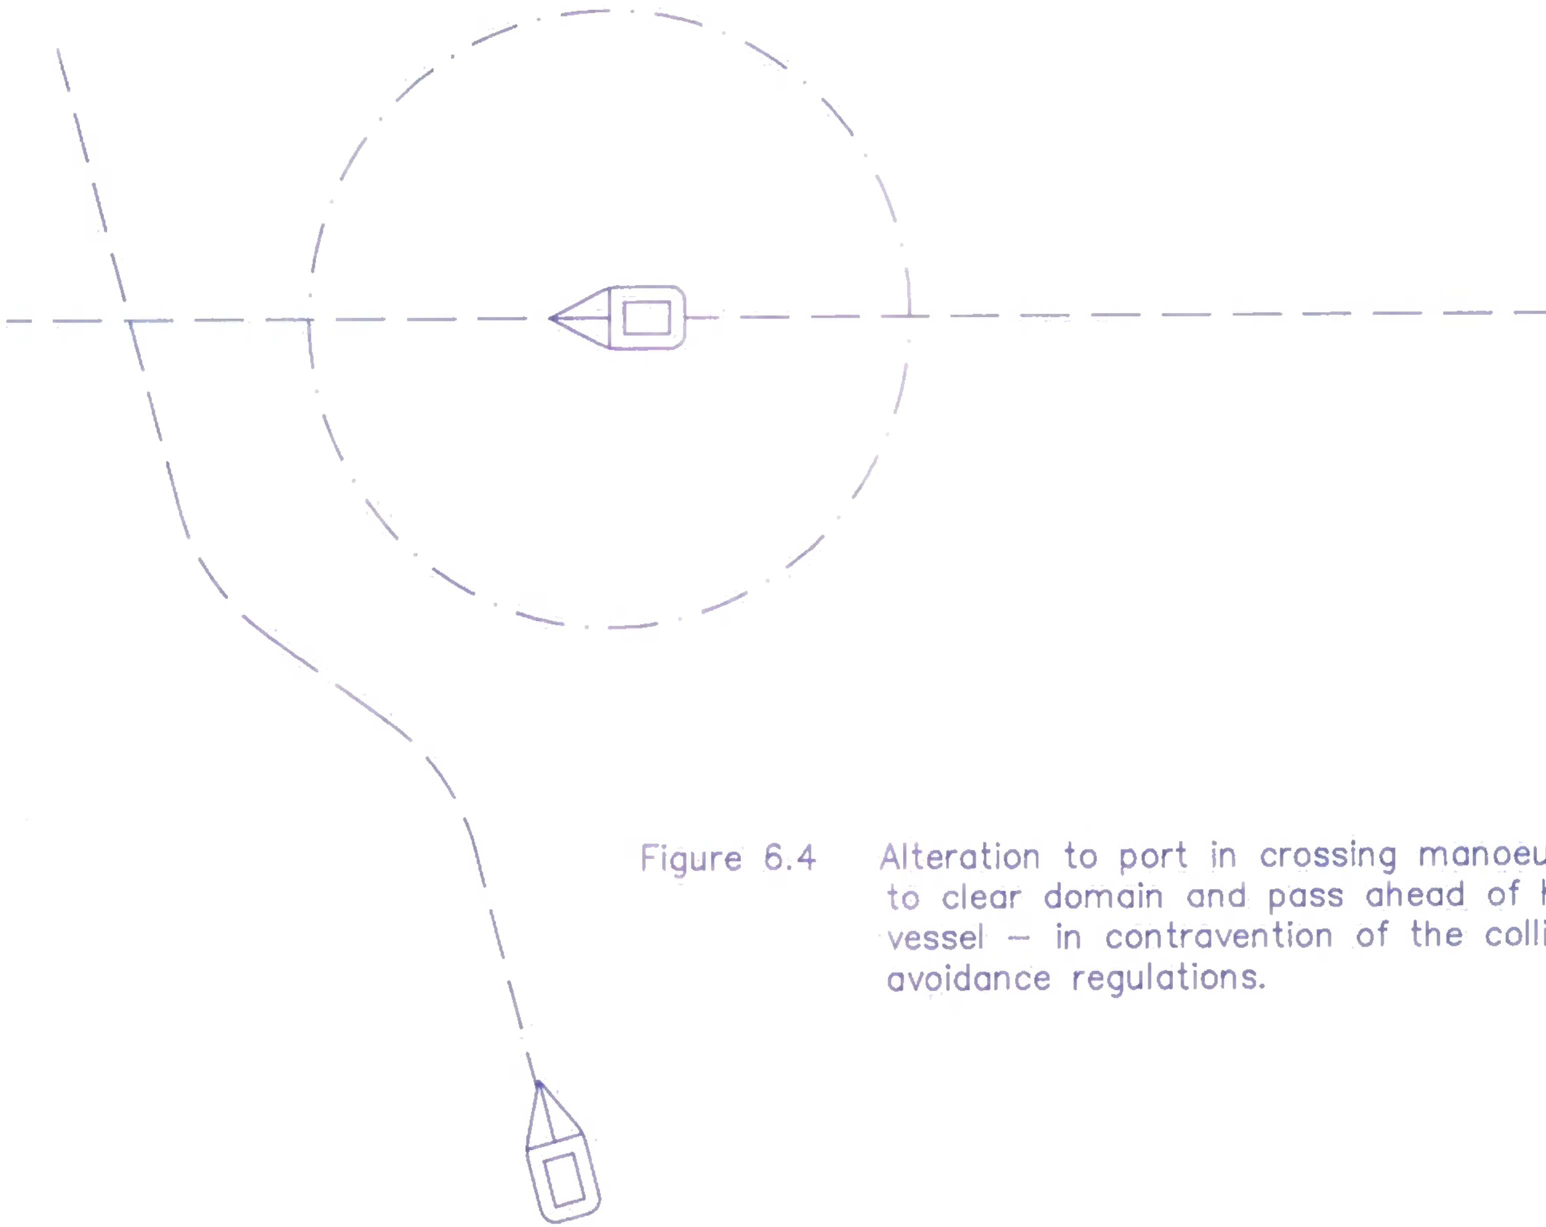

Figure 6.4 Alteration to port in crossing manoeuvre, to clear domain and pass ahead of hazard vessel – in contravention of the collision avoidance regulations.

Konyn [12] has formulated the RDRR+ criterion, which takes cognisance of different types of encounter, giving 'thinking time' and 'action time'. However, this does not really address the question posed by the example in figure 6.2, namely "How does one tailor the appropriate time for avoidance decision/action to the specific circumstances of the encounter?".

Every encounter situation is a unique event, although there is a degree of commonality about types of situations. A possible approach is a trial run-through of a proposed response to the situation, and evaluation of the outcome of this response. Clearly such a trial run must be theoretical rather than practical - such a process (albeit in some cases subconscious) presumably forms a part of the experienced master's decision-making. The alternative of working backward from the end-point seems a non-starter, since the nature of the last phase depends on the outcome of the last-but-one, etc.

In formulating a computer-based solution to the collision avoidance problem, it has proved necessary to address this problem of differential RDRR values for different encounters. An early version of the system, with fixed RDRR, was used on a sample of 500 randomly-generated encounters, as follows:

> For each encounter, two vessels with speeds uniformly distributed between 5 knots and 15 knots were placed randomly, but in close proximity such that domain infringement occurred. The courses of the two vessels were also uniformly distributed, 0-360 degrees.

Each vessel was then 'reversed' along its own track by the same amount (in terms of time), until each was outside the fixed RDRR value with respect to the other. The simulated encounter was then allowed to proceed, having first calculated the projected separation at CPA if no course change was effected; each vessel was given a random turning circle radius in the range 0.2 to 0.7 nautical miles, and both were driven by the expert system logic, so that either or both might conduct avoidance manoeuvres (see chapter 8 for a description of the simulation environment used). The separation at CPA as a consequence of collision avoidance was also noted. The major area of concern was that the system should direct manoeuvres such that all CPA's should fall outside the given domain radius. A secondary consideration was that there should not be excessive deviation from track as a consequence of these manoeuvres (shown by substantially increased separations at CPA).

The statistics for these 500 simulations are shown on the bar chart in figure 6.5. This illustrates that in the vast majority of cases potential domain infringement was avoided without excessive deviation, under the existing rule structure. However, a small number of encounters still have a separation at CPA which is less than the domain radius - the objective of the system was not being achieved in all cases.

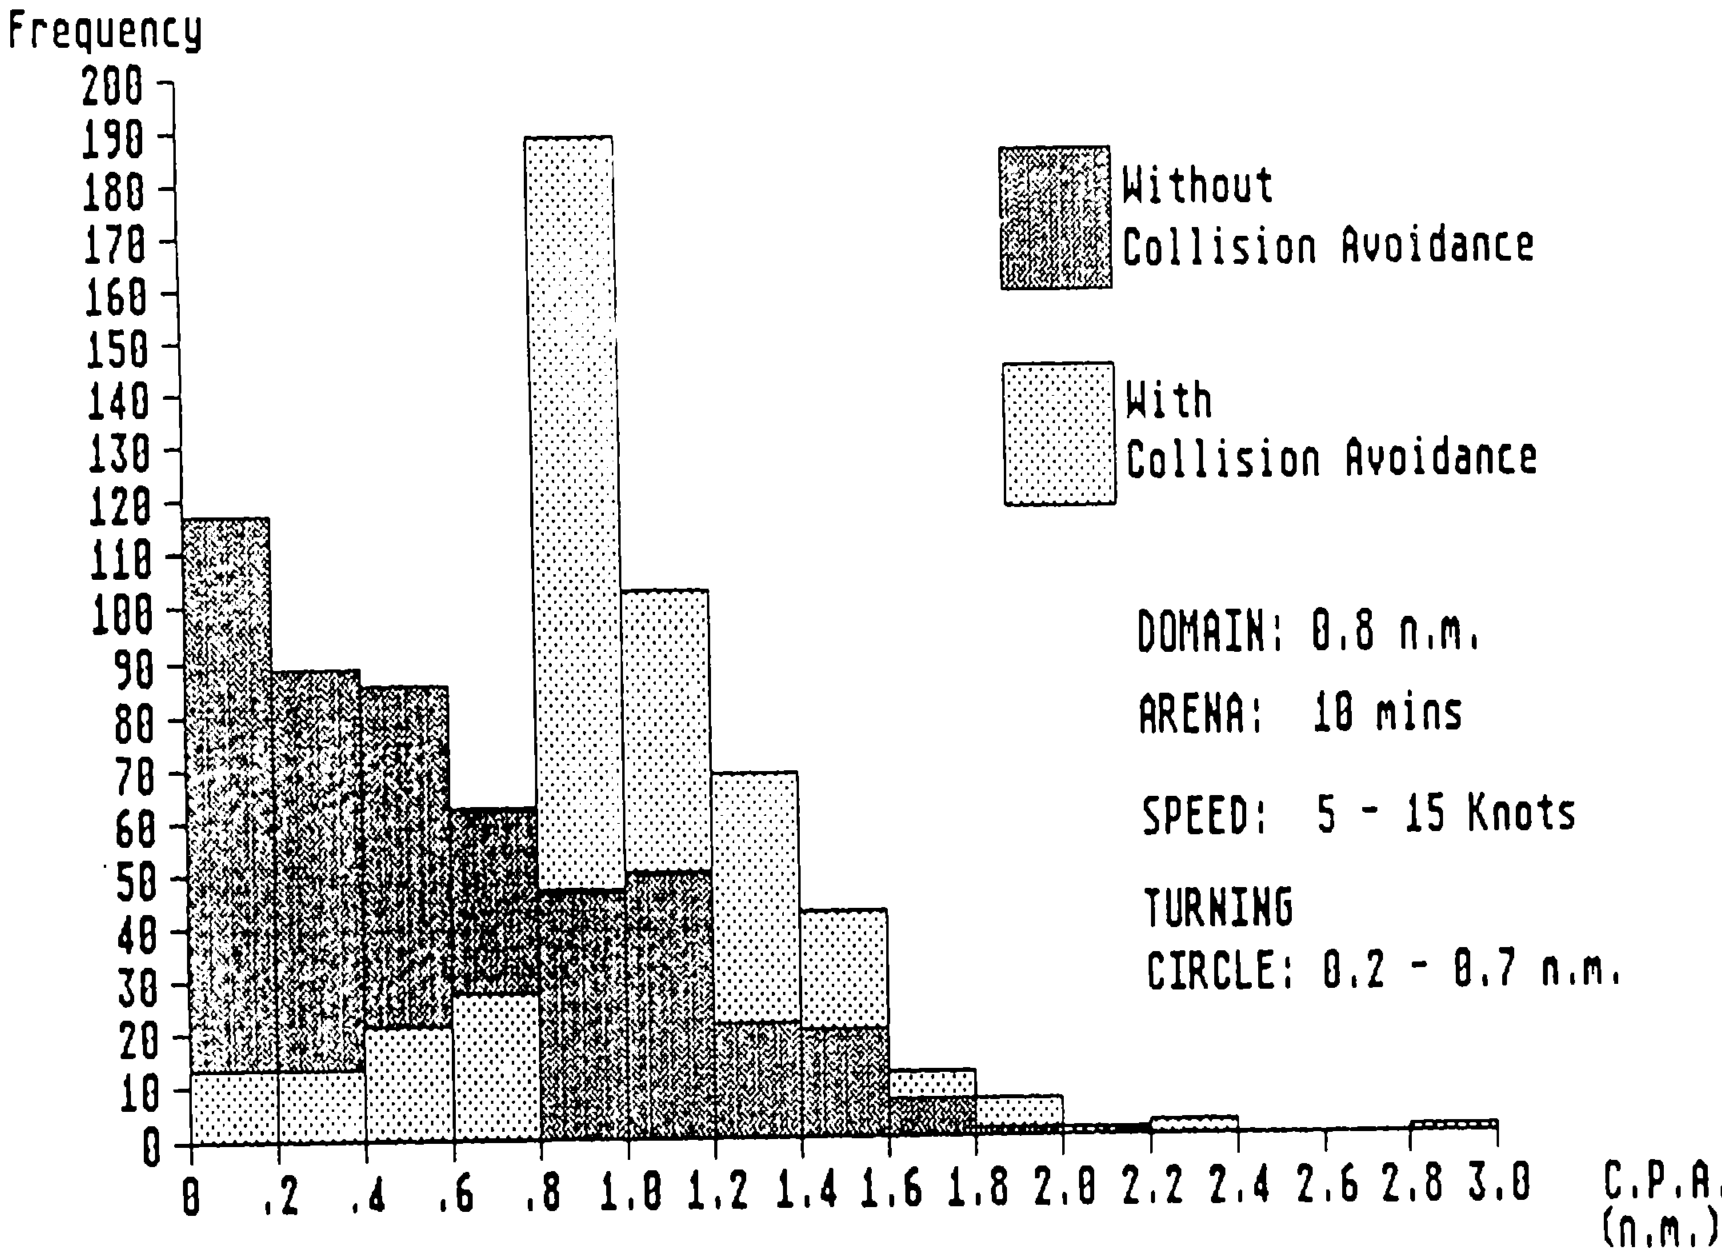

Figure 6.5 Outcome of 500 random encounter simulations on the prototype system, showing effects of collision avoidance strategies.

One possible approach would be to increase the RDRR until all cases cleared satisfactorily. However, this would provoke early, and possibly excessive, avoidance in a large number of cases. Also, as shown in figure 6.2(c), avoidance manoeuvring may in some cases bring vessels closer together before they separate - indicating a need for a substantial increase in RDRR value for such instances. A blanket increase to cover such cases would not appear to be in any sense an efficient strategy.

The radical solution proposed and implemented in this expert system involves that most perceptive of faculties, hindsight. It is quite simple to see, once carried through, whether a particular manoeuvre has proved successful; the disadvantage of such a ploy is that, if it has not (or, less drastically, it has proved inefficient), then it is too late to do anything about it. That criticism no longer holds if the manoeuvre is simply simulated in advance, following through the rules in sequence as they would apply in the actual event. An unsuccessful, or inefficient, outcome may be improved upon by variation of certain parameters for a second or successive simulated manoeuvre. Accurate modelling of ship behaviour, coupled with use of the 'live' rule base for the simulation runs, ensures a realistic picture of the outcome of each manoeuvre.

Just such a forward-looking strategy is used within the expert system to evaluate the logical outcome of a manoeuvre sequence well in advance of the time for that sequence. This is done by simulating the sequence of steps which would be actioned by the rule base, given a specific

RDRR, then assessing the success or failure of that sequence. If the sequence fails, the process may be repeated for a larger RDRR, and if necessary for successively larger RDRR values until a successful manoeuvre is achieved - see illustrated example in figure 6.6. This look-ahead process is carried out by a dedicated module of code within the inference engine, traversing the rule structure exactly as it would in the actual manoeuvre, but on a time-scale which permits evaluation of a complete manoeuvre within a second or so. For this purpose it uses a model of own-ship behaviour as described in section 2.3; the hazard vessel is assumed to continue on a straight course (if hazard is the give-way vessel no standard avoidance action need be simulated, but possible emergency manoeuvring must be allowed for - see section 5.3).

If a hazard is perceived by the system as being on a course which will infringe own-ship's domain, no action is taken until the hazard reaches a time-threshold referred to as the 'ballpark' - currently set at 20 minutes before domain infringement. At this time, the look-ahead simulation module 'advances' both vessels (in simulation) along their tracks to a predefined minimum RDRR value. It then traverses the rule base, following the rules appropriate to the situation at each step, at a greatly reduced time-scale.

Motion of both ships is simulated within the module, including any manoeuvres to be undertaken by own-ship as directed by the expert system (a fixed turning radius is used - section 2.3 deals with this limitation). If at any stage the domain is breached, then the strategy is deemed to have failed; The RDRR is increased by 2 minutes and the

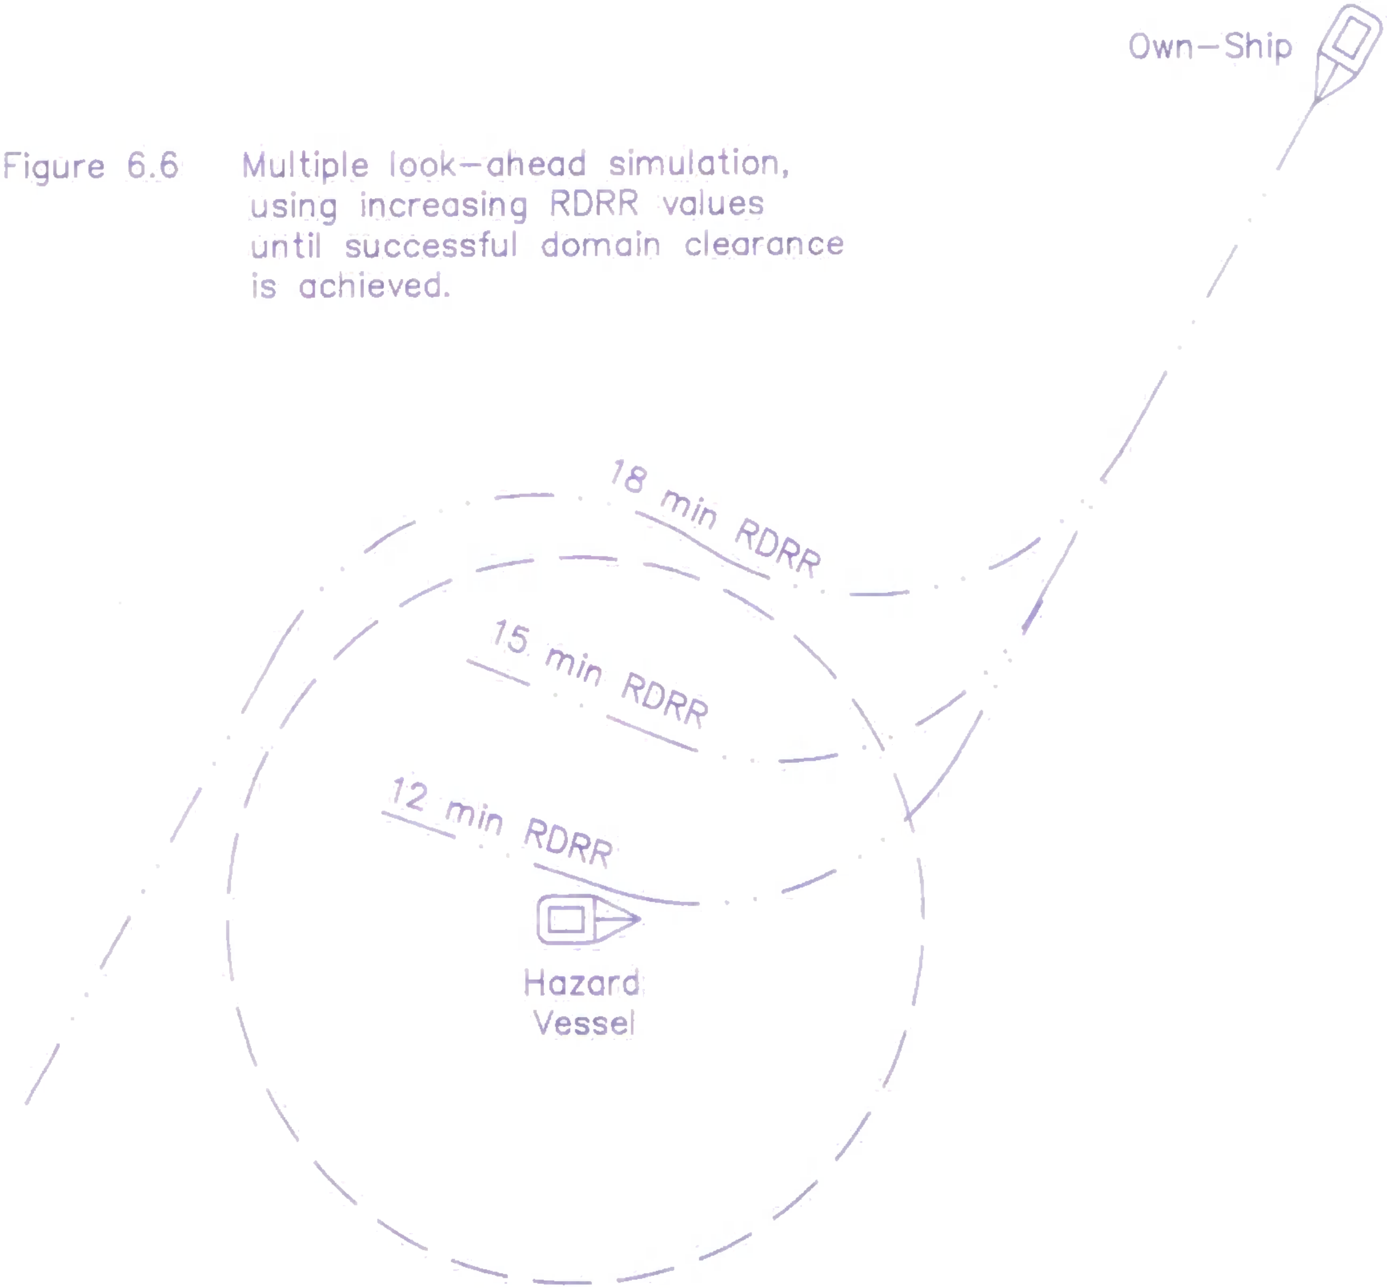

Figure 6.6 Multiple look-ahead simulation, using increasing RDRR values until successful domain clearance is achieved.

process is repeated. This cycle continues until a suitable RDRR is found which achieves domain clearance. This is then registered within the expert system as the effective RDRR for this manoeuvre - also referred to as the PSMT (Predetermined Safe Manoeuvring Time).

If the rule base indicates that own-ship is the stand-on vessel in this encounter, with hazard having responsibility to divert, the look-ahead module will again cycle with increasing values for the RDRR. This will continue until either:

(a) a predefined maximum RDRR is reached, the earliest time at which hazard's master may reasonably have assesses the encounter for avoidance purposes; or

(b) responsibility for the encounter switches from hazard to own-ship (e.g. as a consequence of a situation of the type shown in figure 5.3); n.b. it should be remembered that this process is cycling <u>backwards</u> in time, considering successively earlier times for action.

In case (a) it is concluded that hazard vessel's master will take responsibility for avoidance action - subject to any later indication to the contrary, and consequent need for emergency manoeuvring. This necessitates a further run of the process, this time with reference to the emergency domain, to identify the 'emergency RDRR', i.e. the latest time at which emergency manoeuvring rules may be brought into play if hazard does not avoid, to prevent the emergency domain being breached.

In case (b) own-ship takes responsibility for the avoidance manoeuvre (since this may be the verdict of hazard's master). Identification of the PSMT is then carried out as described, starting from the time when own-ship was reckoned to be responsible for avoidance, and working backward in time (if necessary) until a reliable RDRR - the PSMT - is found.

The success of this look-ahead technique rests upon the fact that the same rule structure is used in the look-ahead simulation as in the practical outworking of the manoeuvre. This ensures that the sequence of steps which are evaluated in the look-ahead process is exactly the sequence used in the actual avoidance process. This guarantees prior knowledge of:

(a) the success of the manoeuvre (barring untoward action by the hazard vessel);

(b) other consequences of the action - distance off track, loss of speed, etc.

Only the former is of relevance in ensuring a safe manoeuvre; considerations of the latter type relate directly to the problem of identification of an optimal strategy.

6.3. Look-Ahead Simulation for Optimal Manoeuvring

The concept of safety is not absolute. In every walk of life, from crossing the road to handling high explosives, any decision on safety is a calculated risk rather than an absolute certainty (in both of the examples cited, the only 100% safe strategy is "don't!"). So it is with marine collision avoidance: in general, the greater the safety margin, the greater the loss of utility. The optimal manoeuvre will be one which incorporates an appropriate level of safety without unduly sacrificing utility.

In any two-ship encounter in open sea conditions, there will be a number of options for the manoeuvre to be undertaken. Even where (as in most cases) the regulations specify a direction for avoidance action, the choice as to timing of that action, and rudder setting, are left to the master of the vessel. This leaves scope for the selection of a course of action which best suits the requirements of the master - and ultimately the operators of the vessel - for that particular situation.

Figure 6.7 illustrates this principle for a crossing encounter, in which ship A is the give-way vessel, approaching ship B on its port side. A is required to alter to starboard and pass behind B, leaving a clear passing distance. This distance is represented by the domain, which for convenience has been shown as a circle around B.

Figure 6.7 Crossing encounter, ship A being the give-way vessel. Six possible alterations to starboard are shown, using three different rudder settings at each of two manoeuvring times. Five of these clear the domain sucessfully, but with different outcomes in terms of diversion from course, resultant distance off track, etc.

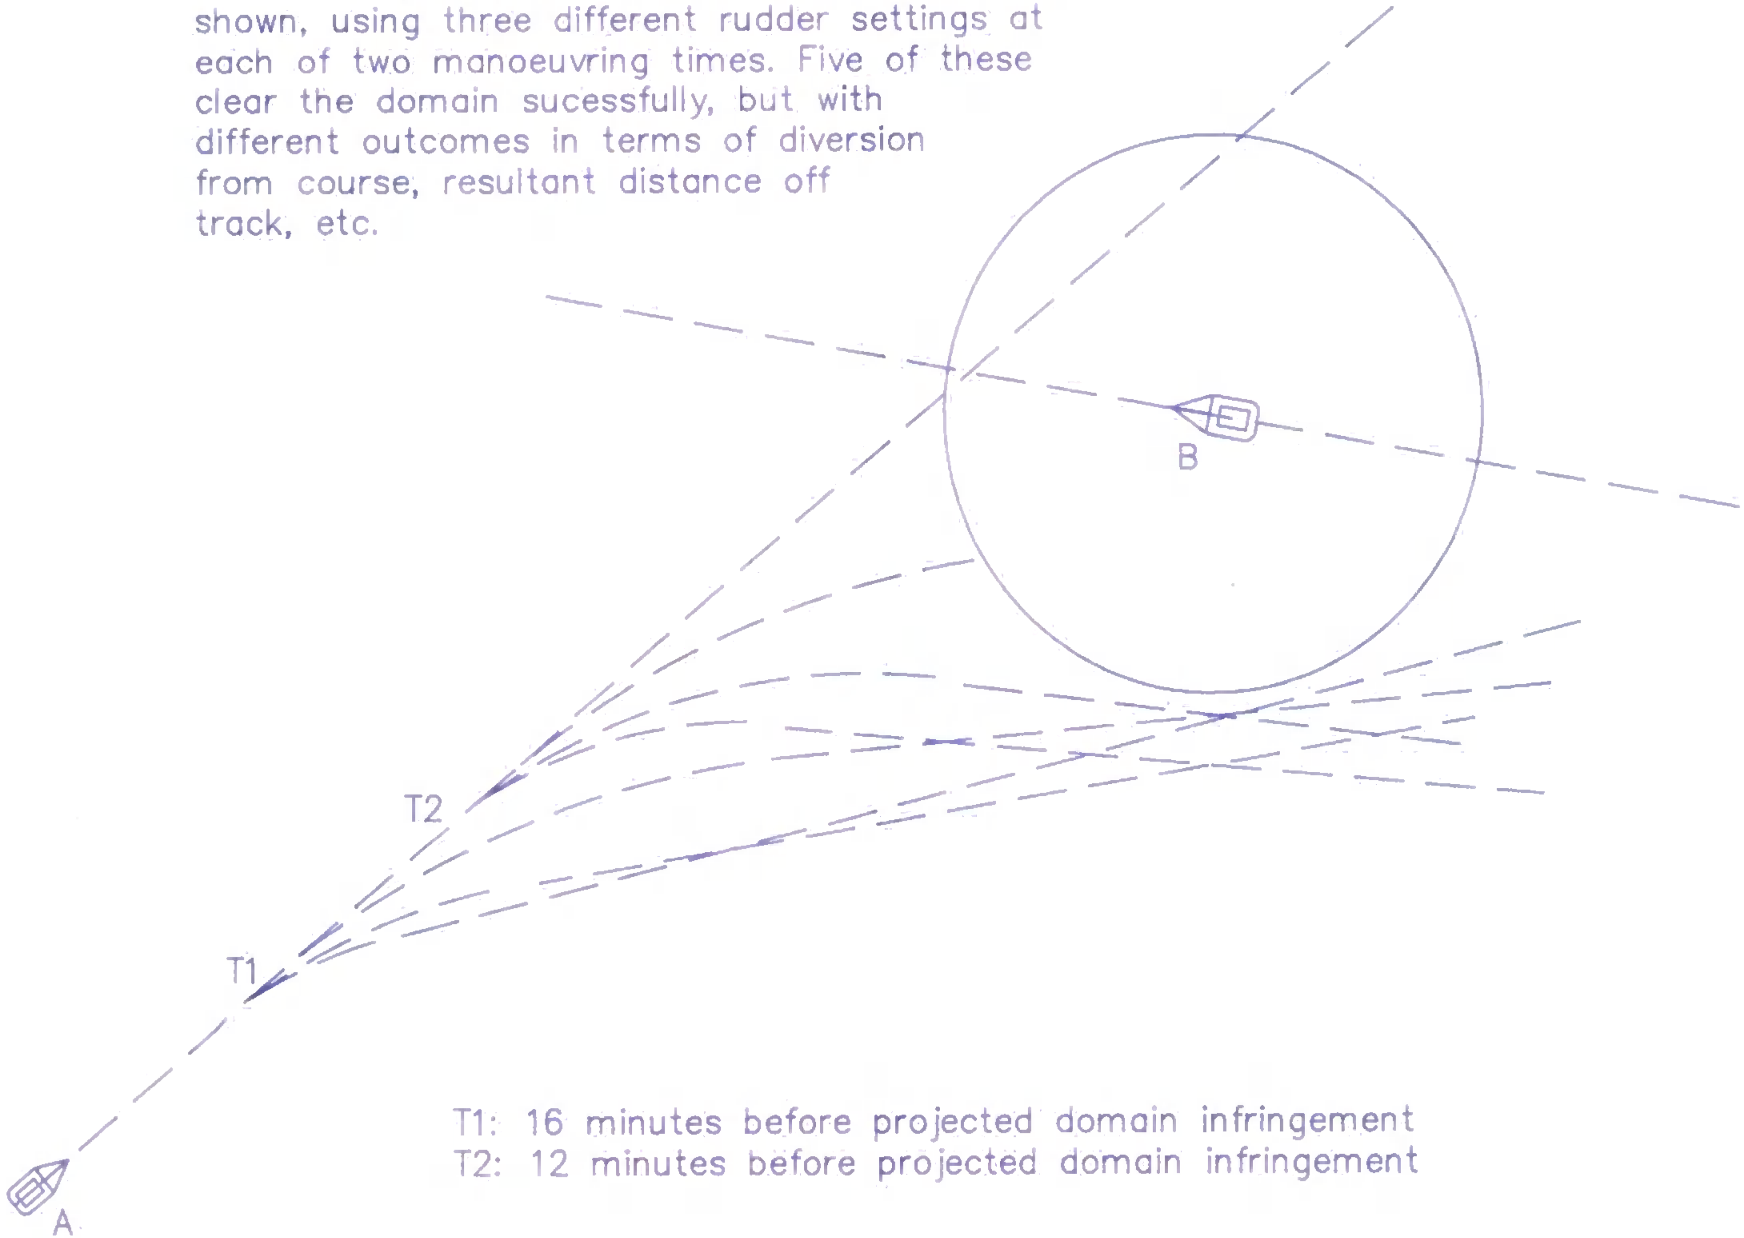

The diagram shows manoeuvres by vessel A at two sample manoeuvring times: 12 minutes and 16 minutes before domain infringement. At each of these times, action on each of three possible rudder settings is illustrated: 1/3 rudder, 2/3 rudder and full rudder. This gives a total of six possible manoeuvring tracks, as shown. The stand-on course for ship A is also illustrated.

The circle around ship B illustrates the position of the domain at the crucial time of projected domain infringement if ship A were to hold its course. The objective of ship A's manoeuvre is to avoid entering that circle, and thus to achieve an acceptable separation at CPA. As can be seen from the diagram, the extent of the turn required is different for each of the potentially succesful strategies.

It will be seen that, on 1/3 rudder, ship A is unable to clear the domain around B by altering 12 minutes before domain infringement; this particular option is therefore not viable. This leaves five viable options out of the six considered, each of which will give different outcomes in terms of: loss of speed; distance diverted off track; time lost in the manoeuvre; excess fuel used; etc. These different criteria may have variable levels of priority depending on the nature of the voyage and type of cargo.

An experienced master will judge such an encounter from some distance off, and determine the best manoeuvre in the circumstances. It would seem appropriate that a computerised expert system should likewise select the manoeuvre which optimises such criteria, rather than just one which avoids the danger. This requires the forward-looking simulation of a number of alternatives (as above) and selection of that one which meets some formal definition of what constitutes an 'optimal manoeuvre'.

That formal definition may be in terms of a 'cost function', which allocates an appropriate weighting to each of the criteria judged to contribute to the 'cost' of a manoeuvre (examples are listed above). The sum of these weighted factors then constitutes an overall cost (in nominal units) of any particular manoeuvre. The costs for a number of manoeuvres may then be compared to determine the optimal manoeuvre in the current circumstances. Clearly, the definition of optimality may be changed by adjusting the weightings applied to each of the cost factors involved.

The look-ahead routine described in 6.2 has been extended to provide exactly this form of discrimination between alternative courses of action. At a time threshold referred to as the 'ballpark' (typically 28 minutes before projected domain infringement), a range of options are simulated and 'costed' as described above. These options include three rudder settings for possible action times in two-minute steps from the minimum RDRR value (specified as a parameter - typically 10 minutes before domain infringement) to 18 minutes before domain infringement. Action is also considered for both the 20 and 25 minute RDRR points; in these cases the manoeuvre is regarded as an early course change, not subject to the regulations, so both port and starboard alterations are considered. It must be stressed that these timings are purely nominal, chosen for system development purposes; they may easily be adjusted if necessary, in the light of further experience and expert advice.

Clearly such a multiple look-ahead and evaluation is substantially more demanding, in terms of CPU usage, than the facility described in 6.2. Consequently a certain degree of 'streamlining' has been incorporated in this process.

Firstly, this routine simulates any manoeuvre only as far as either failure (i.e. domain infringement) or successful avoidance and preparation to turn back onto course (see figure 6.8). Rather than simulating the turn back onto course, a factor is added to the cost function, based on the angular displacement at this stage from the original course; with appropriate weighting, this represents the cost of turning back through that angle. Since all alternatives are judged equally on this 'abbreviated' manoeuvre, a fair comparison may still be made, and the optimal manoeuvre selected.

Secondly, certain sections of each manoeuvre may be calculated analytically, rather than looping through the associated rules numerous times. This must, of course, be done in such a way as to exactly emulate the actions of those rules. Both the initial turn and the straight-line section are susceptible to this approach to some extent, as described in appendix A section 3.

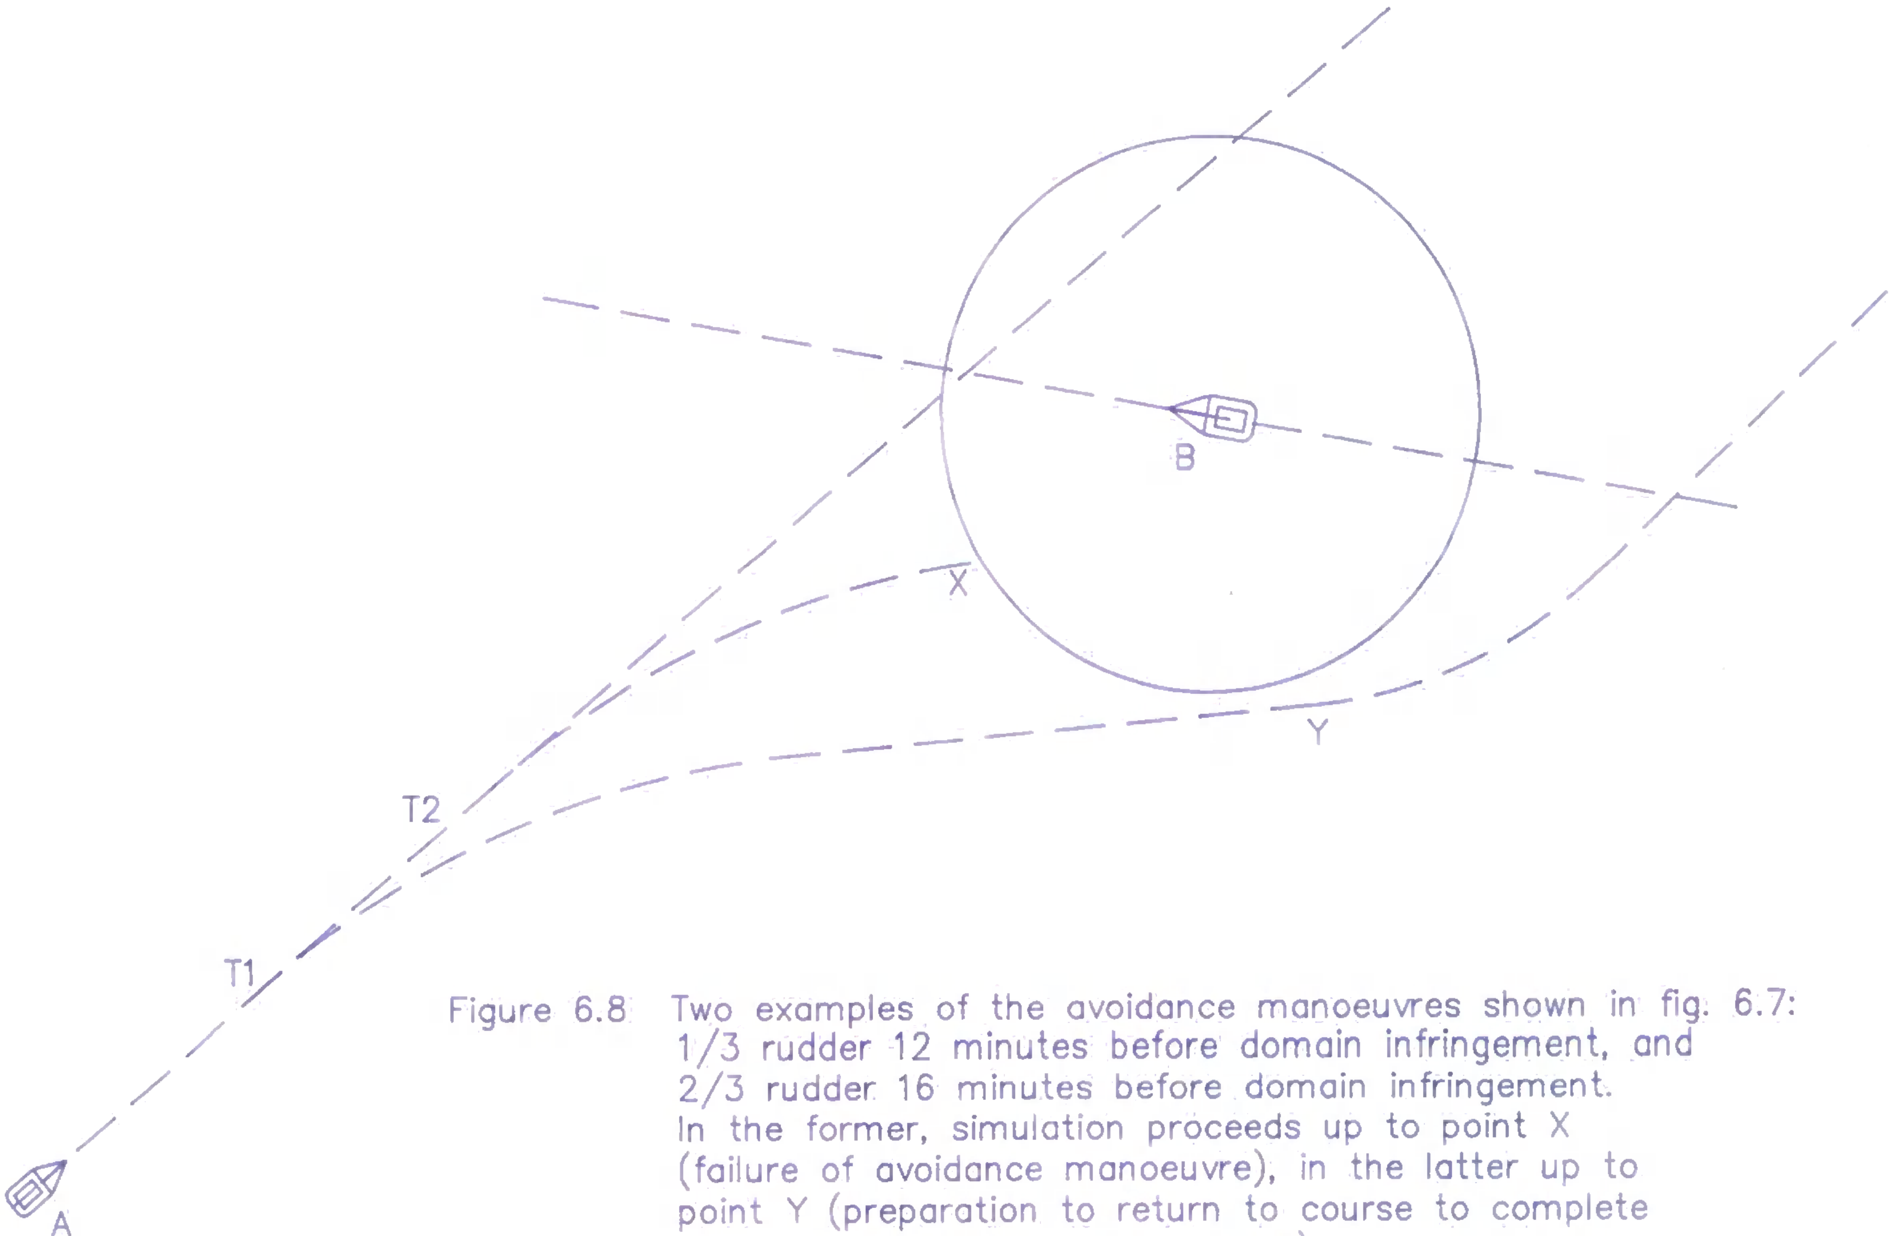

Figure 6.8 Two examples of the avoidance manoeuvres shown in fig. 6.7: 1/3 rudder 12 minutes before domain infringement, and 2/3 rudder 16 minutes before domain infringement. In the former, simulation proceeds up to point X (failure of avoidance manoeuvre), in the latter up to point Y (preparation to return to course to complete a successful avoidance manoeuvre).

CHAPTER 7

THE USER INTERFACE AND USER INTERACTION

"The medium is the message."

Marshall McLuhan (1911-1981) Canadian Sociologist

In a real-time safety-critical interactive system, clarity of communication is of paramount importance. There is little point in generating advice on hazard avoidance, to be acted upon as a matter of high priority, if that advice is embedded in a screenful of text which takes minutes to read and obscures the message by its verbosity.

Life on board ship may vary between two extremes, in terms of the information being received and analysed in respect of prevailing conditions. On exit from, or entrance to, a port, or whilst negotiating a busy seaway, there is a plethora of information arriving at the bridge from a wide variety of sources - radar, radio, navigation systems, communication from other parts of the ship, observations from those on watch. This information must be collated and translated into appropriate action. Such a complex situation can result in some crucial factor being overlooked until it is too late, possibly giving rise to a collision or near-miss (with associated disruptive emergency action).

At the other extreme, there are periods on a sea passage where no hazardous situation presents itself for a long time; no other vessels or charted features give cause for concern, and the vigilance of port entry or exit may be relaxed somewhat. It is at just such times that an unexpected situation may catch the unwary mariner off his guard, and disaster ensue - witness the collision between the trawler 'Dionne Marie' and the stationary tanker 'Rose Bay' on Cup Final Day (!) 1990.

An electronic aid to safe navigation must be capable of assimilating all pertinent information in either of these extreme situations (and all intermediate cases) and extracting significant detail. It should then provide a concise summary of the relevant facts and conclusions, in an easily-digestible form. Graphics and colour, as well as split-screen techniques, may be usefully employed to enhance the level and speed of communication. A balance must be struck between paucity of information and excessive clutter on-screen; simple controls for user selection of optional features offer greater flexibility in display formats.

The display requirements detailed above all point to an advanced WIMPS (Windows-Icons-Mouse-Pointer System) environment, such as that available on the Acorn Archimedes microcomputer. A number of windows (sub-screens onto different aspects of the process) may be open on the screen simultaneously, displaying text and/or graphics; any window may optionally incorporate controls for re-sizing, and for horizontal and/or vertical scrolling. Colour and graphics effects may be used for visual impact.

Pop-up menus may be invoked by use of mouse buttons, offering user options (e.g. re-scaling of graphics display) also selectable by mouse buttons. Icons (representative symbols) may be displayed on screen for mouse-and-pointer activation or termination of processes. 'Dialog boxes' may be designed to display informative text under certain conditions, usually with user response; this may be by a mouse-click on a response box (possibly selection of one of several) or by text input. Such dialog boxes simplify the two-way communication between system and user.

Much consideration has been given to the ergonomics of the system under discussion here. It would seem appropriate to display some graphical representation of the current scenario, including circumstances leading up to that situation; this would allow the user to visualise possible future developments. An up-to-the-moment summary of the salient features of an encounter - range and bearing of hazard, time to projected domain infringement, course and speed of both own-ship and hazard - would also aid in objective assessment of the situation. Not least, such an advisory system must advise, preferably in clear, succinct statements; a rationale for the advice given should also be available upon request.

Figure 7.1 shows how the screen display for the expert system is split into three windows, meeting the needs outlined above. The following sections give detail of the content of each window.

7.1. The 'Tracks' Window

This square window displays graphically the tracks of own-ship and any hazard in the vicinity (figure 7.1). The track of own-ship is shown in green, that of the hazard in orange. In further multi-ship developments [7], various other colours are used for additional hazard vessels. Each

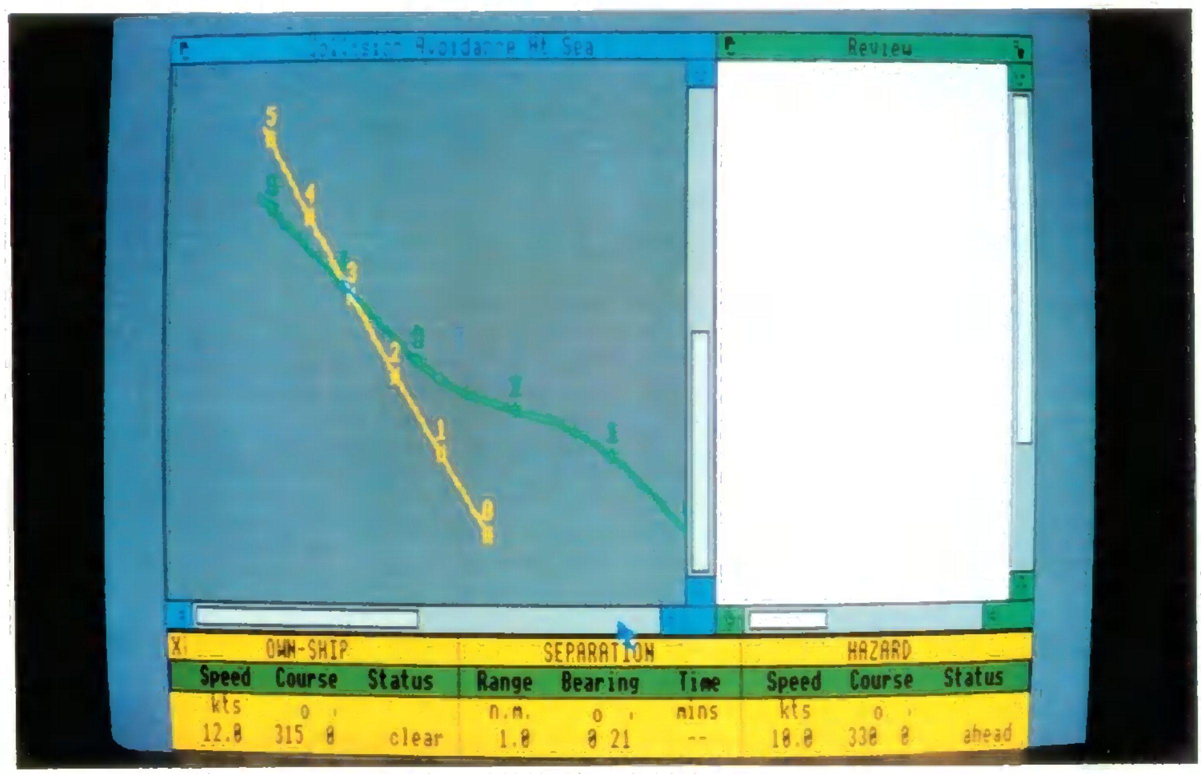

Figure 7.1 Screen display for expert system, showing windows for:

(a) tracks of vessels involved (upper left);

(b) optional review of encounter so far (upper right);

(c) continuously-updated status information (bottom).

track is marked at 8-minute intervals (real time - see chapter 8 on simulated encounters) with a symbol and a 'time-stamp', 0-9. Thus corresponding points in time can be compared on the tracks of different vessels. It is also possible to optionally add lettered markers to the track of own-ship, to identify critical points in the development of the encounter as related in the 'Review' window (see section 7.3).

Scroll bars on the base and right-hand side of this window permit movement of the displayed portion around a sea area four times that size (at normal scaling). Alternatively, the scale of the display may be reduced by a factor of two, enabling the whole of that larger sea area to be viewed within the window at one time. A one-mile grid may optionally be superimposed upon the display, to give a feel of scale. These features are illustrated in figures 7.2, 7.3 and 7.4.

The current format of this display window has proved adequate for the development work undertaken so far, but further enhancement of this display technique is planned to cover the needs of longer voyages. As it stands, the display tracks the passage of of the vessels involved across the window for the duration of the situation being considered - usually a potential collision or near-miss encounter. If left for a longer period - for example, a complete voyage - then vessels initially shown on the display (including own-ship) will ultimately disappear off one or another edge of the display window. Clearly, it is a prerequisite that own-ship should at all times be indicated on this display, preferably in a fairly central position. This requirement, and various approaches which may be taken to satisfy it, are considered in section 11.2.

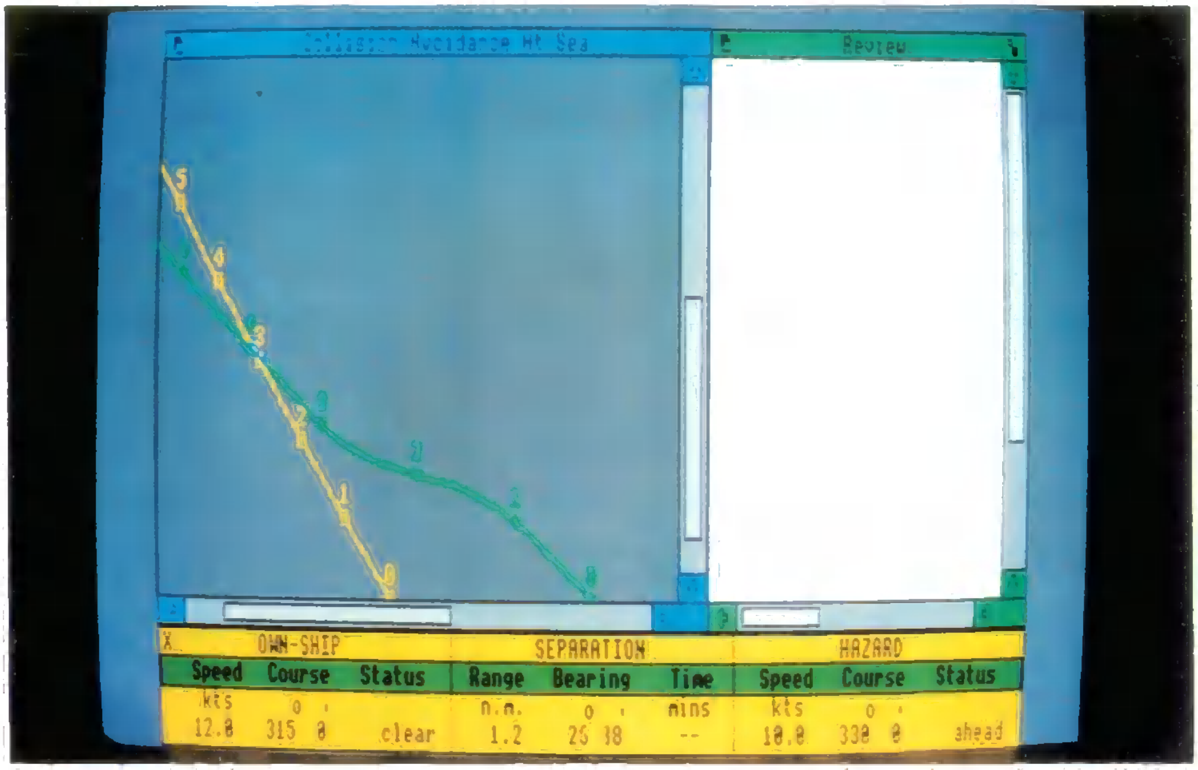

Figure 7.2 Screen display, showing use of scroll bars on 'tracks' window to move display up/down and to left/right.

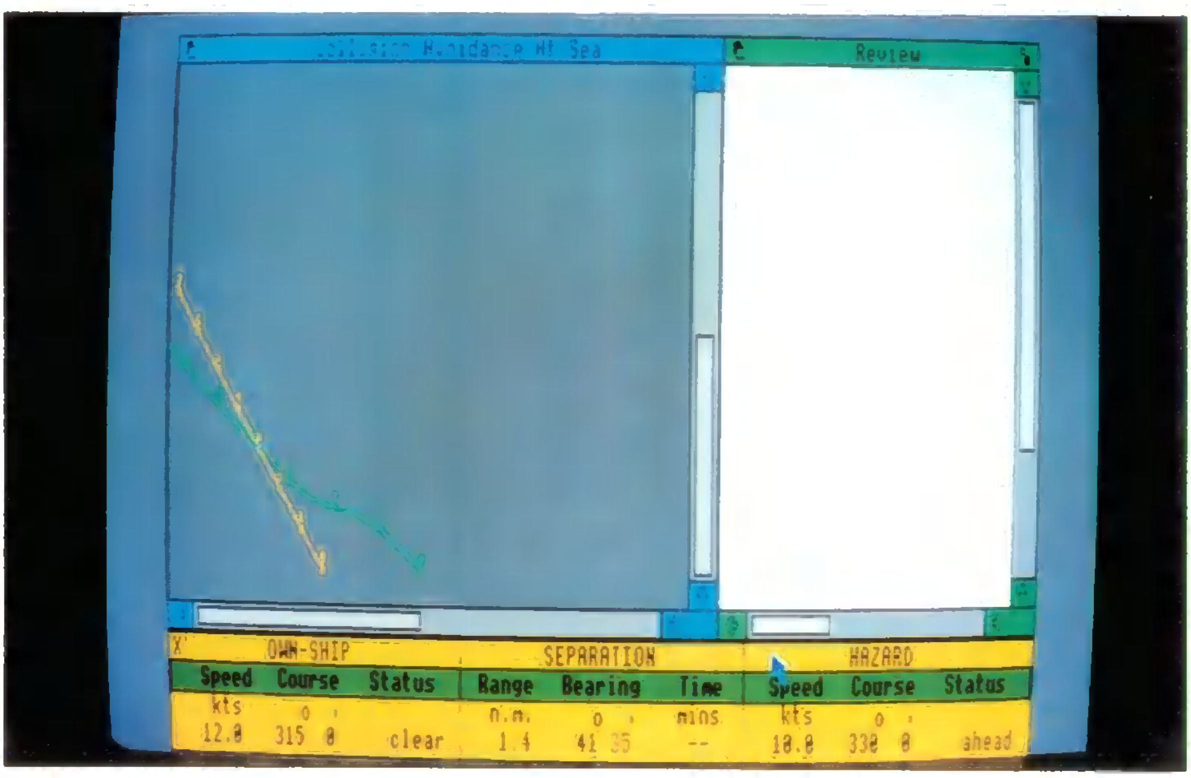

Figure 7.3 Screen display with scale of graphics window reduced by a factor of two to show full 'tracks' area.

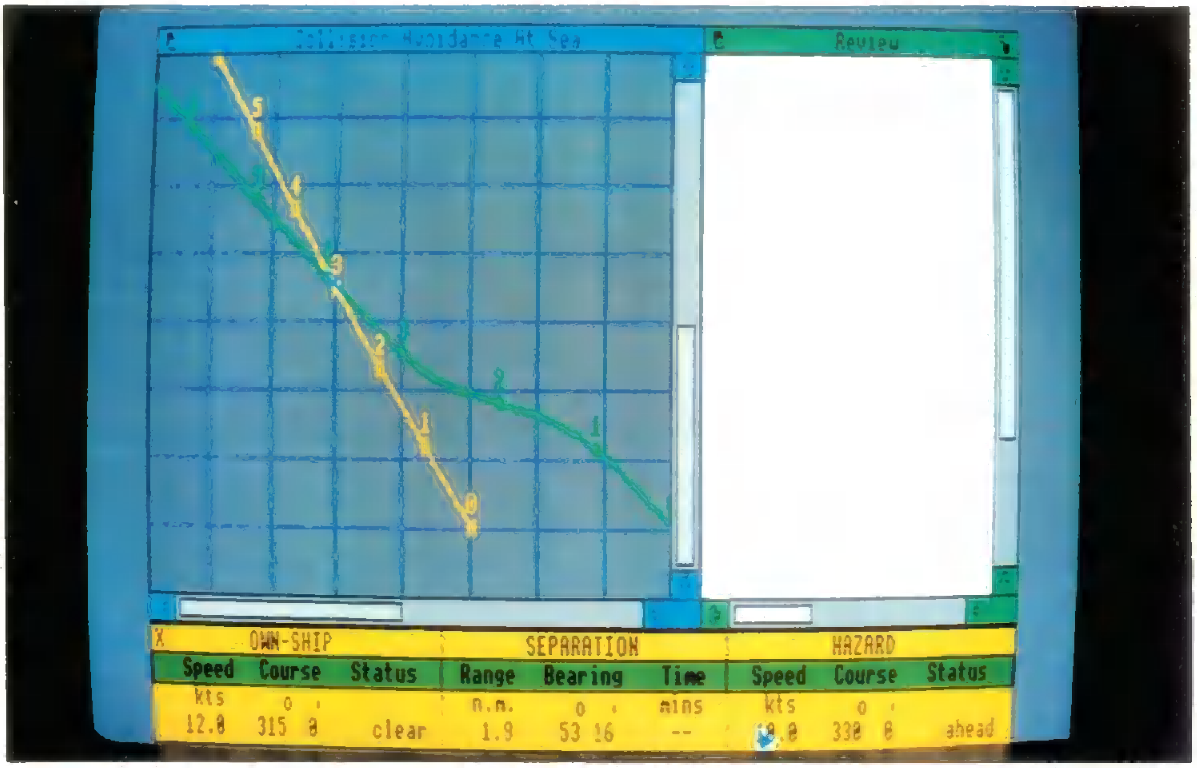

(a) Large scale graphics display

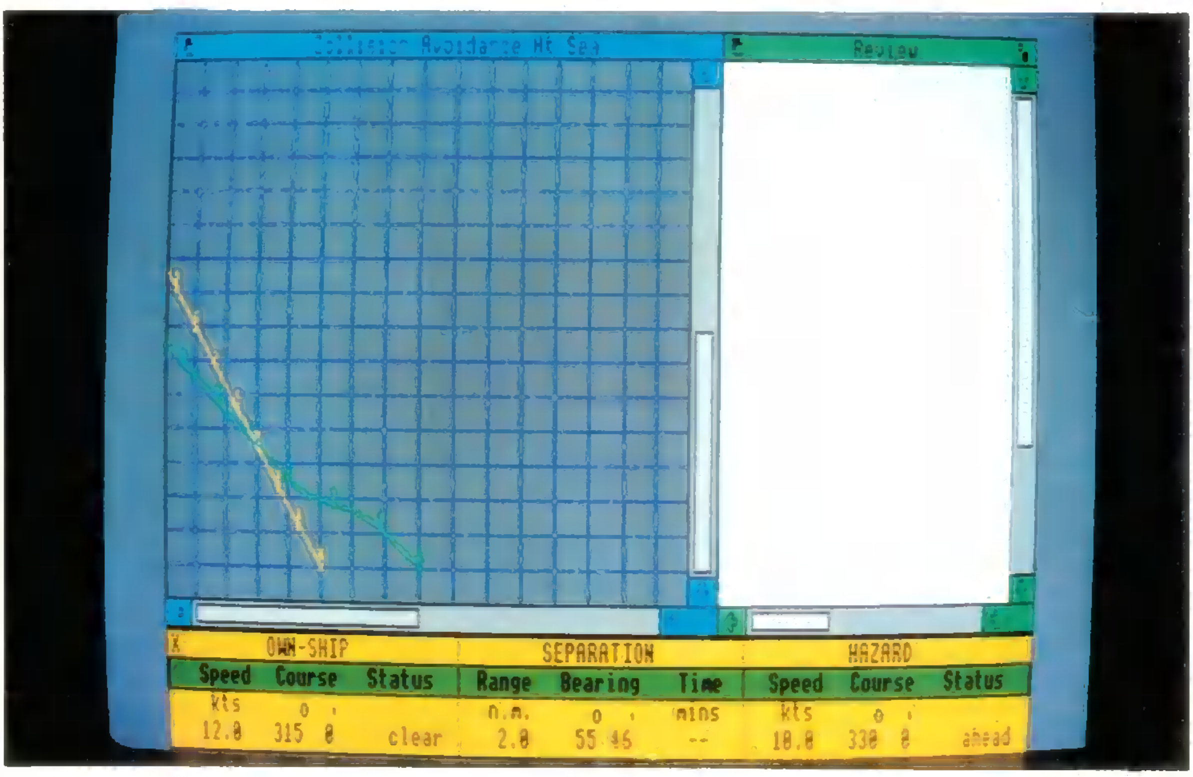

(b) Small scale graphics display

Figure 7.4 Screen display showing 1 n.m. grid on 'tracks' window.

7.2. The 'Status' Window

This window is optional, being selected from a pop-up menu activated by clicking on the expert system icon (see section 7.4); it may be closed at any time by activating the 'close' box provided (top left of window). The window is situated along the bottom of the screen, covering the 'icon bar' which holds the icons for all currently-loaded processes. Consequently it is necessary to close the window (temporarily, at least) if access to any of those icons is required, e.g. to reach an icon menu.

The status window is updated at every time-step (20 seconds, real time), giving a summary of the current state, in text and numerical form. Own-ship's speed and course are shown, as is the corresponding data for the hazard vessel (the most threatening hazard, in later developments). The current status of own-ship with regard to any encounter situation (held as part of each 'rule' node) is also shown - 'clear', 'cross-A' (first stage of crossing encounter) etc; observed state of the hazard vessel - 'ahead' or 'turning' - is likewise displayed. In addition, data relating the hazard to own-ship is shown - range, bearing and time to projected domain infringement on current courses.

The colours for this window have been selected to give clear presentation of detail without distracting the user's eye from other parts of the display. The amber background is generally recognised as the colour least conducive to eyestrain, hence the use of amber screens in many commercial and industrial applications (notably radar screens).

The dark green text shows clearly on this background, again without drawing attention away from other information on the screen.

The use of such muted colours contrasts with the intentionally bright background and contrasting foreground colours in the 'tracks' graphics display window. This has the effect of directing the user's attention to the latter, for quick assimilation of the current scenario; supporting textual and numeric detail may then be referred to by a conscious switch of attention.

7.3. The 'Review' Window

A common feature of rule-based intelligent systems is their capacity to provide a rationale for the decisions taken or recommendations made. This usually takes the form of inferential logic based on current information in the knowledge base:

> "A is so, therefore B is so."
> "Since B is so, and C is not so, it follows that D is so."
> "Either D or E is so, consequently F must be so."
> "Having decided that F is so, and knowing that G is so and H is so, the logical conclusion is _ _ _ _ _ _".

The rule-based system under consideration here is somewhat different conceptually, since present decisions/recommendations are related to a time sequence of events, originating some time in the past and culminating at some point in the future:

> "Hazard vessel has moved from that point to this point in the past 10 minutes. Continuing on that course, it will infringe our domain 28 minutes from now."
>
> "If we evaluate the encounter situation 12 minutes before domain infringement (i.e. 16 minutes from now), we will prove to be the give-way vessel in a crossing encounter."
>
> "If we apply half-rudder to starboard 16 minutes from now, our clearance of the other vessel will be greater than the domain radius. This is therefore a safe manoeuvre".

Any review of the decision-making process <u>must</u> therefore incorporate a time dimension. More than that, these time-stages should be somehow linked to the graphical information in the 'tracks' window, so as to illustrate both the causes and the effects of the decisions.

The 'review' window covers approximately one-third of the screen, on the right above the 'status' window. Technically it is a 'text object' in RISC-OS parlance, i.e. a window to hold text which may be added to at any time. The window itself incorporates controls for scrolling horizontally and vertically, also for expanding the size of the window.

As the encounter proceeds, each key position on own-ship's track is noted (but not yet marked); the text for the reasoning and decision associated with that point in time is added to the stored review text (but again, not yet displayed). Each such text item is a succinct statement of the pertinent fact(s) and the inference drawn.

At any time, the user may ask for the review text to be displayed (see 7.4 - menus). The text so far accumulated is shown in the text window, successive text items being labelled with successive capital letters, one for each key point. The track of own-ship, in the 'tracks' window, is also marked with corresponding letters, indicating those points at which decisions/conclusions were made (see figure 7.5). The review display, with track lettering, may be updated at any time by repeated menu selection of the 'review' option.

As previously indicated, decisions by this expert system are time-based, being subject to time constraints as well as physical constraints. It follows that the review text must include a temporal context, as well as explanation of spatial considerations in the 2-dimensional frame of reference. This must be done in such a way as to facilitate understanding and give a meaningful summary which may be quickly and easily assimilated; at the same time, the user must be given the opportunity to examine the logical decision process in detail, if so required.

The present implementation provides a brief, succinct statement for each key step in the decision process. It is intended, at a later date, to provide optional elaboration of each of those steps; the format of the review window and text item have been designed with this enhancement in mind. Details of this planned development are given in section 11.3.

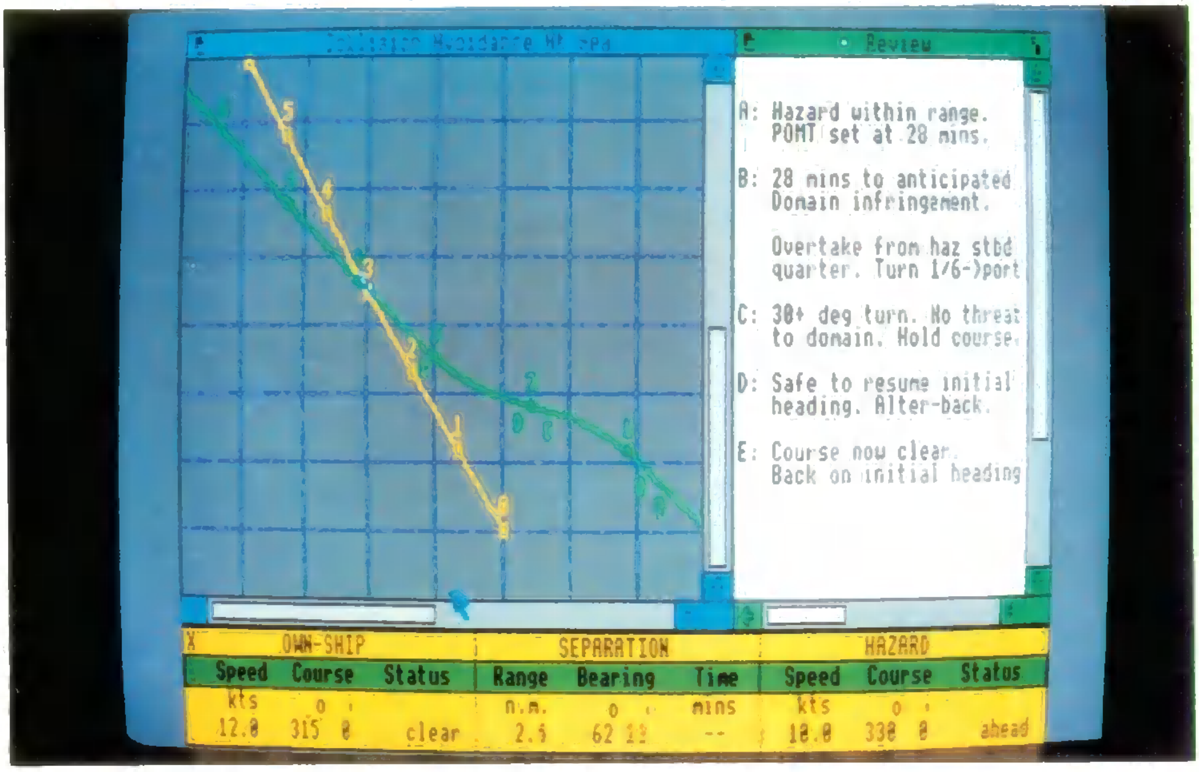

Figure 7.5 Screen display, showing review information displayed in upper right window. Own-ship's track (graphics window, upper left) is marked with lettering corresponding to review text lettering, to indicate decision/action points.

Each item of review text is held within the corresponding rule node (see section 3.2) for which that text is appropriate. Not every rule carries a section of review text; likewise, only certain review texts are 'flagged' to have sequential lettering attached to them.

Some of the review texts carry a 'wildcard' symbol (#) followed by an identifier (e.g. #P). This combination is replaced at action time by the current value of an item of dynamic data within the system. For example, the text:

PSMT set at #P minutes

may be shown at run-time as:

PSMT set at 16 minutes

if 16 is the value determined by the look-ahead simulation process.

As yet, this facility is used only sparsely, but a look-up table at the end of the 'rules' file makes extension of the symbol list a very simple process, when required in future extension of the rule structure.

7.4. Menus and Dialog Boxes

In addition to the three facets of the screen display detailed above, there is at times a requirement for user interaction, when the user has the opportunity to provide information pertinent to the functioning of the system. This may be mandatory, where the expert system prompts the user for required information; alternatively it may be optional, where the user elects to provide new system parameters or select alternative options.

The WIMPS environment provides two means of such communication, namely pop-up menus and dialog boxes.

7.4.1. Pop-Up Menus

A menu is activated by positioning the pointer over a specific icon, or within a specific window, and then pressing the 'menu' button on the mouse (the middle button, in the case of the Archimedes). The current expert system makes use of two such menus, which between them illustrate the uses of this facility.

One menu is attached to the expert system task icon, the green ship on the icon bar at the bottom of the screen (sometimes hidden by the 'status' window). This menu provides the user with the following options (as illustrated in figure 7.6):

Information about this application
Display 'status' window along bottom of screen
Save current 'tracks' window to disk
Quit this task

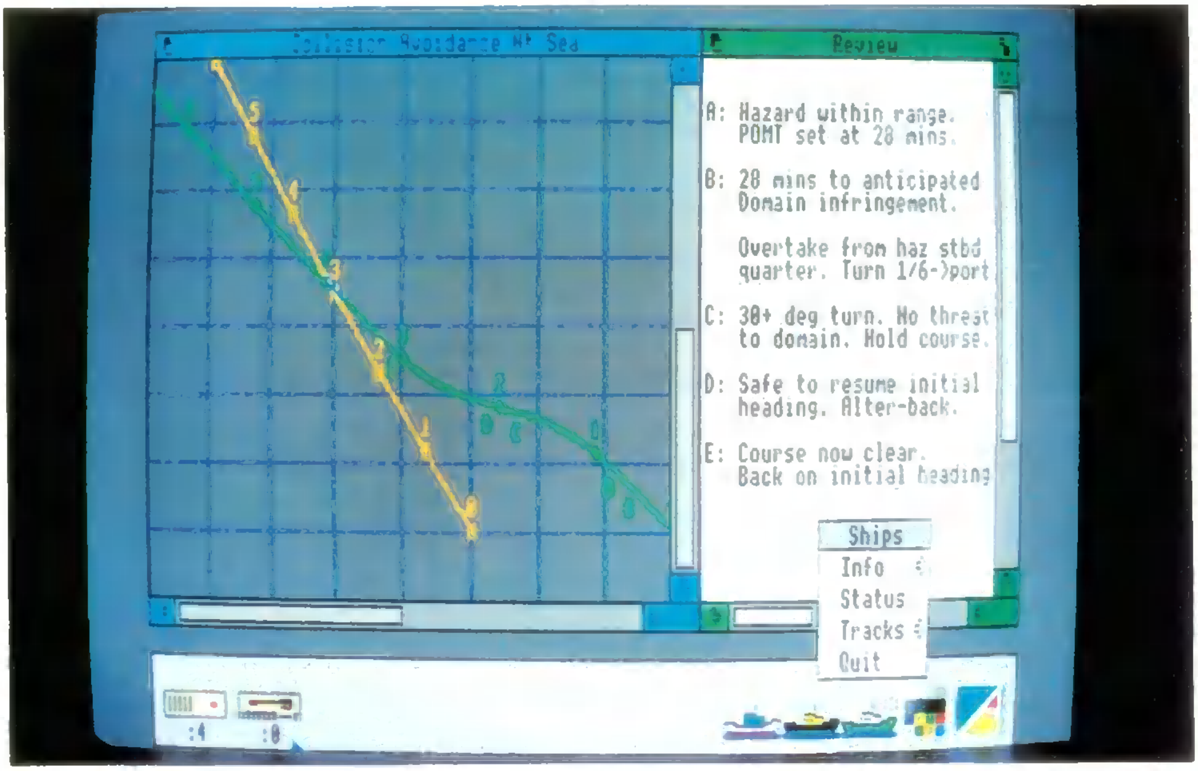

Figure 7.6 Expert System icon pop-up menu.

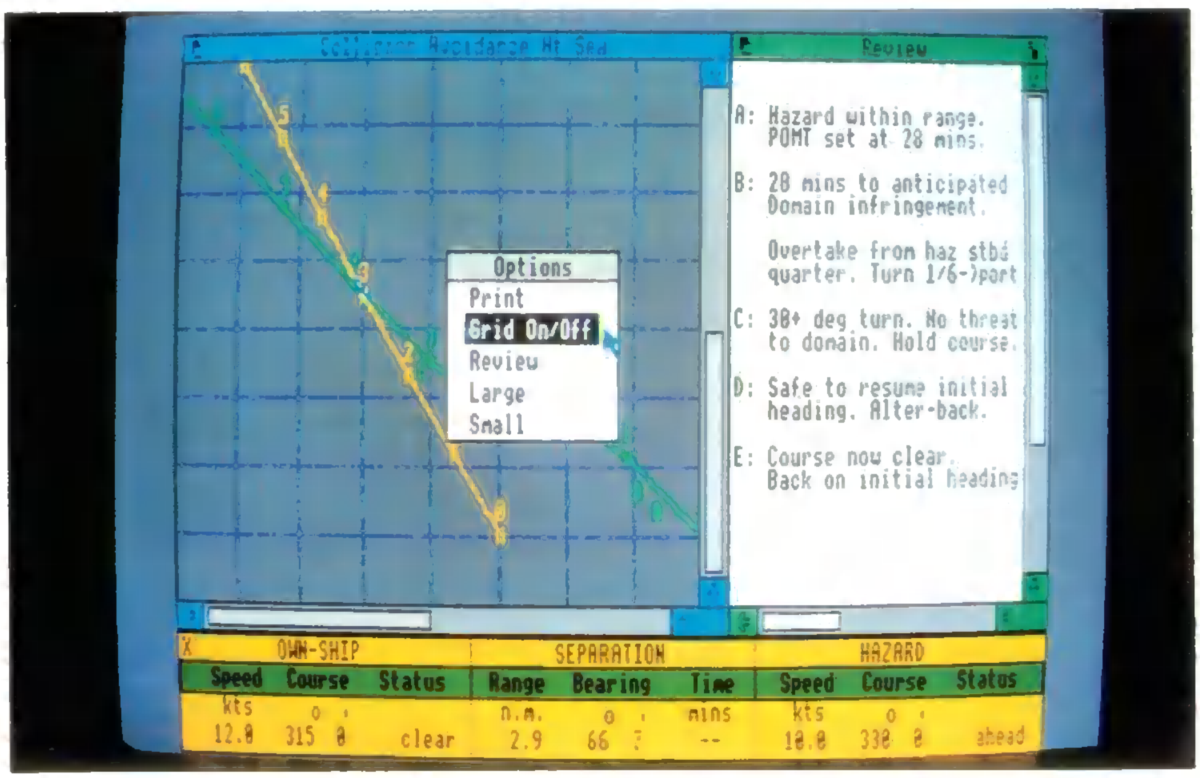

Figure 7.7 'Tracks' window pop-up menu.

As the pointer is moved up or down the menu, each option in turn is highlighted. Pressing the left-hand mouse button will action the highlighted option, and clear the menu; if the option is not a terminating process (such as 'quit'), it may be actioned without clearing the menu by pressing the right-hand mouse button. If it is decided that none of the options is required, the menu may be cleared from the screen by simply moving the pointer away from the menu and pressing the left-hand mouse button.

The other menu in the present system is attached to the 'tracks' window, and offers the user options in relation to that window. The 'tracks' window menu is shown in figure 7.7. The options provided comprise:

Print a hard-copy of the current screen display

Toggle on/off a 1-mile square grid superimposed on the 'tracks' window.

Display review of encounter so far, and mark own-ship's track with corresponding lettering.

Switch display to small/large scale (as appropriate)

It is a fairly simple programming task to extend an existing menu, or to attach a new menu to an icon or a window which is currently not 'menued'. Both of these enhancements may prove appropriate in future extensions of the system.

7.4.2. Dialog Boxes

A dialog box is an interactive facility which is programmed to appear on the screen under specific circumstances. It provides the user with information relating to those circumstances, and requires a response of the user. It may cause 'freezing' of the associated task (and possibly other concurrent tasks, if it is not 'well-behaved') until the user has provided a satisfactory response. Alternatively (depending on the nature of the circumstances), the task responsible for the dialog box may continue unimpeded, changing its course if appropriate once the user has responded.

The form of user response required by such dialog boxes varies widely. The box may be purely informative, requiring only an acknowledgement by the user in the form of clicking the (left-hand) mouse button whilst the pointer is over a 'button' icon (labelled 'OK', or similar) in the box. It may prompt the user for selection from two or more options - again, by clicking on the appropriate 'button' icon. Or it may require text input to one or more text icons within the dialog box. A single dialog box may incorporate any combination of these features.

A number of dialog boxes are used in this system, and it is anticipated that further dialog boxes will be incorporated as the system increases in complexity. Figure 7.8 illustrates the dialog box used to input the initial expert system parameters, on activation of the task by clicking on the expert system task icon; the task does not commence until all entries to the box have been completed, but the presence of this box does not impede any other tasks which may be running concurrently. Figure 7.9 shows the system response if an attempt is made to activate the expert system icon a second time; this dialog box disables all other user interaction for the expert system task until the appropriate response is made by the user, but does not affect other concurrent tasks. Figure 7.10 illustrates the dialog box which is displayed if the 'about this application' menu option is selected from the task icon menu (see above).

Between them, menus and dialog boxes provide powerful tools for two-way communication between the system and the user. With expansion of the system to incorporate alternative hazard ship models, electronic chart data (ECDIS), more complex manoeuvring (for multi-ship and emergency situations) with the possibility of user selection from alternative strategies, it seems likely that such communication will be even more widely used.

Facilities exist, should they prove necessary, for hierarchical structures - sub-menus, dialog boxes called from within dialog boxes, etc. Management of such structures, and ease of use by the non-computer-specialist, are clearly important considerations should this interaction process reach an advanced level of complexity.

Dialog boxes and menus have also been used within the tasks supporting the simulated test environment. Chapter 8 gives further details on this.

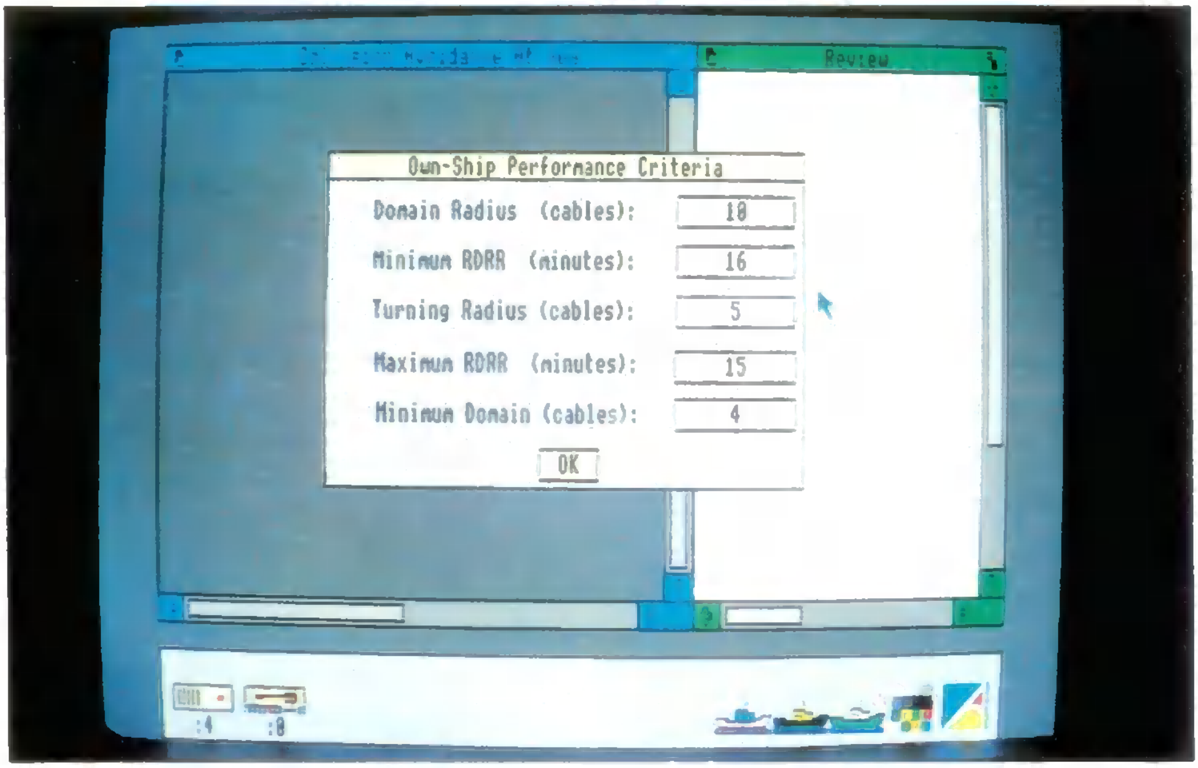

Figure 7.8 Dialog box for input of expert system parameters.

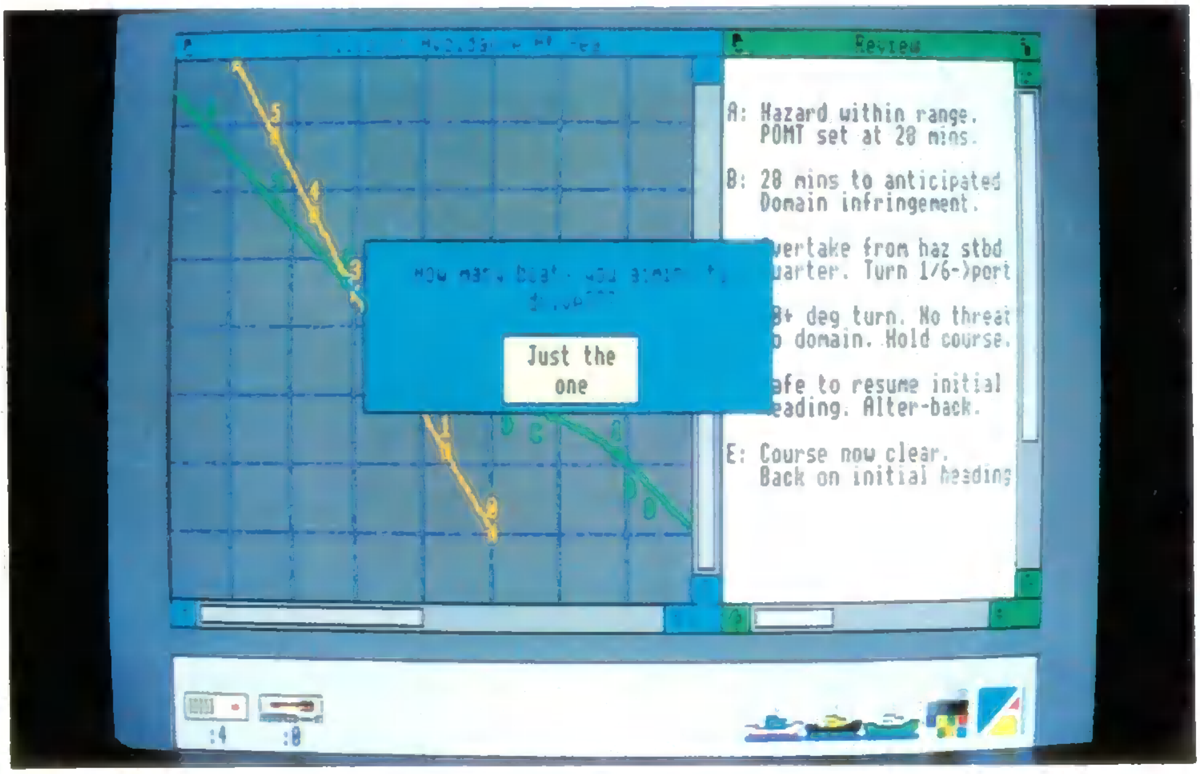

Figure 7.9 Response to repeated activation of expert system.

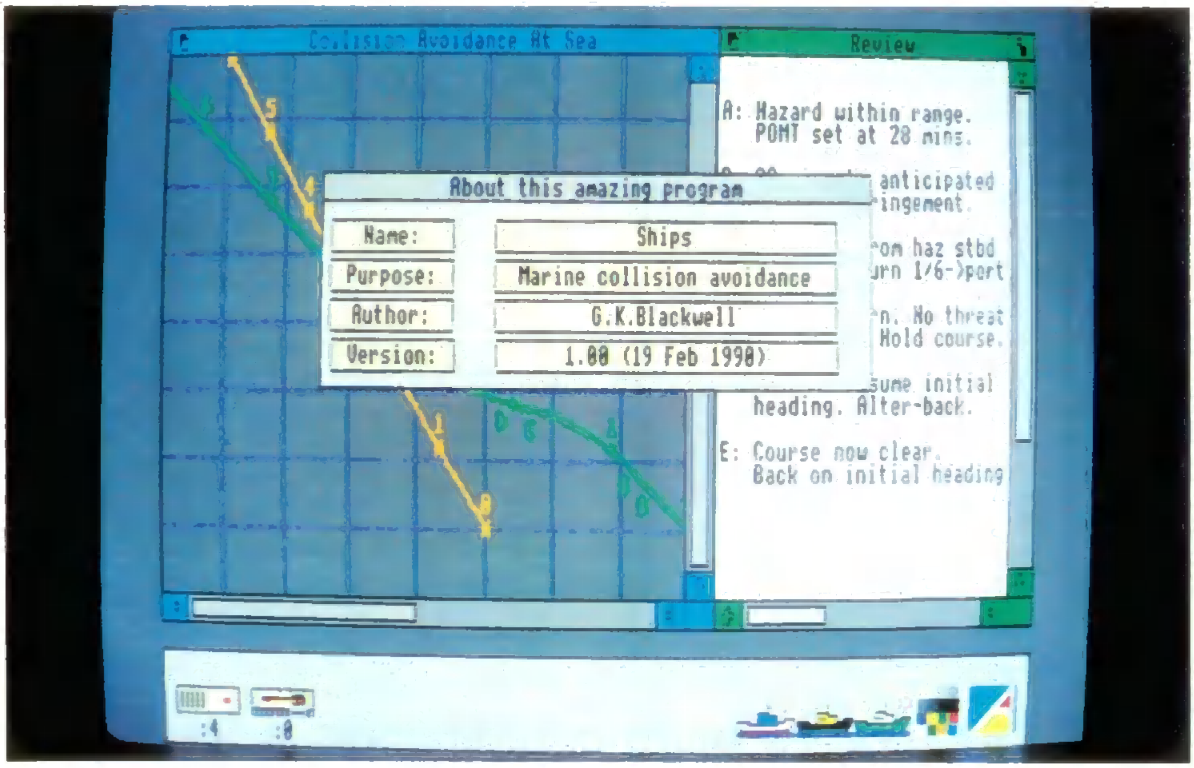

Figure 7.10 Dialog box giving information about the program.

CHAPTER 8

SIMULATED ENVIRONMENTS FOR DEVELOPMENT AND TESTING

"To protect my secret - and myself - I retreated to the safety of this totally artificial Universe and hid myself away in a forgotten cruise liner."

Douglas Adams: The Restaurant At The End Of The Universe (1980)

The development of this expert system has progressed through a number of stages, of increasing complexity. At each stage it has been necessary to provide an appropriate test environment, to study the behaviour of the system under simulated conditions suited to its current level of performance. Each stage is described in some detail below.

8.1. The Integrated Two-Ship System

As explained in section 1.3.2, early prototypes of this expert system were developed on an Atari ST microcomputer. The first of these was a single program which incorporated the expert system decision logic and the drivers for two simulated ships. A schematic of this first prototype is shown in figure 8.1. In this early implementation, no distinction is drawn between the two vessels, apart from labelling them as 'own-ship' and 'hazard'; they could equally well be termed 'ship A' and 'ship B'. The same decision logic is used for both vessels, and both are driven by the same ship simulator code.

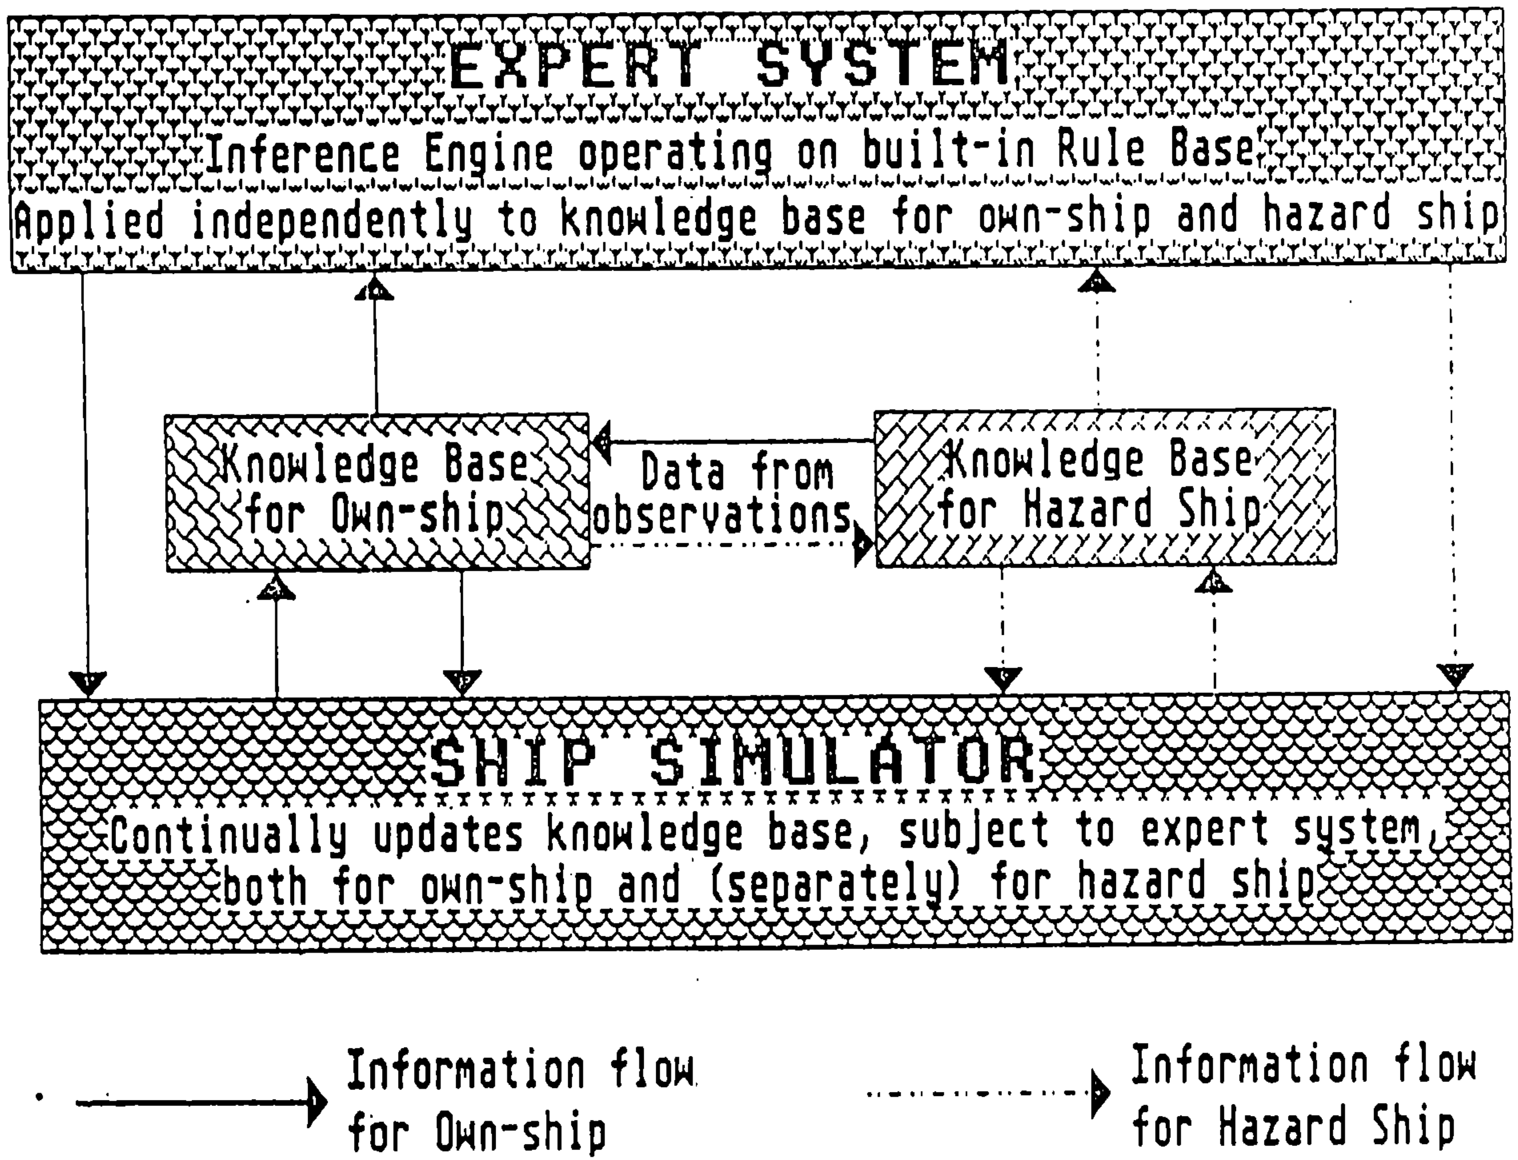

Figure 8.1 Outline structure of prototype system, showing interaction between expert system, knowledge bases for two vessels and ship simulator used to model own-ship and hazard.

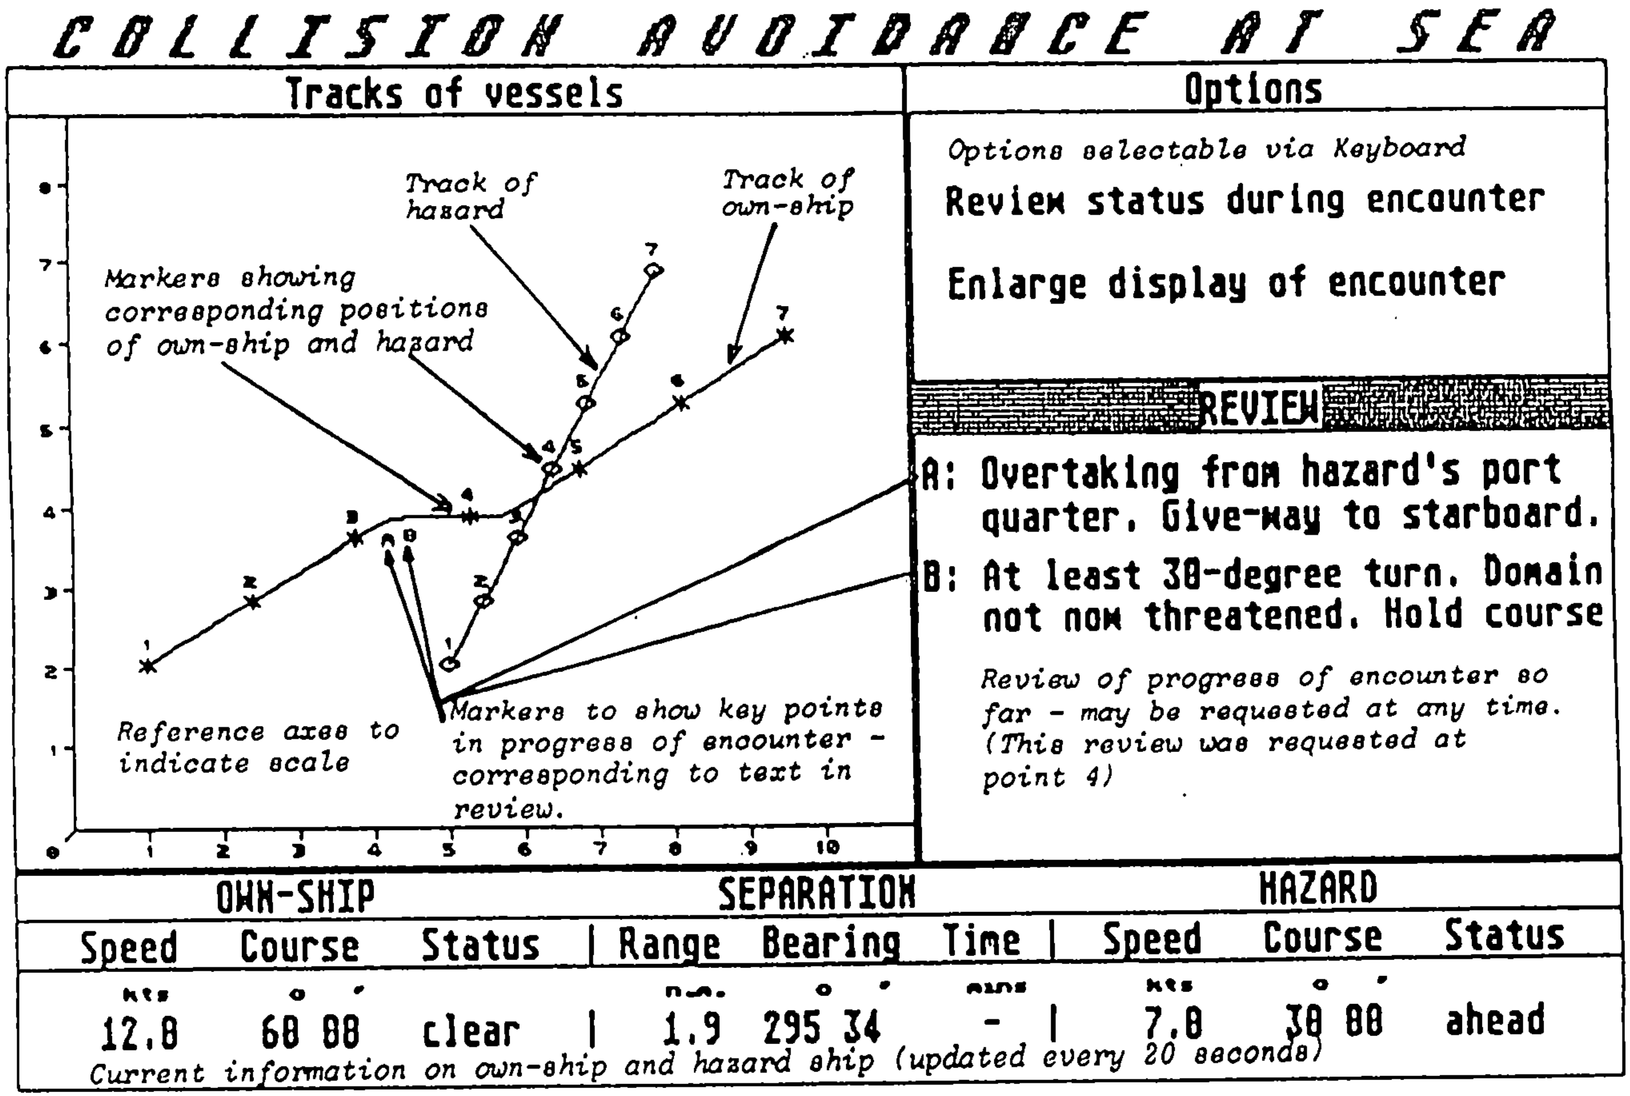

Figure 8.2 Annotated example of screen display generated by prototype system.

Figure 8.2 is an annotated example of the screen display generated by this first prototype. It may be noted that in certain respects the display is asymmetric with respect to the two vessels: the status of 'own-ship' is given definitively, whereas the status of 'hazard' is given only as 'ahead' or 'turning' (as might be deduced from radar observation); the review text, and corresponding track lettering, relate to the situation as seen from 'own-ship'; and the 'separation' details given in the status window are also as observed from own-ship. These are purely features of the user interface, and do not reflect the essential symmetry of the simulation process.

This first prototype was not in any sense intended as an expert system in its own right, but rather as a test-bed for the rule structure and decision logic to be used in later versions. In this version, the rules are actually embedded in the program code, as opposed to being held in a separate file available for simple modification. Also, the separation within the program between the 'expert system' element, the 'ship simulator' element and the 'user interface' element is not explicitly drawn, but is rather an implicit grouping of the functions related to each of these tasks.

As illustrated in figure 8.1, each of the two vessels is represented by a separate 'knowledge base', each holding all of the static information and current dynamic information for its allocated vessel. There is no direct link between these two knowledge bases - the two vessels are not given unrestricted access to information on each other. However, there is a limited interchange of information, comparable to that available through radar in a 'live' situation, as described in the outline of the simulation process below.

At the initialisation of each simulation run, initial parameters for each of the two vessels are input: speed, course, position (on a standard x-y grid), turning circle radius; also domain radius (common to both vessels) and RDRR value to be used for this simulation run. The simulation then proceeds in steps of 20 seconds, real time - simulated by 1/2-second intervals of processor time. In this context, the expert system acts as a controller for each vessel in turn.

At each time-step the system performs the following stages:

1) transfer data on speed, course and position of each vessel to the knowledge base of the other (as one would expect to be available in the 'live' situation through conventional radar and ARPA systems);

2) apply expert system logic to own-ship knowledge base (including acquired data on hazard vessel) and record appropriate action for this time-step; perform corresponding function on knowledge base for hazard vessel;

3) simulate motion of own-ship for one 20-second time interval, and update own-ship knowledge base (taking due account of indicators set by expert system logic in stage 2); perform corresponding function for hazard vessel;

4) output display information to screen for this time-step; record any 'review' information relevant to this time-step, for display on request; send copy of current 'status' information to datalogging file for possible later analysis.

The expert system logic used in this early version consists of the standard rules as described in section 5.1, evaluated on the basis of a fixed RDRR value (input as a parameter); no emergency manoeuvring was included in this implementation. Section 6.1 summarised the results of a sample of 500 randomly-generated encounter situations simulated under this system; that summary is shown as a bar chart in figure 6.5.

8.2. The Modular Multi-Processor System

As indicated in 8.1, the first prototype was not intended as an expert system, rather as a proving-ground for expert system concepts. The next phase was to encompass those concepts in a self-contained body of code and associated data structures, with formalised interfaces for input of information and (possibly) output of control data.

For this purpose it was necessary to explicitly define which sections of the earlier version actually comprised the 'expert system' element of the composite, and what constituted the inputs and outputs for that element (This included the user interface, and one ship knowledge base, for own-ship, as well as the rules with their associated driving logic). It was also necessary to explicitly identify the 'ship simulator' element of the original composite.

The expert system was now to be housed in a dedicated microcomputer system (as intended for the final product), with hardware interfaces for input and output of information from/to associated systems. It was therefore necessary at this stage to add communications modules to the software structure, to enable input and output via the RS232 serial port and MIDI (Musical Instrument Digital Interface) parallel port - these being the two ports available for such purposes on the Atari ST.

A secondary consideration, in isolating the various elements of the original prototype, was to separate the knowledge bases of the two ships. It was felt that shared hardware might be felt to compromise the independence of the two knowledge bases. For this reason, the ship simulator for own-ship was programmed on a second Atari ST microcomputer and the simulator plus driving logic for the hazard vessel on a third ST. This latter package comprised the whole of the original program, but run for one ship only; the data for the other vessel (own-ship) was to be communicated via the interface port.

Thus the configuration shown in figure 8.3 was arrived at. The expert system was connected via the MIDI parallel port to an Atari running a simulator program modelling behaviour of own-ship, and via the RS232 serial port to another Atari which modelled an 'intelligent' hazard vessel (i.e. behaving in accordance with the rules, as incorporated in the original system).

Figure 8.3 3-processor configuration used for second prototype. Expert system running in dedicated microcomputer communicates with simulation environment via serial and parallel links.

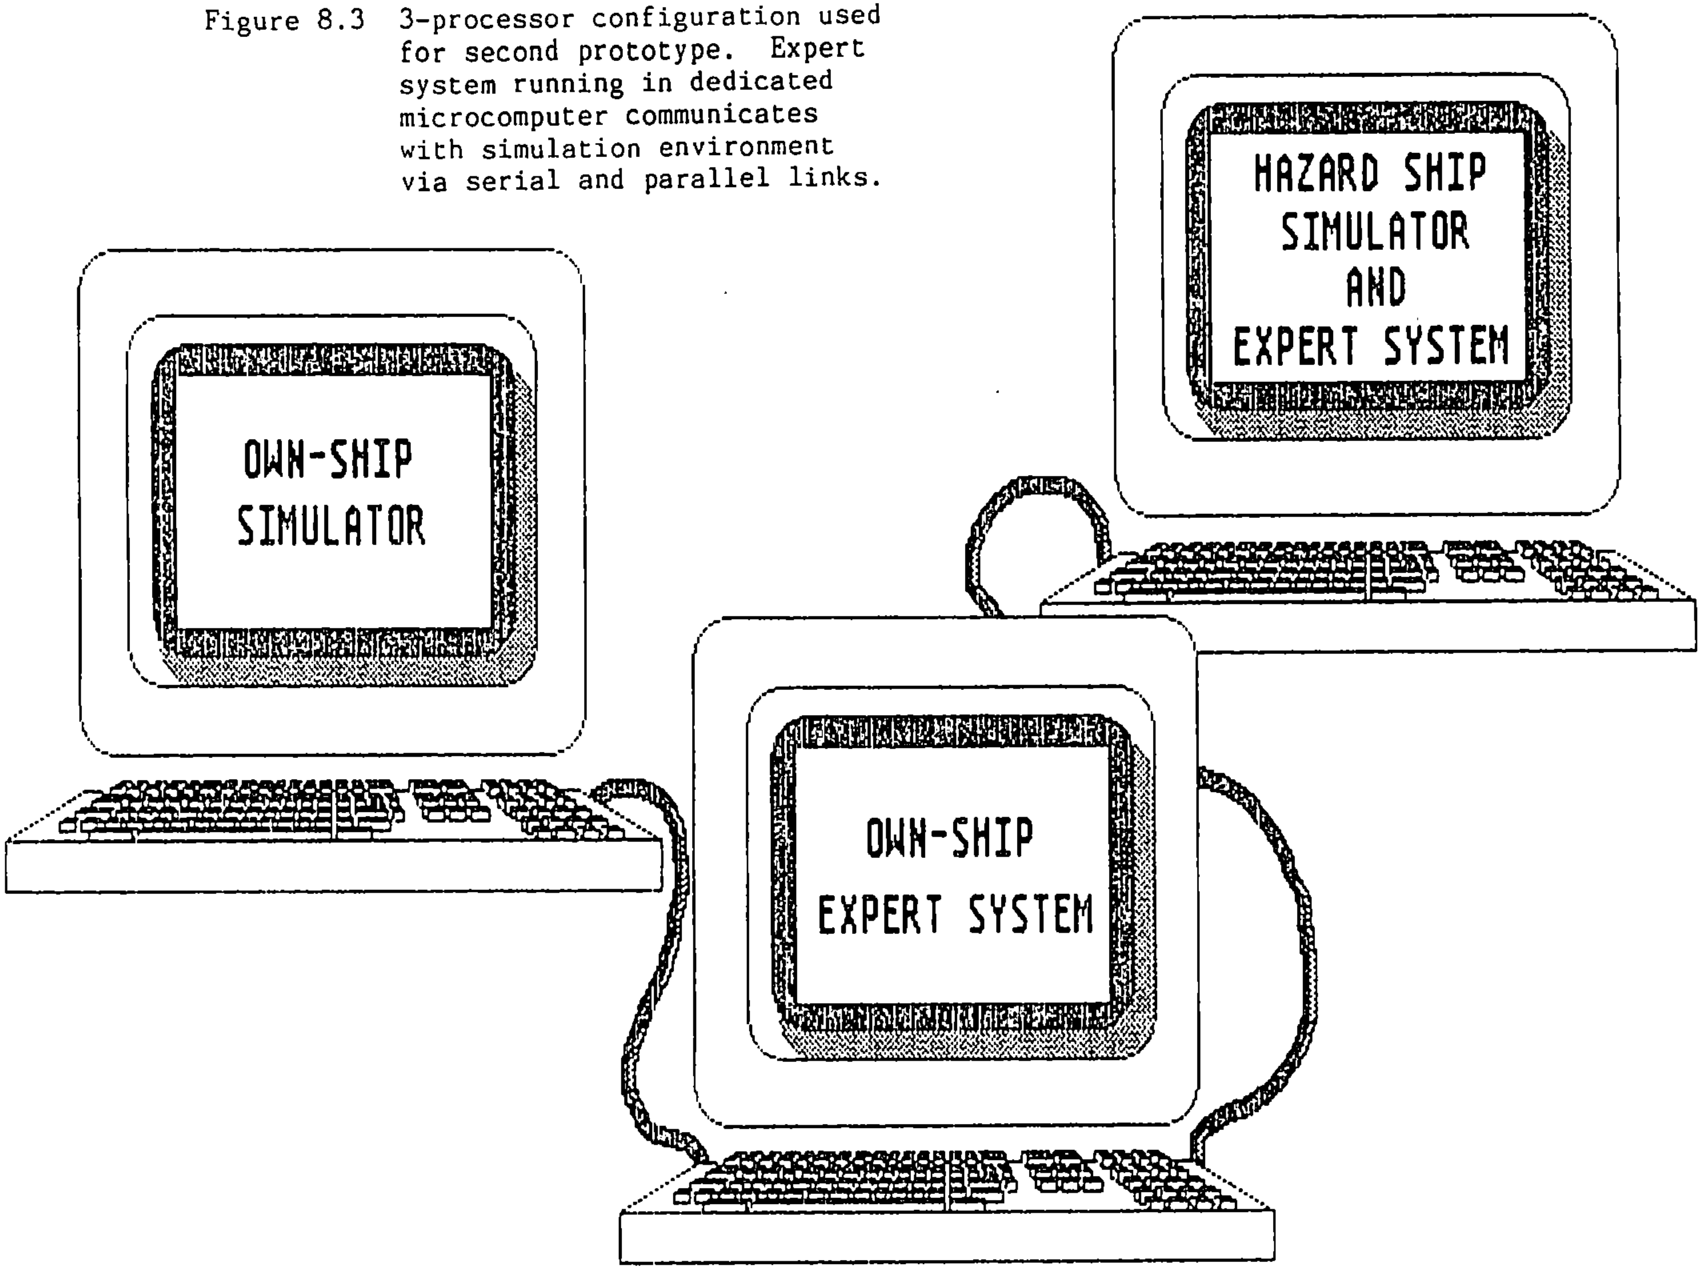

At the beginning of each simulation run, each computer system would be initialised separately, and would prompt for the initial parameters appropriate to its particular function. Domain radius and RDRR value were again common to both vessels, being input to the expert system and communicated to the hazard system via the RS232 link. As each of the simulator systems was initialised, it would signal its state of readiness to the 'expert system' processor via its communication link.

Once all systems were initialised, activities of the three processors would be synchronised by signals from the 'expert system' processor. In all other respects, the functioning of this development configuration corresponded identically to the processing in the first prototype; the screen display on the 'expert system' Atari VDU (Visual Display Unit) was indistinguishable from that on the first prototype for the same input data.

This phase of development was considered to have served its dual purpose in:

(a) defining the distinction between the expert system itself and the peripheral elements which constitute the test environment;

(b) achieving a physical separation between the knowledge base for own-ship (part of the expert system structure) and the knowledge base for the hazard vessel (part of the test environment) - thus assuring the integrity of the information input process for the expert system.

This paved the way for the transfer to the Archimedes multi-tasking system, and the 'multi-task' version of the expert system and supporting test environment, as described below.

8.3. System Development in a Multitasking Environment

Having validated the underlying conceptual basis for the system, a long-term strategy was required which would enable fuller development of those concepts. This necessitated a development environment with potential for expansion to meet the increasing capabilities of such an incrementally-enhanced system. The processing power of the Atari ST was limited; steady accumulation of more Ataris to simulate further hazards and other features seemed less than satisfactory; another viable course of action was required.

As an alternative to the linking of numerous microcomputer systems, a powerful multitasking system appeared to offer a number of benefits:

1) each new task could be added without having to increase the number of computer systems in use;

2) communications between tasks could be effected by internal messaging protocols, rather than external hardware links; this offered benefits in both simplicity and reliability of such communications;

3) datalogging for the various tasks could be performed on a single disk, with printouts to a single printer if required.

It was also felt that, now separation of the various tasks had been established, such separation could be maintained under a multitasking regime, giving the same benefits as the previous multi-processor configuration.

As explained in section 1.3.2 the Acorn Archimedes was chosen for its suitability on a number of fronts. One of these was the multitasking capability, coupled with the facility for opening a 'virtual screen' (i.e. a window) onto any of the current tasks. This makes it possible for the user to monitor various tasks concurrently, and to communicate with any of those tasks (via menus and dialog boxes) at any time. Such an environment is clearly ideal for the intended purpose - particularly when supported by a powerful processor and substantial memory.

The creation of an 'application' to run concurrently with other tasks under the WIMPS environment is a substantial undertaking, requiring strict adherence to system protocols [28]. Each application must first initialise its various event-handling routines, indicating what 'events' (mouse moves, key presses, timer interrupts etc.) are of interest to it, and what action should ensue in such cases. On 'acquiring' the CPU through such an event, a task should conduct all necessary processing and relinquish control of the CPU within a very short time interval; if the CPU is 'locked up' by one task, processing for all other concurrent tasks is suspended for the duration.

Windows should be 'well-behaved', i.e. capable of being overlaid by other windows, and regenerating tidily when overlays are removed. Dynamic memory allocation to such features as menus, dialog boxes and windows must be carefully managed, taking memory from available pool (and later returning it) in such a way as to not conflict with other applications or deplete the pool unnecessarily. Last (but not least) any application which is closing down its operation must do so in an orderly manner, ensuring that no fragments of their display are left littering the screen, and that all files, specialist functions and subsidiary processes (if any) are closed or cleared in accordance with system protocols.

The expert system was transferred to the Archimedes, and coded as a WIMPS application, as just described. At the same time, the look-ahead simulation process (described in chapter 6) was added, with concomitant enhancements to the rule structure and extensions to the 'knowledge base' data structure; emergency manoeuvring, as described in section 5.2, was added a short while later. A representative icon (figure 8.4a) was created to symbolise the active presence of this application within the multi-tasking environment.

Own-ship simulator was developed as a separate WIMPS application, with its own icon (figure 8.4c) and conforming to system protocols. The hazard vessel, operating under the logic of the original prototype, was likewise converted to a WIMPS application (with icon - figure 8.4b). It was anticipated that, with minor modifications, the 'hazard ship'

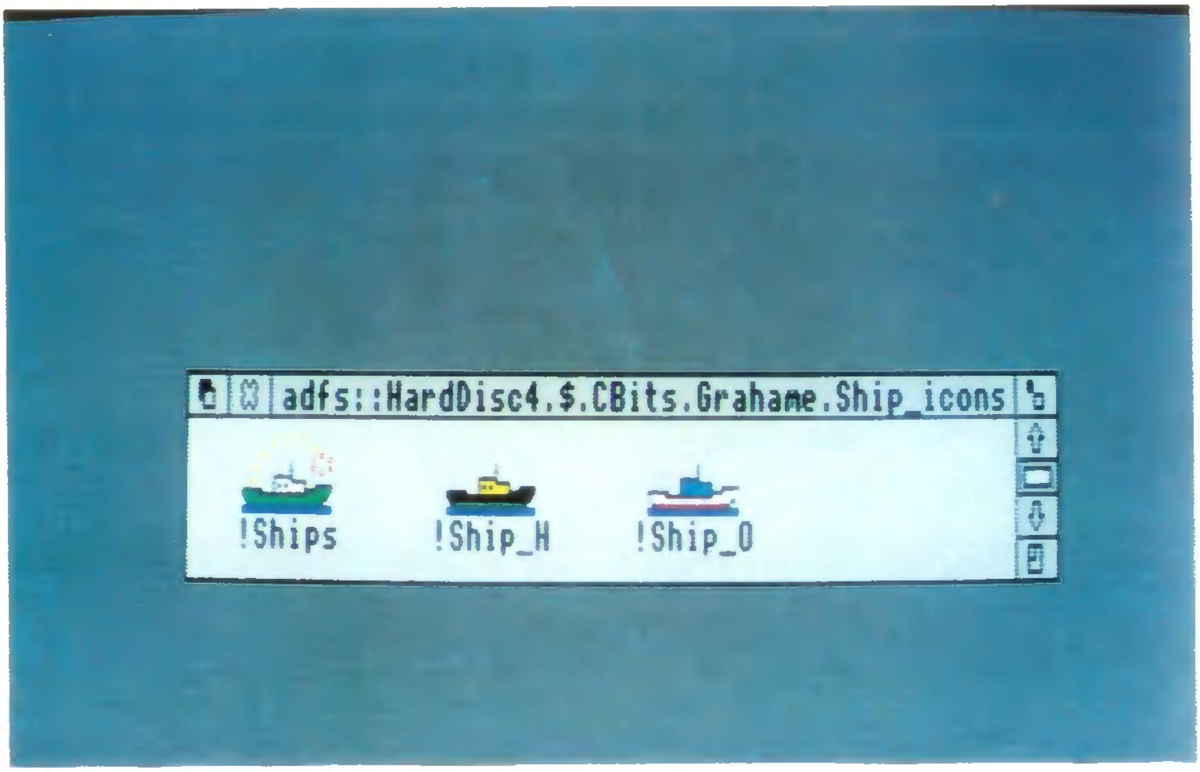

(a) (b) (c)

Figure 8.4 Icons representing expert system and ship simulator tasks

(a) expert system icon

(b) hazard ship simulator icon

(c) own-ship simulator icon

application could be replicated within the multi-tasking environment to represent a number of hazard vessels; this has since been done [6] (see also chapter 10).

The Archimedes operating system, RISC-OS, incorporates a message-handling facility that covers a broad range of functions, passing messages between concurrent tasks or actioning/generating messages between a task and the operating system itself. A message of any sort to an application constitutes an event to be processed by that application.

Messaging protocols must be strictly observed by any application if confusion is to be avoided. Even when observing these protocols, some difficulty was experienced in flushing 'used' messages from the system; continuous repetition of messages (contrary to system documentation) caused serious problems for some time. These problems were eventually overcome by 'serial numbering' each message, and ignoring any message which appeared a second or subsequent time. This is understood to be a 'bug' in RISC-OS, which should be corrected in a later release.

The messaging facility has been used extensively within this development environment, performing the following tasks:

1) upon activation of the 'expert system' application, it broadcasts a message to all peripheral applications currently installed (at present this is only 'hazard ship' application and 'own-ship simulator' application); this causes each to become active and enter its own initialisation routine, as detailed below;

2) periodic time-initiated event messages , generated by the operating system, advise the expert system of each successive time interval for processing purposes (20 seconds in real time, variable - currently 1 second - for simulation exercises);

3) at each time-step, a broadcast message from the 'expert system' application prompts all other related applications to perform their periodic processing functions; a directed message from the 'expert system' application to the 'own-ship simulator' application communicates control information directing activity for this time-step;

4) as each peripheral application completes its processing for one time-step, it communicates relevant information back to the 'expert system' application.

Functioning of the 'expert system' application is described in detail in previous chapters; the following text is concerned with the functioning of the other tasks which comprise the support/development environment for that application.

8.3.1. 'Own-Ship Simulator' Application

The 'own-ship simulator' application models behaviour of the vessel being guided by the expert system; in the development environment, that guidance takes the form of control information transmitted by the 'expert system' application and actioned by this application. As for the expert system, the ship model currently in use is very simplistic, using a combination of straight-line segments and constant-velocity circular arcs. However, a ship model of any desired complexity could be installed in this application without difficulty, giving added realism to the test environment without altering the expert system in any way.

On activation of this process by a broadcast message from the 'expert system' module (see above), a dialog box is displayed. This prompts for and accepts the following data on own-ship:

a) current position (on standard x-y grid);

b) current speed;

c) current course;

d) turning circle radius.

This data is then held by this application for calculation of future situations, as well as being transmitted to the 'expert system' application for processing, display and response. The functioning of the ship model is totally mechanistic, continuing in a straight line or following an arc of a turning circle, as directed by the expert system. A menu attached to the icon for this application enables the closing-down of the application at any time by selection of the 'quit' option. This effectively stops own-ship, without closing down the expert system.

8.3.2. 'Hazard Vessel' Application

This application models the behaviour of an 'intelligent' hazard vessel, i.e. a vessel acting in conformity with the collision avoidance regulations (as implemented in the first and second prototypes). Upon activation of this application by a broadcast signal from the 'expert system' application (see above), functional parameters for this application are prompted for and input via a dialog box. These include:

a) current position (on the standard x-y grid);

b) current speed;

c) current course;

d) turning circle radius;

e) domain radius;

f) RDRR value to be used;

These parameters include the necessary data for modelling ship behaviour, as well as data for reference by the expert system decision logic of the application.

At regular intervals this application is triggered by a message from the 'expert system' application; this message includes current position, speed and course of own-ship (representing radar data). It then performs the following sequence:

1) current situation is evaluated by inbuilt expert system logic, and course of action defined for this time-step;

2) ship simulator module implements decisions made in (1), and calculates hazard vessel's new situation (position, speed, course) after this time-step (20 seconds, real time);

3) a message is sent to the 'expert system' application, giving this new information (representing data available via radar).

At any time this application may be closed down, by selection of the 'quit' option in the pop-up menu attached to the 'hazard vessel' icon - see figure 8.5. This has the effect of stopping the hazard vessel, without in any other way affecting the 'expert system' application.

The menu shown in figure 8.5 also offers an option of manual control of the hazard vessel. This enables the operator to bypass the expert system logic in this application, forcing the vessel to continue ahead, or to turn either to port or starboard. In this way, the hazard vessel may be held on a collision course when it should by rights give-way (according to the regulations), or even directed into such a collision course at a late stage.

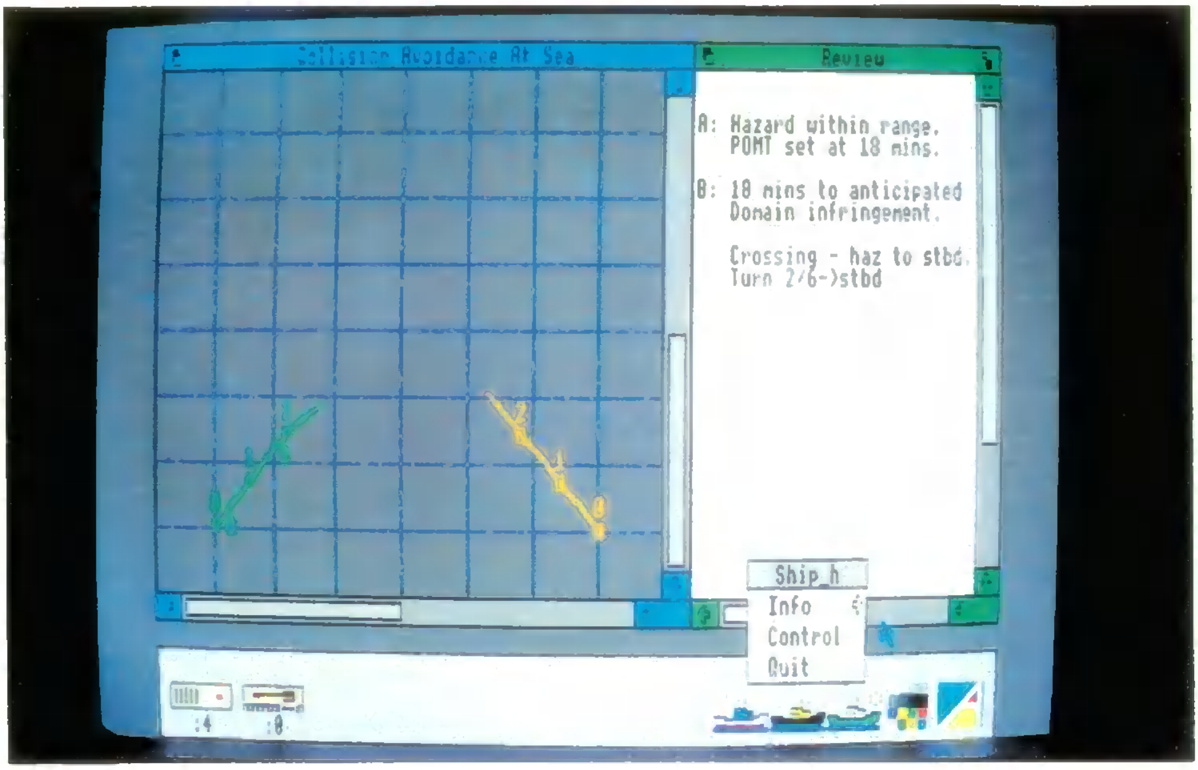

Figure 8.5 Hazard Vessel Menu

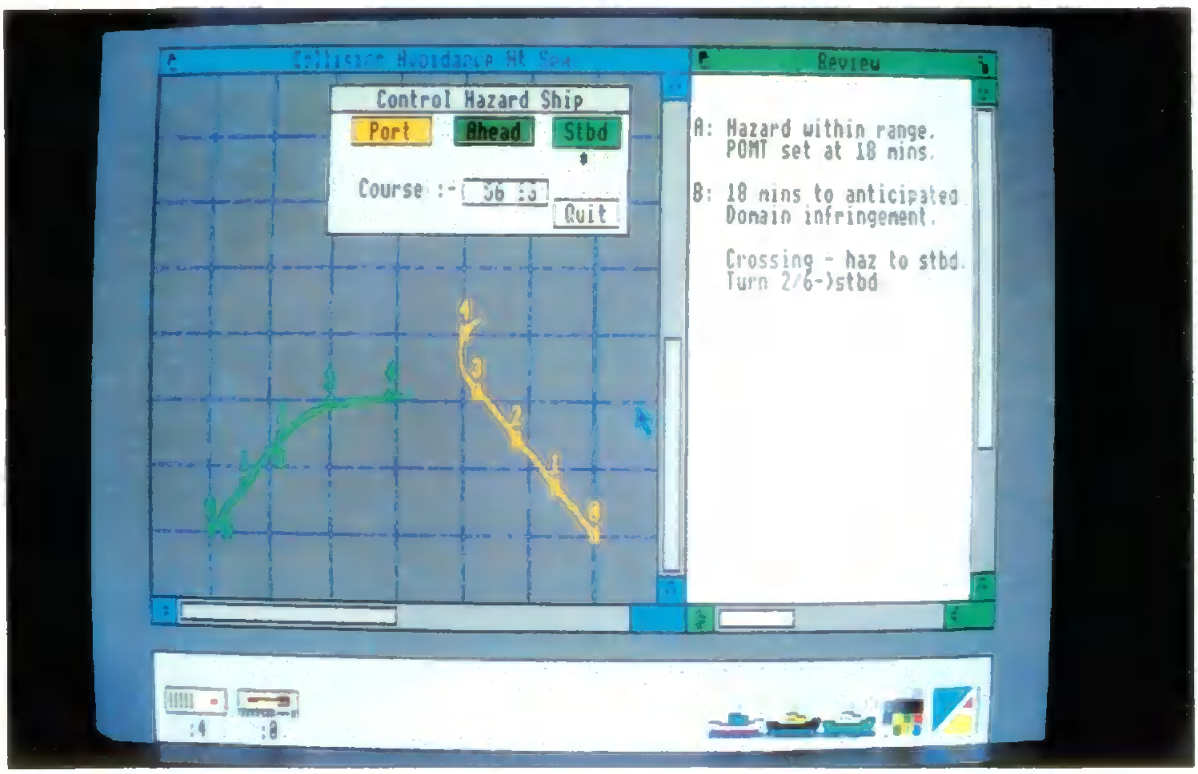

Figure 8.6 Hazard Vessel Control Box

This facility has been provided to enable the operator to create emergency situations, to test emergency manoeuvring strategies in the 'expert system' application. This requires the introduction of such 'rogue' hazard vessels, which do not conform to the collision avoidance regulations. On selection of this menu option, a dialog box is displayed (figure 8.6) in which the user may select 'port', 'ahead' or 'starboard' as required. This 'control box' remains on display for the duration of any manual intervention. A star below the selected course indicates the operator's action. On completion of any manually-directed manoeuvre, selection of the 'quit' option in this dialog box closes the box and returns the vessel to automatic control.

The addition of the manual control box to the hazard vessel is illustrative of a general principle of this development environment. It is possible to enhance, extend or even completely restructure any part of that environment, as long as the format of communication with other parts of the total system is maintained. It should also be possible to introduce new elements to this environment as separate applications, with the relatively simple addition of message-handling routines in each case.

CHAPTER 9

SYSTEM VALIDATION

"Every man's work, whether it be literature or music or pictures or architecture or anything else, is always a portrait of himself."

Samuel Butler (1835-1902) The Way Of All Flesh

In order to confirm satisfactory performance at the current stage of development, the system must be validated against three criteria:

(1) effectiveness of the rule base in ensuring non-violation of the domain;

(2) appropriateness of avoidance manoeuvres in various situations, as judged by competent mariners;

(3) viability of the system on-board ship.

Each of these criteria has been addressed in the validation phase, by respectively:

(1) simulating one thousand randomly-generated encounter situations and measuring the separation at CPA without, and then with, collision avoidance directed by the expert system;

(2) presenting a variety of simulated encounter situations to groups of experienced mariners, and analysing their responses to the system's handling of those situations;

(3) installing the system on-board a research vessel (see appendix B) equipped with ARPA, setting up a number of encounter situations involving a sister vessel (the 'hazard'), and following the directions of the expert system in each case to assess its efficacy.

The results of these three processes are detailed in the following sections.

9.1. Random Simulation of a Large Number of Encounter Situations

A driver program was written, capable of generating a specified number of encounter situations, subject to the following conditions:

(1) own-ship was initially placed at a fixed position, and a second position randomly generated for the hazard vessel, such that the range between the two vessels was less than the domain radius (0.8 nautical miles for this exercise);

(2) each vessel was allocated a random course and a random speed in the range 5 - 15 knots; own-ship was also given a random turning circle radius in the range 0.2 - 0.7 nautical miles;

(3) a 'time-reversal' process was then applied whereby each vessel was drawn back along its track for a distance corresponding to equal time intervals, until the hazard vessel was outside own-ship's 'ballpark' (20 minutes to domain infringement, in this case);

(4) the separation at CPA without collision avoidance was calculated (assuming both straight-line tracks); the encounter was then allowed to proceed under the direction of the expert system logic, and the new separation at CPA recorded.

Notes on the multiple simulation program:

(a) the inference engine and rule structure for the full expert system were copied without modification, for operation by the driver program; exclusion of all aspects of the user interface, and separation from the WIMPS environment, enabled considerable speeding-up of the simulation process;

(b) the hazard vessel was not provided with any avoidance capability; encounters in which the hazard was recognised as the give-way vessel were aborted, and not included in the final result (or the count of number of encounters);

(c) the 'early course change' option was included at the 20-minute RDRR.

As might be expected, early trials with this program generated encounter situations which had not been envisaged in previous 'one-off' encounters set up manually. A few of these highlighted loopholes in the rule structure, falling broadly into three categories:

(i) some rules required two criteria to be satisfied before moving on; in exceptional circumstances these criteria proved to be mutually exclusive, and so the manoeuvre was never completed;

(ii) a manoeuvre which should have been feasible without domain infringement proved to have infringed the domain - sometimes quite badly; this was found to be due to the operation of certain rules in borderline cases, deferred operation of such rules forcing at least two executions (possibly then occasioning significantly more) where only one execution of the rule was required;

(iii) the return-to-course segment of the manoeuvre was calculated on the basis of a diagram which was believed to represent a fully general case; it transpired that a small number of encounter situations did not conform to this diagram, and consequent erroneous calculations led to premature returning-to-course, and hence domain infringement.

Having identified these exceptions, and their root causes, the following corrective measures were taken:

(i) where one rule was found to contain two requirements which were on occasions mutually exclusive, an extra rule was inserted so that the requirements could be checked sequentially, rather than at the same time; in such cases, fulfilment of one criterion followed by fulfilment of the other has proved satisfactory, both in meeting the needs of the rule structure and in resolving this problem;

(ii) excessive operation of a particular rule (notably 'oversteering') was simply remedied by changing certain rule linkages from 'deferred' to 'immediate' status;

(iii) the problem of the 'return-to-course' segment was due to relating the circle geometry involved to a specific diagram - which did not fully represent all possible cases; this was replaced by a totally different form of analysis, based on the binary halving concept (see appendix A); this approach, which makes no suppositions as to the juxtapositions of the two vessels, overcame that limitation.

Following the above amendments to the rule base logic, further trials have been carried out using the multiple simulation program. Table 9.1 gives the results of a run of 1000 randomly-generated encounters, also shown in the form of a bar chart in figure 9.1. As can be seen, all of these encounters threatened infringement of the 0.8 n.m. domain area; however, with the application of the expert system rules, domain infringement has been averted in every case. Moreover, the vast majority of avoidance manoeuvres have required the bare minimum of diversion, taking the threat just outside the domain area.

Closest Point of approach (cables)	Number of instances without ACAS	Number of instances with ACAS (Automatic Collision Avoidance System)
0 to 1-	309	0
1 to 2-	180	0
2 to 3-	150	0
3 to 4-	116	0
4 to 5-	101	0
5 to 6-	66	0
6 to 7-	49	0
7 to 8-	29	0
8 to 9-	0	600
9 to 10-	0	167
10 to 11-	0	83
11 to 12-	0	63
12 to 13-	0	34
13 to 14-	0	24
14 to 15-	0	14
15 to 16-	0	8
16 to 17-	0	2
17 to 18-	0	3
18 to 19-	0	2
19 to 20-	0	0

Table 9.1 Results from 1000 randomly-generated encounter situations (see fig. 9.1 for details of parameters used)

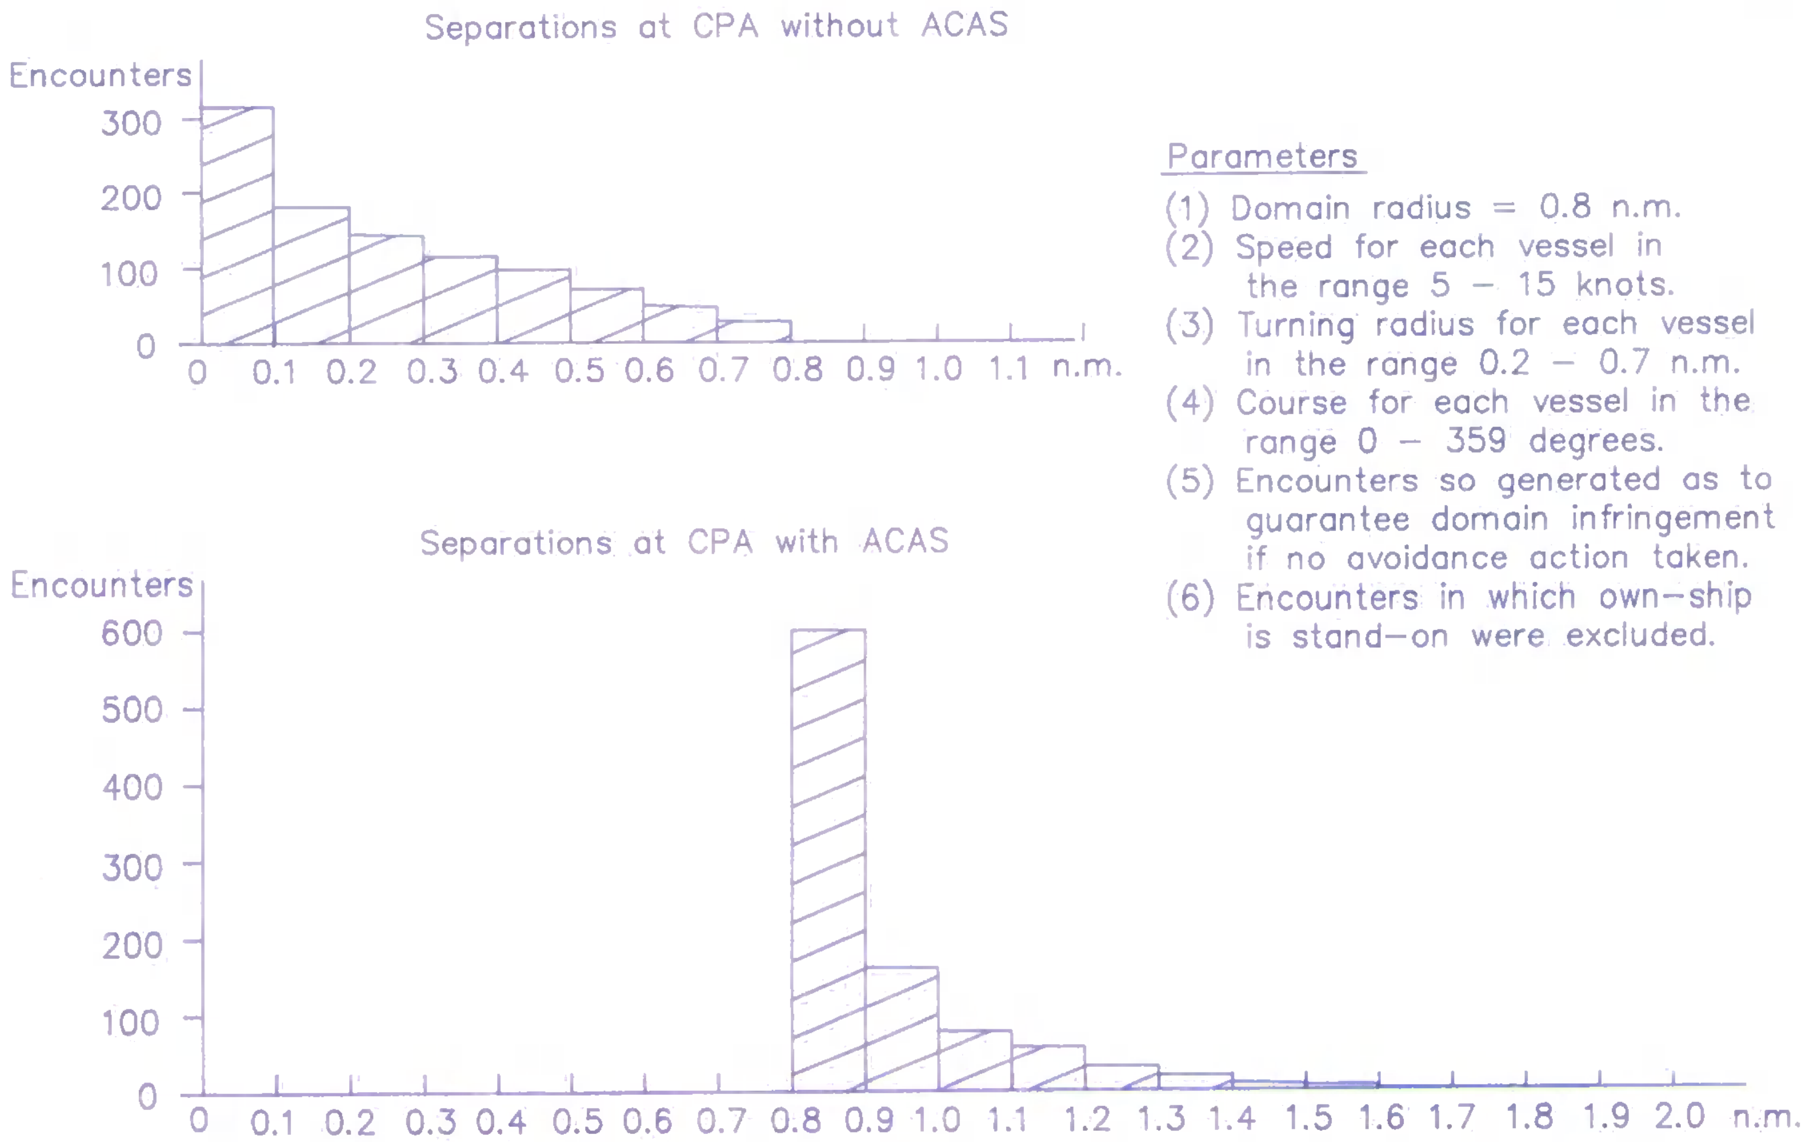

Figure 9.1 Effects of Automatic Collision Avoidance System on a sample of 1000 randomly-generated encounter situations.

Clearly, the requirement for a clear turn, coupled with the large turning radius of some of the simulated vessels, has led to rather greater divergence than necessary. However, this substantial simulation exercise vindicates the expectation that:

(a) the expert system will direct manoeuvres so as to ensure avoidance of domain infringement;

(b) avoidance manoeuvres will be conducted with minimal diversion from course, subject to regulations and other safety considerations.

9.2. Appraisal of System Performance by Experienced Mariners

For this exercise, a sample of 25 experienced mariners was drawn from two different locations and two levels of expertise. First, 14 ship's masters were consulted whilst on a course at Southampton Institute, Warsash. Shortly afterwards, the system was demonstrated to a group of 11 first mates on a master's course at Polytechnic South West, Plymouth. Both groups of mariners were shown the same set of encounter situations, and asked to record their views for each encounter on a standard questionnaire (figure 9.2). As a consequence of observations made by the first group, certain parameters were varied for the second set of tests, illustrating the ability of the system to respond to user requirements without major modifications.

Figure 9.2

QUESTIONNAIRE
(re: Computer-Assisted Collision Avoidance System)

1. Do you find the graphics display comprehensible? Y/N ______

   Comment: ______________________________________

   ____

   ____

2. Do you find the supporting status information

   (a) Comprehensible? Y/N ____ (b) Useful? Y/N ____

   Comment: ______________________________________

   ____

   ____

3. Do you find the review text

   (a) Comprehensible? Y/N ____ (b) Helpful? Y/N ____

   Comment: ______________________________________

   ____

   ____

4. Do you consider that the avoidance manoeuvre shown conforms with good practice? Y/N ____

   Comment: ______________________________________

   ____

   ____

5. Would you consider this system to be a useful training aid?

   Y/N ____

   Comment: ______________________________________

   ____

   ____

As shown in figure 9.2, the questionnaire covers three aspects of the system:

(i) presentation of information to the user (Q.1 - 3), in terms of (a) clarity and (b) usefulness;

(ii) suitability of avoidance manoeuvres conducted in response to specific encounter situations (Q.4);

(iii) appropriateness of the system for use as a training aid (Q.5).

It was agreed with both groups that responses to Q.1 - 3 and Q.5 would not vary from one encounter to another (except insofar as system response to a specific situation might be considered inappropriate - this is covered by Q.4). Consequently, only Q.4 was considered separately for each encounter, responses to the other questions being given only once, on the questionnaire for the first encounter demonstrated.

Prior to demonstrations of specific encounter situations, the system was explained in detail, in terms of its function, to each group of mariners. The following points were made:

(i) the system as demonstrated deals with two vessels (own-ship and hazard) in the open sea; however, it is intended that the system should ultimately be expanded to handle multi-ship encounters in confined waters;

(ii) with (i) in mind, together with other considerations (vessel size, manoeuvring characteristics, etc.), the system incorporates scope for user variation of a number of functional parameters - passing distance, minimum time before encounter at which action should be taken, emergency manoeuvring considerations; more 'fixed' considerations - notably minimum turn for an avoidance manoeuvre - may be varied within the program without difficulty, if considered inappropriate;

(iii) the system is intended to operate in an advisory capacity, with inputs from on-board sensors (rudder, speed log etc.), ARPA and ultimately electronic charts and navigation systems; however, for development purposes the system on display simulates the relevant inputs through two additional processes - own-ship and hazard simulators - and controls the actions of own-ship according to the advice given.

The nature of the system display was shown and explained, including demonstration of user menu options.

Question 1 received a 'yes' response from every one of the 25 mariners involved in this exercise. I.e. the graphical representation of ships' tracks was considered clear and unambiguous.

Question 2 received a very favourable response from the masters at Warsash, less so from the second mates at Plymouth. Of the 14 at Warsash, 11 considered the status information comprehensible and 12

considered it useful; two masters abstained on the first point, and two (not the same two) on the second. It is interesting to note that one 'no' and one abstension on the comprehensibility issue both found the information useful!
The 11 mariners at Plymouth were divided on both the points at issue, being split 5-5 (with one abstension) on comprehensibility and 4-4 (with three abstensions) on usefulness. This time, three who voted 'no' for point (a) still found the display useful!

Question 3 elicited a fairly strong vote of confidence from both groups in terms of clarity of presentation (10 yes, 2 no, 2 abstensions at Warsash, 8 yes, 3 no, no abstensions at Plymouth); however, the reaction to the 'helpfulness' question was again split in both camps - 5 yes, 2 no, 7 abstensions at Warsash; 5 yes, 5 no, 1 abstension at Plymouth. The critical comment voiced by several of the 'no' voters concerned the use of non-maritime jargon (domain, POMT) in the review text. Terms such as TCPA would be much preferred - a valid point which can be met by suitable modifications to the review text. One 'yes' voter added the comment that the situation was "well explained" by the review text.

Question 4: Responses to the various encounters, and the system's reaction to them, are dealt with in detail below.

Question 5: The majority of those canvassed regarded the system as eminently suitable for training purposes, 12 voting 'yes' (one 'no', one abstension) at Warsash, and 8 in favour (three 'no', no abstensions) at Plymouth. One of the Plymouth 'noes' commented that it was OK for classroom work, but unsuitable for use on a ship - perhaps the point of this 'training' question was missed here.

Responses to Handling of Encounters (Q.4)

Three encounter situations were demonstrated, two of which featured own-ship as the give-way vessel and one in which own-ship was stand-on to a rogue vessel, forcing an emergency manoeuvre. The nature of each of these encounters, and the responses of both groups of mariners to each, are detailed below.

Encounter #1

Own-ship heading 040 degrees, speed 7 knots, initial position (1,1) on grid. Hazard heading 320 degrees, speed 7 knots, initial position (7,1). These parameters give a crossing encounter with own-ship the give-way vessel, leading to a collision if no action is taken.

When demonstrated to the masters at Warsash, the domain radius was set at 8 cables, and the minimum RDRR at 8 minutes; the system actually calculated a POMT of 18 minutes. The 'minimum clear turn' setting within the program at this time was 15 degrees, and the look-ahead routine was set to evaluate the situation 30 minutes before domain infringement. The screen display for this encounter, as observed by the Warsash mariners, is shown in figure 9.3(a).

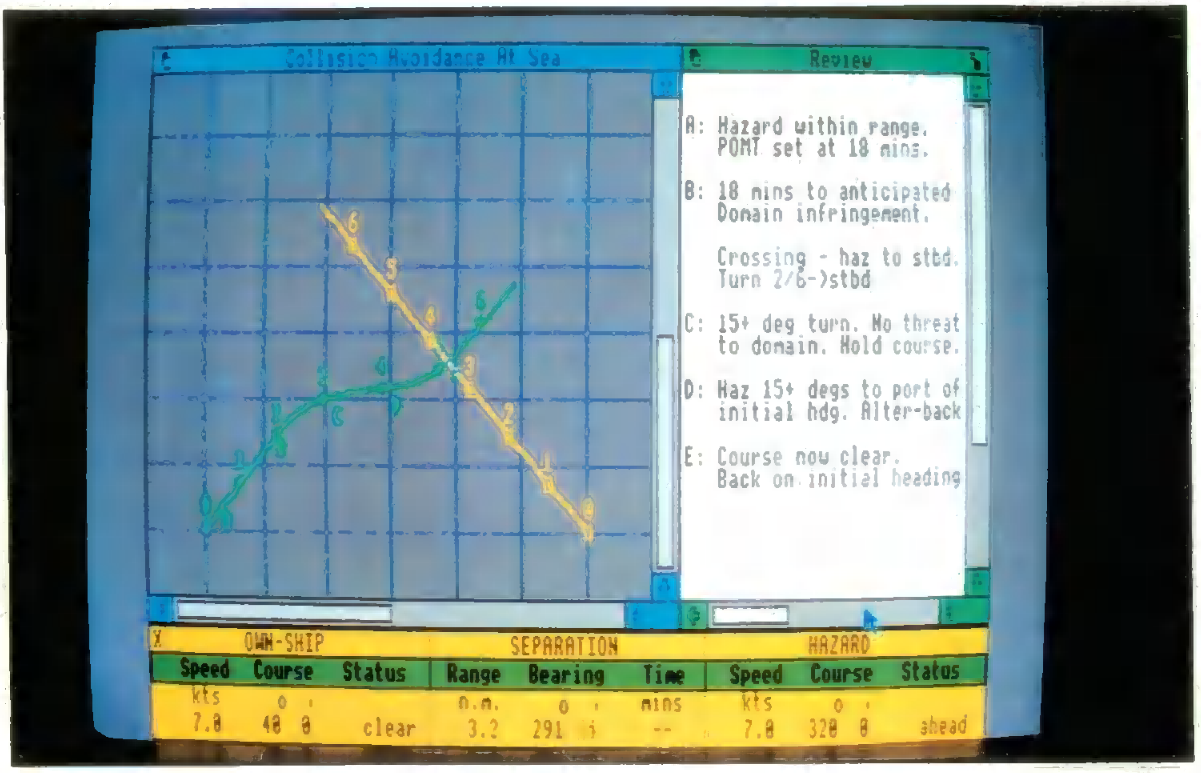

Figure 9.3(a) Encounter #1, as shown to Warsash mariners

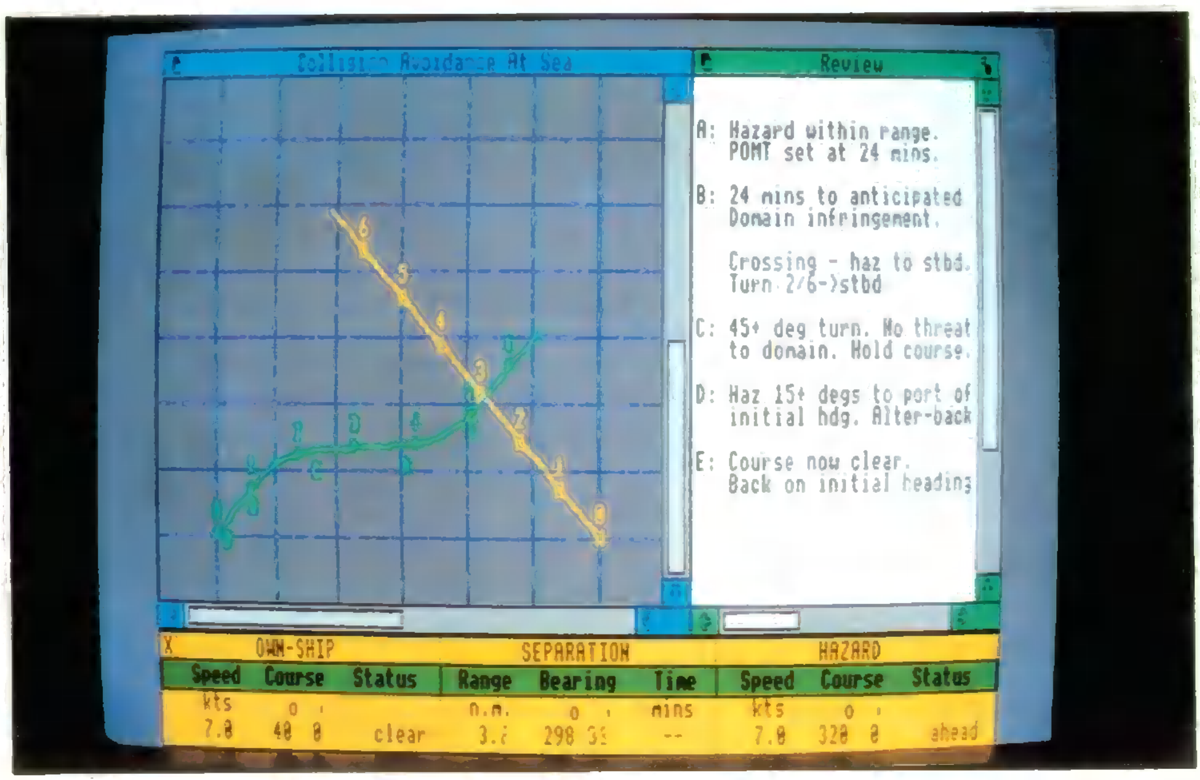

Figure 9.3(b) Encounter #1, as shown to Plymouth mariners

The response to this demonstration was mixed, 6 voting 'yes' on Q.4, 7 voting 'no', and one abstaining. The critical observation made by a number of participants was that the manoeuvre would be better performed earlier in an open-sea situation, and that the amount of turn was inadequate for a clear indication of intent.

It was considered appropriate to act upon these comments, and then ask the second group of mariners to give opinions on the revised approach to this situation; this would effectively validate the ability of the system to respond to user requirements by simple adjustment of system parameters. Consequently, the 'minimum turn' was reset to 45 degrees, the domain radius for this encounter was set to 12 cables, and the minimum RDRR to 24 minutes. The effect of these changes is shown by the screen display in figure 9.3(b).

The response of the Plymouth mariners to this revised approach was substantially more favourable than the Warsash reaction, 9 voting 'yes' on question 4, 2 voting 'no' and none abstaining (both 'no' voters still considered the manoeuvre to be too late, one observing that it would be acceptable in confined waters). Thus a 43% favourable response at Warsash became an 82% favourable response at Plymouth. This was clearly a vindication of the increase in the system parameters, corresponding to "acting in good time", "making a clear turn" and "passing at a safe distance", as (rather subjectively) worded in the Collision Avoidance Regulations.

Encounter #2

Own-ship heading 045 degrees, speed 7 knots, initial position (1,1) on grid. Hazard heading 225 degrees, speed 7 knots, initial position (7,7). These parameters lead to a head-on collision if no action is taken; there is responsibility upon both vessels to alter to starboard.

In the Warsash demonstration, the domain radius was set at 10 cables and the minimum RDRR at 8 minutes; the system again evaluated the POMT at 18 minutes. The screen display for this encounter is shown in figure 9.4(a) - as can be seen, hazard took no avoiding action, having evaluated the situation rather later when own-ship was already clear.

Owing to time constraints, only 7 of the Warsash masters observed encounter #2 (and 7 encounter #3). Of these, 5 considered that this manoeuvre conformed with good practice, 2 not so; none abstained. Given the majority opinion in favour of this manoeuvre, the domain and minimum RDRR were not changed for this manoeuvre in the Plymouth trial; the modification from 15 to 45 degrees minimum turn was, however, retained - figure 9.4(b) shows the minor variation so caused.

The Plymouth group was also reduced to 7 for encounter #2 (and #3), due to other commitments. Their response was identical to Warsash: 5 'yes', 2 'no', no abstensions. One 'no' voter commented that action time should be based on time to CPA, not distance ("What if both ships were making 30 knots?"); since it had been explained that POMT is a time-based constraint, it must be assumed that this participant had misunderstood.

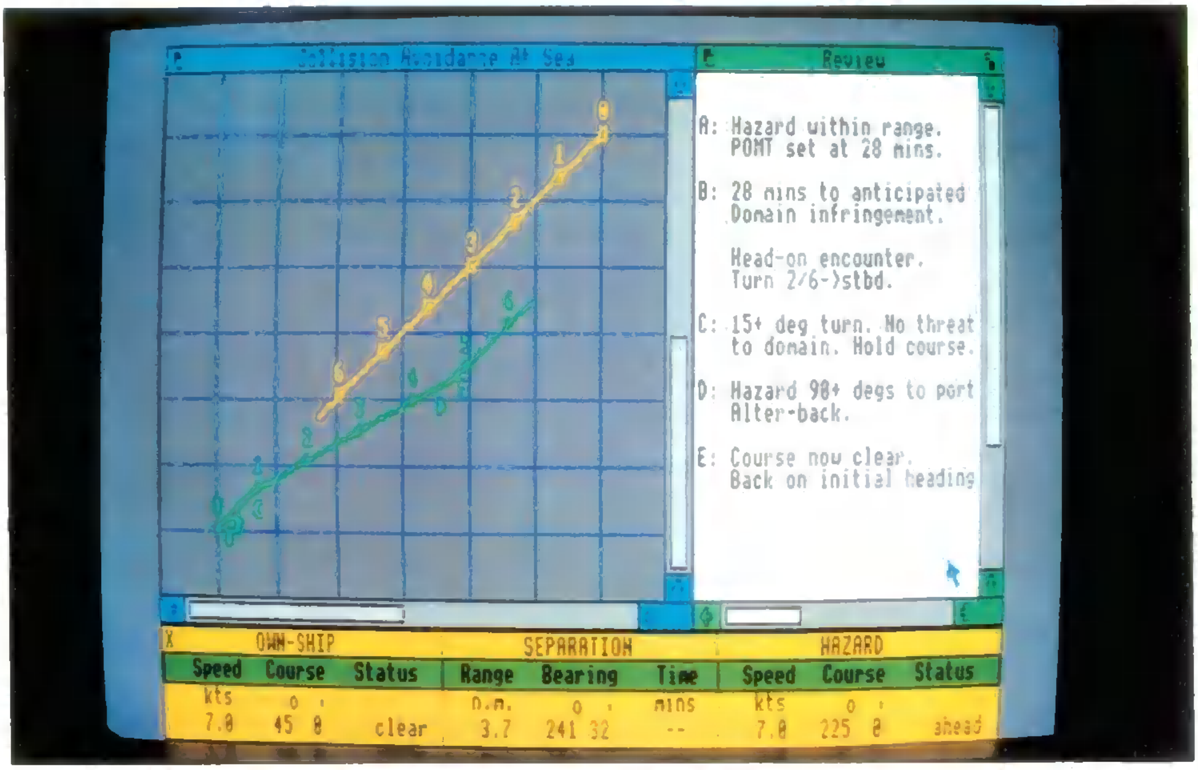

Figure 9.4(a) Encounter #2, as shown to Warsash mariners

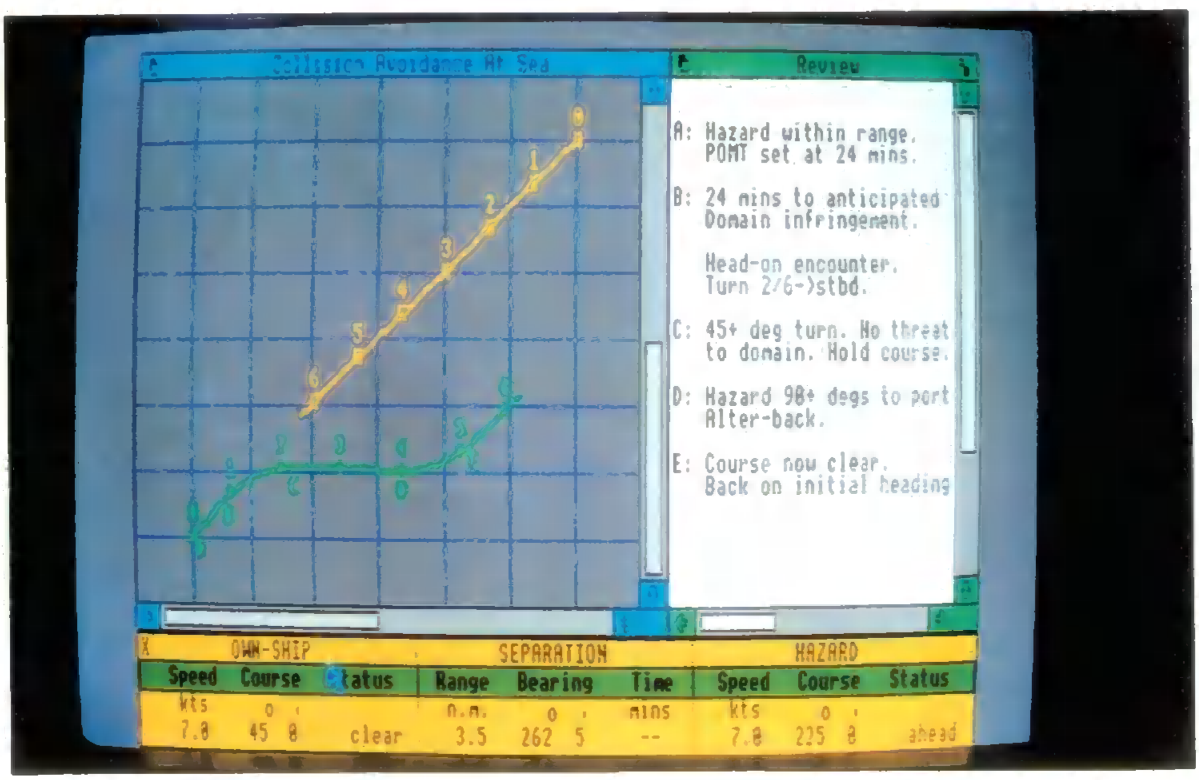

Figure 9.4(b) Encounter #2, as shown to Plymouth mariners

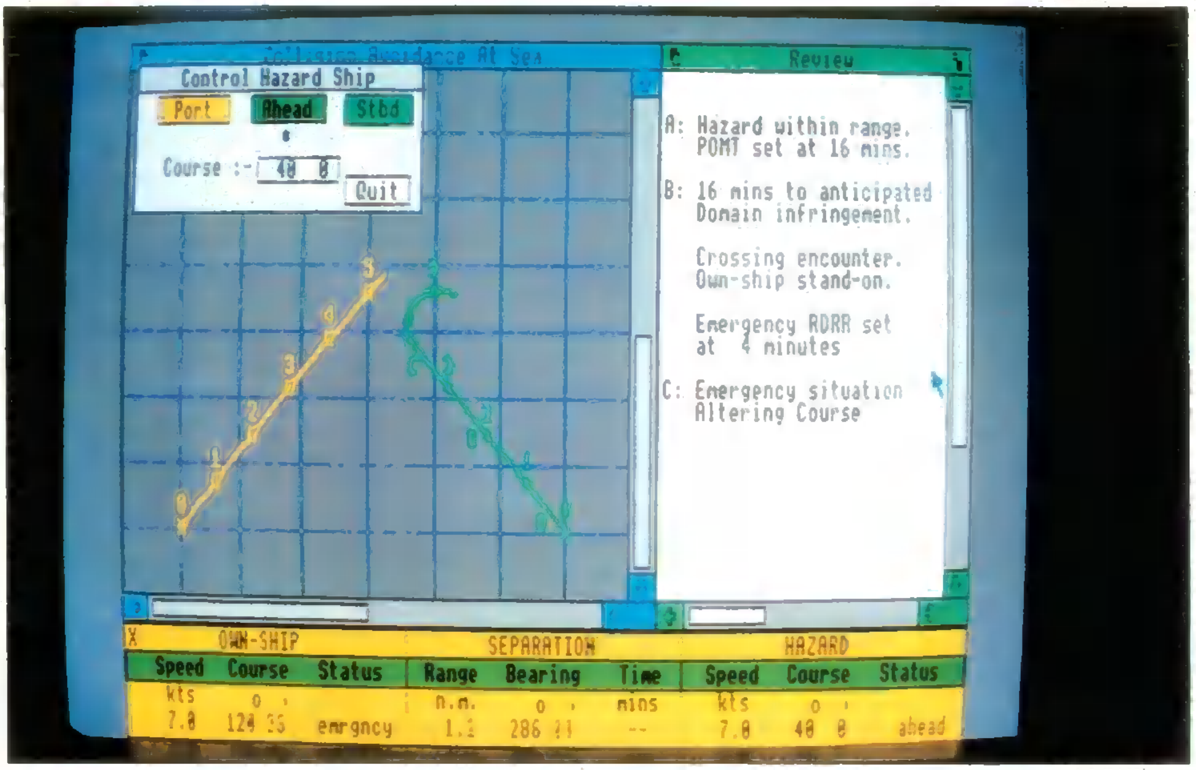

(a) Commencement of emergency manoeuvre

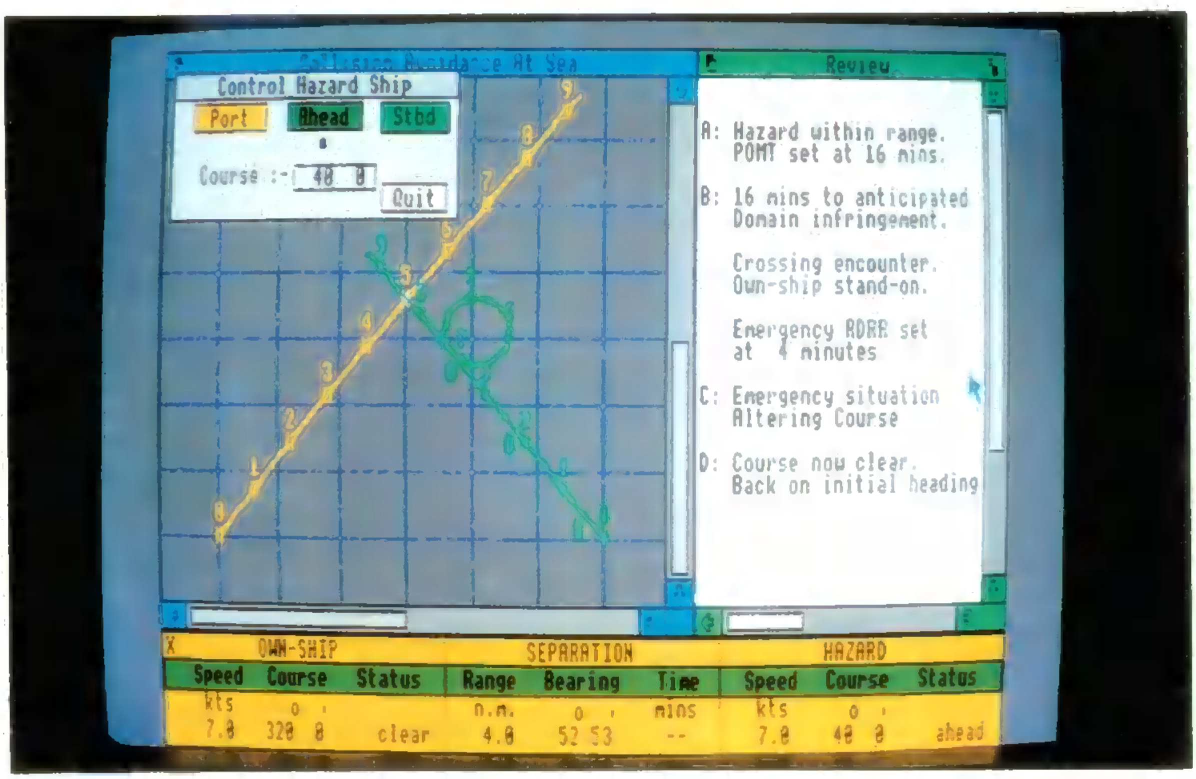

(b) Back on course, having safely completed manoeuvre

Figure 9.5 Encounter #3, as shown to both sets of mariners

Encounter #3

Parameters for both vessels as for encounter #1, but with roles reversed i.e. own-ship becomes the stand-on vessel in a crossing encounter with potential collision. The hazard vessel was controlled so as to prevent it altering to starboard (as it should), forcing own-ship to instigate emergency avoidance action at a late stage. It was explained that the only emergency manoeuvre included in the rule base at the moment was a full-turn, but that more detailed manoeuvres could (and would) be added in due course. It was also explained that, in an emergency, the system functions so as to preserve an emergency domain, in this case set at 4 cables radius; this would, in turn, determine an emergency RDRR at which time action must be taken to preserve that domain.

As indicated above, 7 mariners from Warsash and 7 from Plymouth each evaluated this manoeuvre, which was identical for both groups (figure 9.5) - the 'minimum turn' criterion being irrelevant in a full-turn situation. The response at Warsash was unanimous approval, whilst at Plymouth 6 of the 7 voted 'yes'; the one 'no' added the comment "It is an alternative, but not my preferred one".

Observations on Mariners' Responses

It is worth noting that there was a degree of 'consumer resistance' to the system from the start in both cases. This appeared to be based upon two premises:

(i) that the need for such a system in some way cast aspersions on the technical competence of experienced mariners ("None of us has ever been involved in a collision");

(ii) that such a system might threaten the livelihood of qualified mariners ("Something like this could do us out of a job").

There was also a clear intention to dispel any notion that this task could be fully understood by a non-seafarer, still less handled by some electronic device ("This gear could never do what we have to do" - this quote from the same mariner as (ii) above, within the space of around 30 seconds).

The responses of the two sets of mariners are summarised in table 9.2, which should be considered in conjunction with the foregoing analysis of those responses and associated comments. The results of this exercise support the hypothesis that this system fulfils its objectives, as a flexible working decision system capable of adaptation and/or extension according to user specifications.

Table 9.2 Mariners' Responses To Questionnaires On Demonstrations Of Simulated Encounter Situations

	Question 1						Question 2							Question 3							Question 5					
	YES		NO		ABSTAIN			YES		NO		ABSTAIN			YES		NO		ABSTAIN		YES		NO		ABSTAIN	
	no.	%	no.	%	no.	%		no.	%	no.	%	no.	%		no.	%	no.	%	no.	%	no.	%	no.	%	no.	%
Warsash (14)	14	100	0	0	0	0	(a)	11	79	1	7	2	14	(a)	10	71	2	14	2	14	12	86	1	7	1	7
							(b)	12	86	0	0	2	14	(b)	5	36	2	14	7	50						
Plymouth (11)	11	100	0	0	0	0	(a)	5	46	5	46	1	9	(a)	8	73	3	27	0	0	8	73	3	27	0	0
							(b)	4	36	4	36	3	27	(b)	5	46	5	46	1	9						
Total (25)	25	100	0	0	0	0	(a)	16	64	6	24	3	12	(a)	18	72	5	20	2	8	20	80	4	16	1	4
							(b)	16	64	4	16	5	20	(b)	10	40	7	28	8	32						

Responses To Question 4

	Encounter #1							Encounter #2							Encounter # 3						
	TOTAL	YES		NO		ABSTAIN		TOTAL	YES		NO		ABSTAIN		TOTAL	YES		NO		ABSTAIN	
		no.	%	no.	%	no.	%		no.	%	no.	%	no.	%		no.	%	no.	%	no.	%
Warsash	14	6	43	7	50	1	7	7	5	71	2	29	0	0	7	7	100	0	0	0	0
Plymouth	11	9	82	2	18	0	0	7	5	71	2	29	0	0	7	6	86	1	14	0	0
Total	25*	15	60	9	36	1	4	14*	10	71	4	29	0	0							

(* Not comparing like with like -- see section 9.2)

9.3 Evaluation of Expert System on Sea Trials

Sea trials were conducted aboard a 40 foot picket boat from Britannia Royal Naval College, Dartmouth (see appendix B). A second picket boat was employed as a target vessel.

For the purpose of these trials, it was not intended at this stage to interface the expert system to any instrumentation aboard the research vessel. Rather, the expert system was to be installed adjacent to the controls and instrumentation, but operating independently. For each specific trial encounter, own-ship and target vessel would each be given an initial position, speed and course to steer. These parameters would be fed into the expert system, which would then emulate the behaviour of both vessels as the encounter proceeded. At various points in the encounter, the expert system would issue advice on the preferred course of action for own-ship, and continue modelling the encounter on the assumption that this advice had been followed. Assuming accurate modelling of both vessels by the system, implementation of the advice given (i.e. in terms of practically controlling the research vessel) should hopefully lead to successful resolution of the encounter situation. By tracking the target vessel on the research vessel's ARPA set, a continuous comparison could be drawn between the actual progress of the encounter and the expert system's perception of events.

Arrangement of the trials was not a trivial exercise. The picket boats at BRNC Dartmouth are in use fairly consistently for training purposes; taking two of them out for such an exercise required substantial coordination between research staff, technical staff and boat-handling crew. Superimposed upon this was a requirement for good weather, essential if meaningful observations were to be made from (and of) small vessels at separations of several miles. Additionally, the simple presence of a coastline imposed constraints which had not previously required consideration: simulated trial manoeuvres had tended to be run South-to-North, without concern for fixed obstructions; trials at Dartmouth would require rather more thought if one was to avoid having a manoeuvre disrupted by an approaching land-mass. With maximum speeds of around 10 knots, setup times had also to be considered in the planning schedule, in addition to the expected duration of the encounter.

Having watched and waited out the wet and windy winter weather, a suitable 'window' presented itself in March, when conditions seemed favourable and staff and boats were available for a day. The Archimedes computer was installed beside the ARPA on picket boat 9 (fig. 9.6), and boats 9 and 5 set out from Sandquay (Dartmouth) on flat calm water under clear, sunny skies (fig. 9.7). However, on leaving the Dart estuary a substantial swell was encountered, a locally moderate/rough sea state having been left by the North-Easterly winds of the previous few days. The first noticeable consequence was a requirement to lash down the Archimedes, to prevent it spontaneously transferring from the chart table to the cabin sole.

The first planned exercise was a crossing encounter, requiring an initial separation of 3 nautical miles. Picket boat 5 was directed to head 2 n.m. due West, whilst PB 9 headed 1 n.m. due East. Unfortunately,

(a) ARPA and Archimedes aboard picket boat 9

(b) Detail of the ARPA display

Figure 9.6 Preparation for sea trials

(a) Picket boat 9 leaving Sandquay

(b) Picket boat 5 heading out of Dartmouth

Figure 9.7 Picket boats 9 and 5 under way at about 1 n.m. separation, PB 5 disappeared from the ARPA display due to wave interference, the waves being of the same order of height as the boats. Clearly, the exercise could not continue as planned, since ARPA tracking of the target vessel was not feasible. Despite rendering the system inoperable, this was not considered to be a failing of the system, but rather a limitation of the use of small boats for testing a system designed for substantially larger vessels.

Consideration was given to other methods of tracking the target vessel. In particular, a regular radio call from PB 5 of its Decca position could be compared with the GPS position recorded aboard PB 9 to give range and bearing. However, this was considered too uncertain to justify the complexities involved, since any results so gained would not give clear proof of the efficacy of the expert system. Regretfully, the trials were abandoned for the day.

The next opportunity for such sea trials did not arise until late May, at which time picket boats 9 and 5 were again available for an afternoon and weather conditions ideal. The boats were taken well clear of the land and then separated, PB 5 again heading due West of PB 9.

The first planned encounter was a crossing, with PB 9 approaching PB 5 from the port side, making PB 9 (i.e. own-ship) the give-way vessel (see fig. 9.8(a)). At a separation of 3 nautical miles PB 5 was directed to proceed at 5 knots on a heading of 140 degrees, whilst PB 9 made 5 knots on a heading of 220 degrees. This put the two vessels on a collision course, with projected domain infringement a little over 20 minutes away. The expert system was given the course and speed of each vessel, at grid positions which put them 3 n.m. apart East-West.

As the encounter proceeded, the target vessel was tracked on the ARPA, and readings compared with those given on the status line of the expert system display. As the expert system advised on course changes for the avoidance manoeuvre, that advice was translated into action at the helm. The rapid turning capability of the picket boat was not, of course, fully utilised; rather, the gradual turn indicated by the expert system was followed, emulating the slower response of a large cargo vessel.

Throughout the progress of the encounter, the information given on the two displays (ARPA and expert system) corresponded, to the accuracy given on the computer screen. PB 9 safely maintained the integrity of its specified 0.8 n.m. domain, and came back onto course precisely as planned. The trial was judged by all present to be 100% successful.

The next planned trial was an overtaking encounter, PB 9 approaching PB 5 obliquely from the starboard quarter (see fig. 9.8(b)). The vessels first separated to a distance of 1.75 n.m., PB 5 being due East of PB 9. For this exercise, separation and speeds were kept relatively small, to avoid covering excessive distances in the course of the manoeuvre. PB 5 was directed to proceed on a course of 145 degrees at 2 knots, whilst PB 9 made 5 knots on a heading of 110 degrees. These figures were again fed into the expert system, and the output from the system compared with the ARPA display as events proceeded.

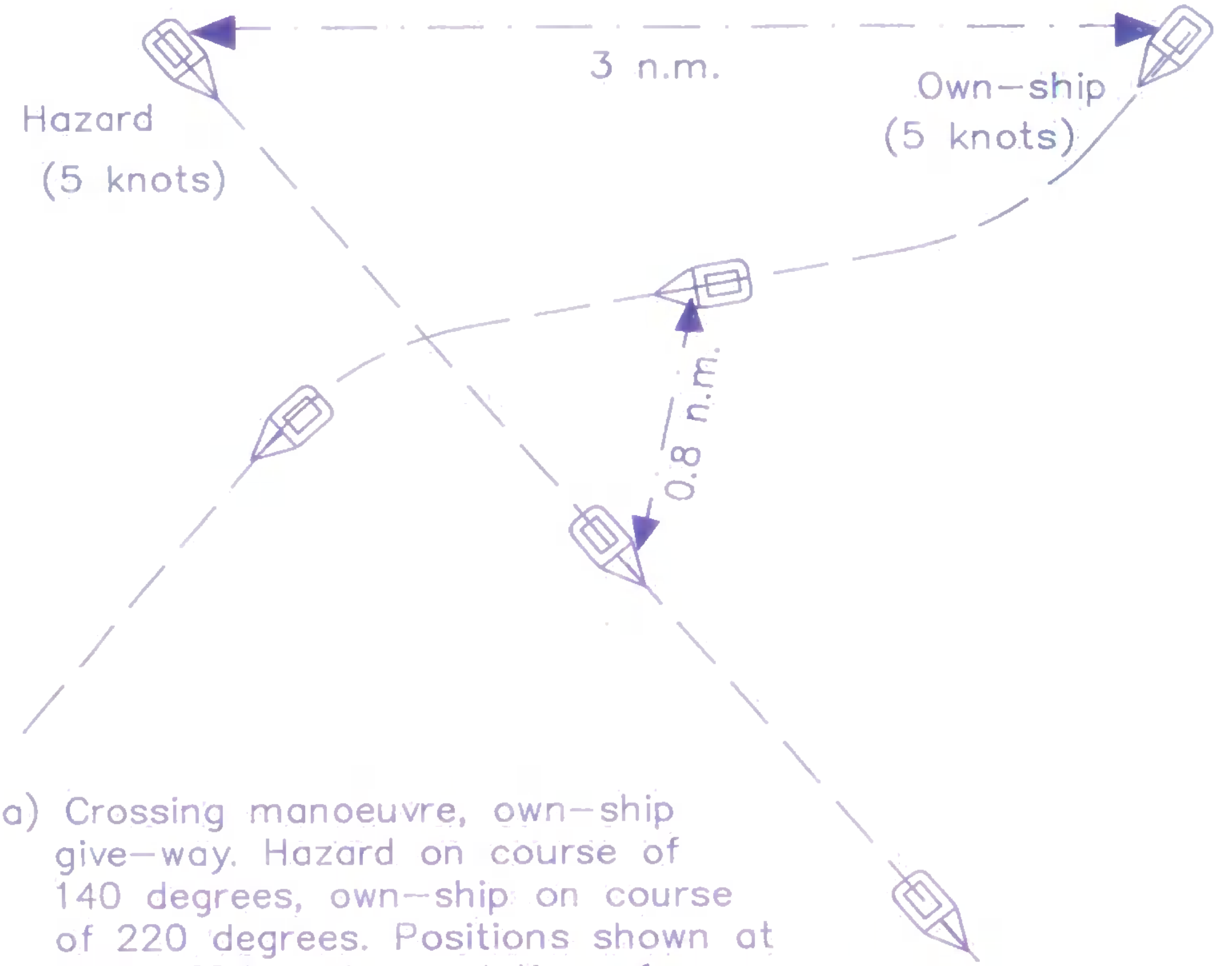

(a) Crossing manoeuvre, own–ship give–way. Hazard on course of 140 degrees, own–ship on course of 220 degrees. Positions shown at start, CPA and completion of manoeuvre.

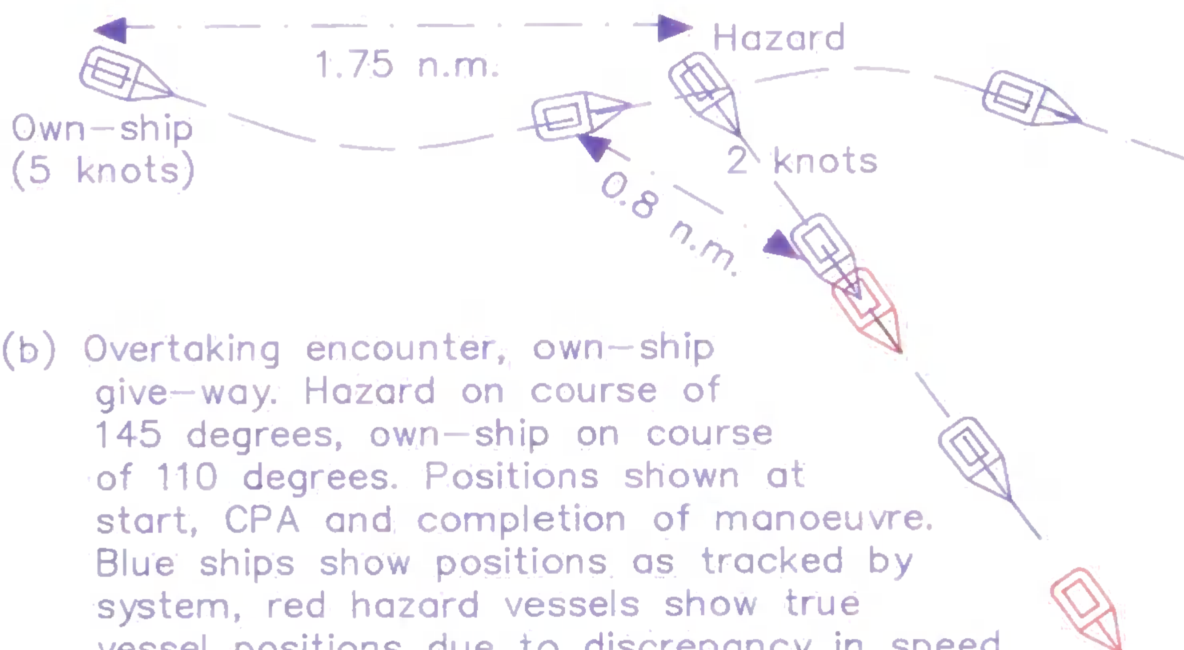

(b) Overtaking encounter, own–ship give–way. Hazard on course of 145 degrees, own–ship on course of 110 degrees. Positions shown at start, CPA and completion of manoeuvre. Blue ships show positions as tracked by system, red hazard vessels show true vessel positions due to discrepancy in speed.

Figure 9.8 Two encounters as performed on sea trials.

The encounter proceeded smoothly, and the directives issued by the expert system were followed as before. However, as PB 9 turned to port to pass behind PB 5, some discrepancy was observed in the separation distances given by the two systems. It became apparent that PB 5 was in fact moving rather faster than her intended 2 knots; as a consequence the gap between the two vessels was opening up rather faster than expected - the range at CPA was approximately 1 n.m., rather than the 0.8 n.m. shown by the computer. The master of the vessel confirmed that it was difficult to hold such a vessel at a consistent slow speed, since speed was gauged from engine revolutions; at low speeds, only one of the two engines was in use, and such a translation was less precise. Given these considerations, it was again judged by all present that the expert system had correctly diagnosed the situation and implemented appropriate evasive action. Had the speed of PB 5 not exceeded its specification, it is considered that the actual encounter would again have followed the pattern presented by the computer.

This third stage of validation has highlighted certain issues related to marine deployment of such equipment, not apparent in simulated land trials: reliance on electronic equipment, matching of ARPA display to circumstances as visually observed from the bridge, constraints imposed by coastal features, effects of weather and sea state. At the same time, it has proved the efficacy of the expert system in its intended environment - at sea, within the command centre of a marine vessel. In these respects the sea trials achieved their intended objectives.

CHAPTER 10

MULTI-SHIP ENCOUNTERS

"Believe me, my young friend, there is nothing - absolutely nothing - half so much worth doing as simply messing about in boats."

Kenneth Grahame: Wind in the Willows (1908)

The expert system described here deals with only one hazard vessel at a time. However, the ultimate requirement is a system which will 'acquire', deal with (whatever that may entail) and then ultimately 'drop' potential hazard vessels, carrying several in the knowledge base at any one time, at whatever level of detail is deemed necessary. Current research [7] is concerned with just such an approach: as any vessel comes within ARPA range, it is evaluated as a potential threat, if necessary entered in the knowledge base for detailed consideration (otherwise simply monitored for any change in circumstance); as an encounter approaches, a strategy is devised which takes account of all vessels recorded in the knowledge base; once a vessel is heading away from own-ship, and unlikely to be involved in any future manoeuvre, it is deleted from the knowledge base.

The intended filter for selection of potential hazards is an enlarged version of the Davis domain referred to in section 2.1. Figure 10.1 illustrates the principle: if a target vessel is within, or scheduled to enter, the 'super-domain', then it may pose a direct threat, or alternatively be in a position so as to affect a manoeuvre around

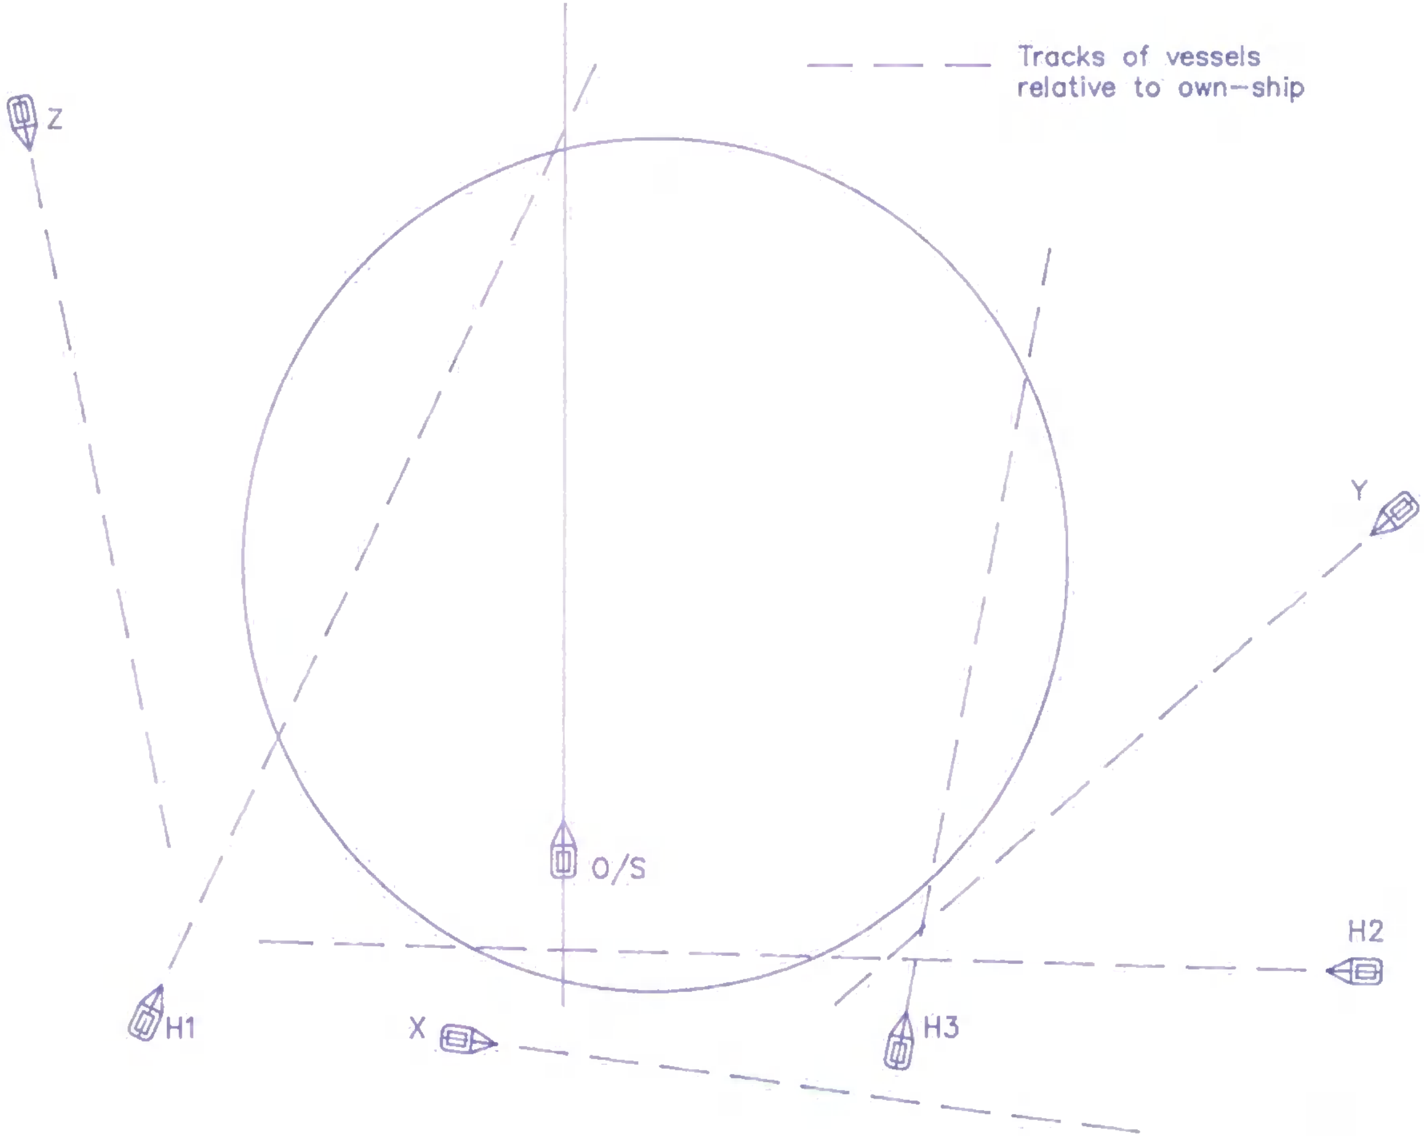

Figure 10.1 The 'Super-Domain', adapted from the Davis Domain (fig. 2.1(b)), used to identify those vessels which must be considered in relation to any avoidance manoeuvre – H1, H2, H3. Vessels whose tracks do not intersect the super-domain – X, Y, Z – are not considered such a threat.

another vessel at some future time; otherwise it is considered unlikely to be in any way relevant to future manoeuvring decisions. Once a vessel is outside the 'super-domain' and heading away from own-ship, it need no longer be considered and may therefore be deleted from the knowledge base.

The derivation of such a strategy for a multi-ship encounter leaves much scope for individual interpretation of the situation. Essentially, a multi-ship scenario may be seen as a series of two-ship encounters, a single event requiring avoidance of a number of ships by a single manoeuvre, or any variant of these two positions.

It might be possible, by consideration of separations and relative velocities, to group two or more hazard vessels for avoidance purposes, and consider them as a single unit. Such a 'super-ship' could be expected to change shape and size with time, and ultimately break up into discrete entities - each comprising one or more vessels. Likewise, vessels originally seen as individuals may at some future time 'merge' (from the point of view of avoidance strategy) into such a composite. The concept of a 'super-ship', comprising a variable number of hazard vessels is clearly very fluid and transient, not easy to tie down. Rules to accomodate such an entity may be devised at some future time.

A much simpler concept, for analysis purposes, is that of successive two-ship encounters. Here too, however, there are certain constraints and considerations which must be given due credence in any derivation of strategy. It is not enough to consider an encounter with hazard A in isolation, without regard for other traffic, then having negotiated that hazard to turn one's attention to hazard B in like manner. Clearly, one's position vis-a-vis vessels B, C, D etc is substantially affected by the way one handles the encounter with vessel A.

At the same time, a multiple look-ahead routine covering a complex series of manoeuvres through a number of hazard vessels is out of the question. For example, let us suppose that a particular strategy for a single manoeuvring stage considers each of three possible rudder settings for any of four action times - twelve options in all (see figure 10.2). It is clear that, as the number of required manoeuvring stages increases, the number of possible options rises exponentially as a power of twelve. Figure 10.3 illustrates an encounter in which just three manoeuvring stages are needed. Each stage may permit any of twelve alternative strategies, giving 12x12x12 = 1728 alternative paths through the encounter. Clearly, if each path is completely evaluated by the look-ahead optimisation routine before settling on the optimum route, the processing time required will vastly exceed the 'thinking time' available before the encounter becomes a disaster.

The solution would appear (as always) to be a compromise, believed to correspond closely to the attitude of the human master in such circumstances. The most threatening hazard (the one with shortest time to domain infringement) is given full attention with regard to an optimal manoeuvre. However, the outcome of that manoeuvre is tested (briefly, by a single attempt) for feasibility of the next stage; if the next stage is not feasible, then the option producing that outcome is

Figure 10.2

Multiple look-ahead for evaluation of POMT, showing effect of each of three rudder settings (R1, R2, R3) for each of four action times (T1, T2, T3, T4).

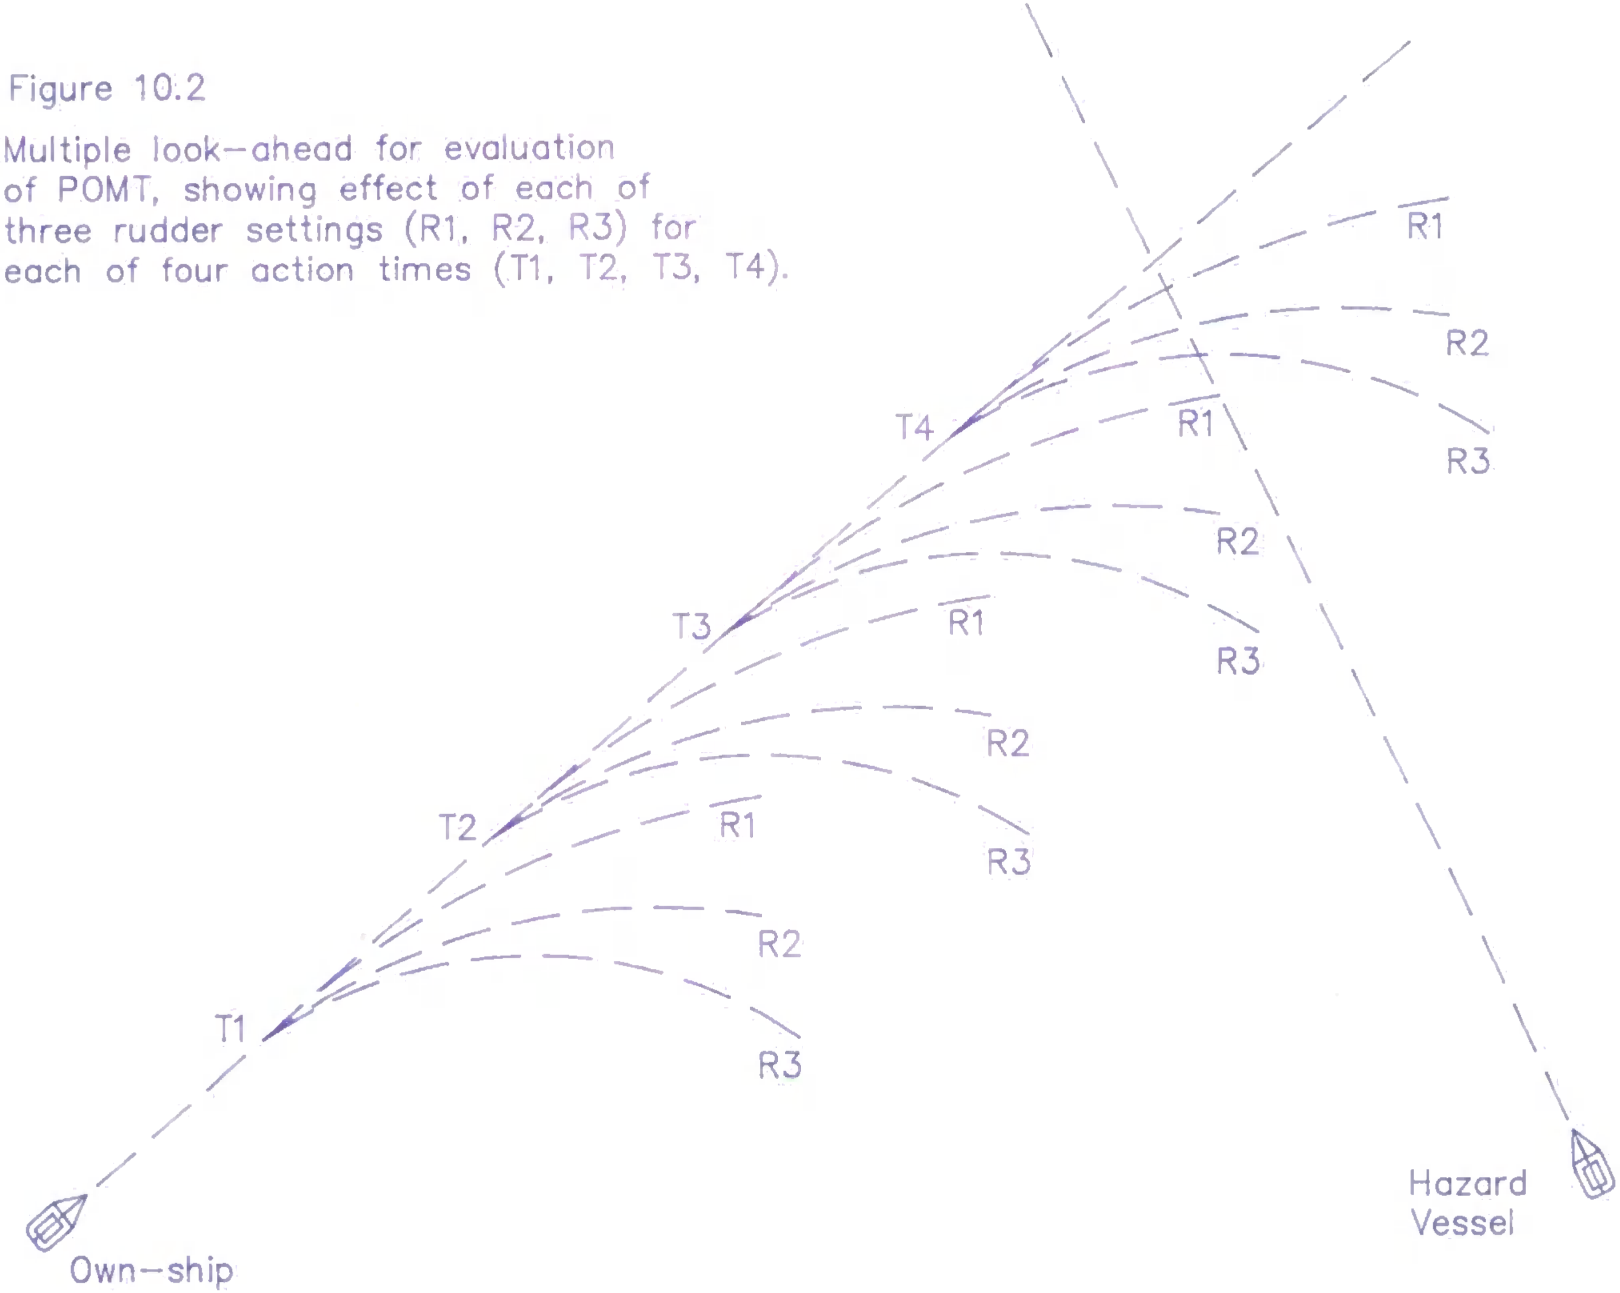

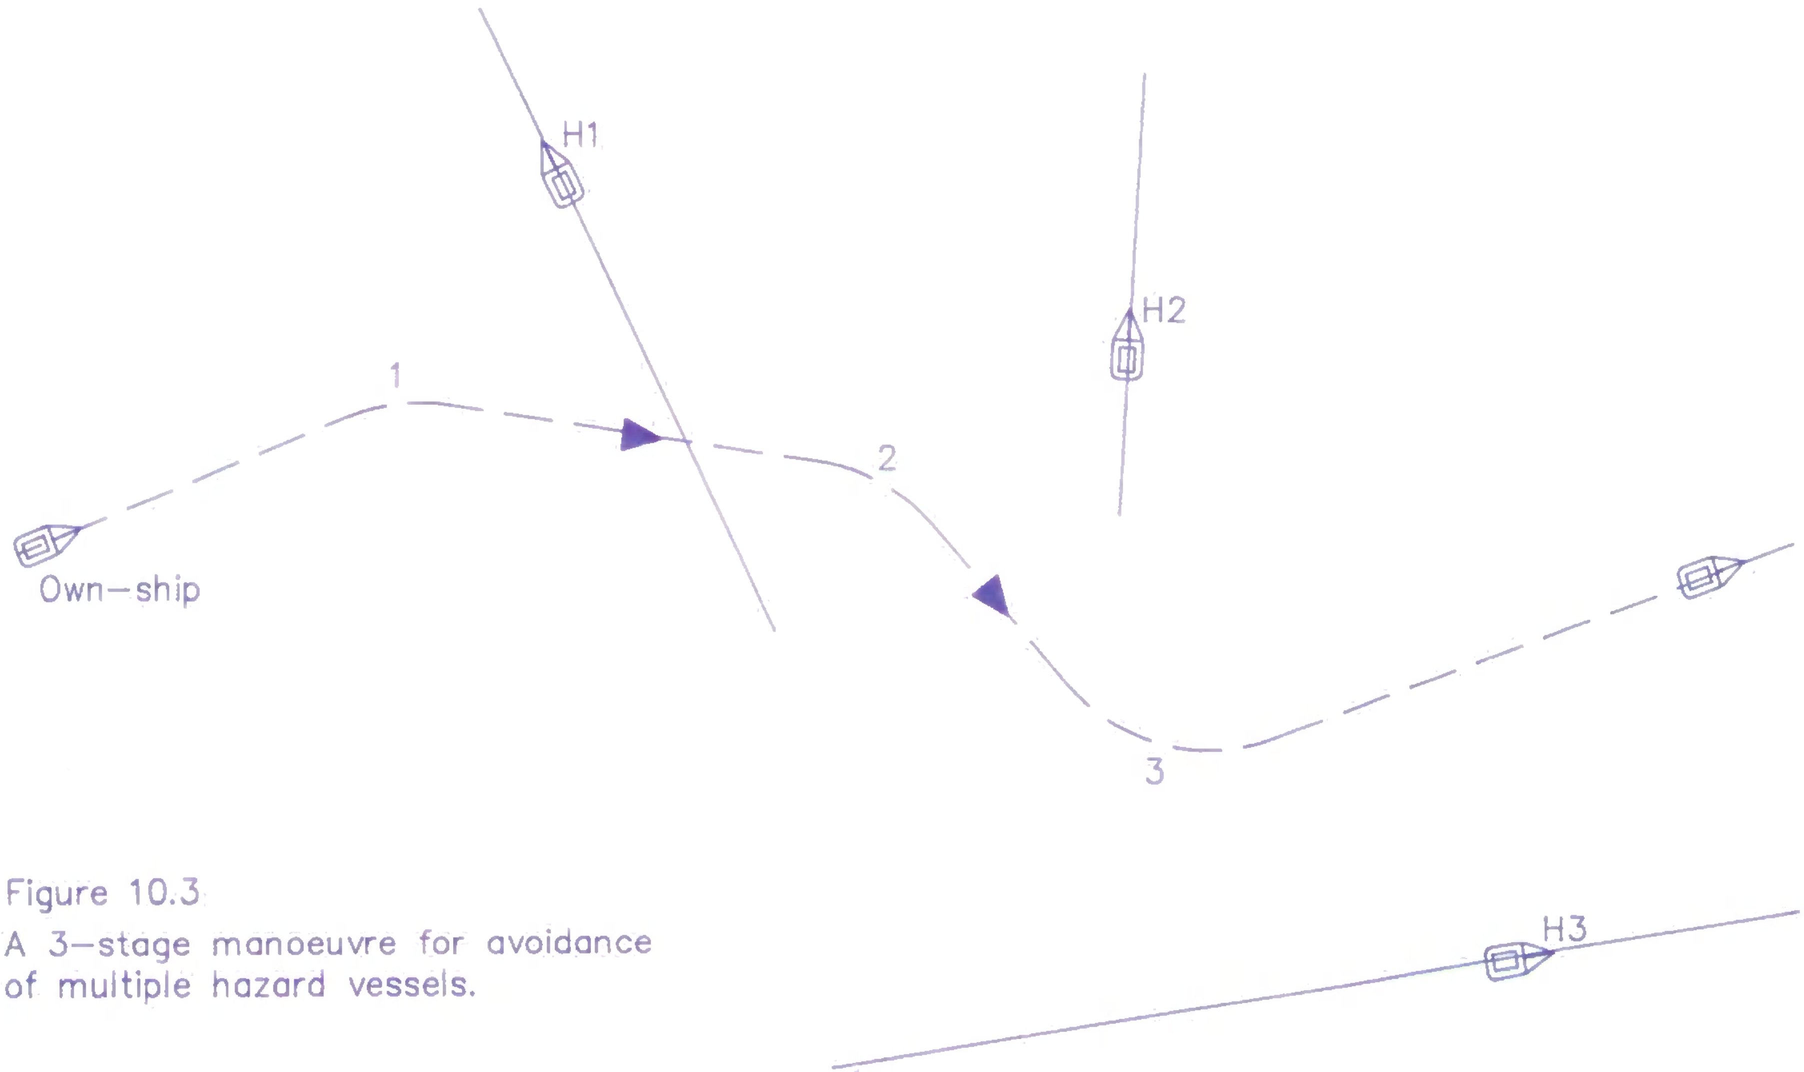

Figure 10.3

A 3-stage manoeuvre for avoidance of multiple hazard vessels.

rejected; if it is feasible, but in some way unsatisfactory (e.g. creating a new encounter situation in which another vessel is now forced to divert), then a costing is added to the optimisation cost function to reflect this.

One further refinement would be to 'merge' successive stages of a composite manoeuvre if such action seems appropriate. Figure 10.4 illustrates this point: own-ship has altered to starboard to avoid hazard 1; after a small turn, this first stage is complete, and own-ship would continue on a straight course for a time were it not for hazard 2; since avoidance of hazard 2 will, in due time, necessitate a further turn to starboard, it makes sense to continue the initial turn as shown, rather than straighten-up and then alter again shortly afterwards.

Obviously, it would not do to carry such action too far. If subsequent hazards were to force own-ship even further off course, then the initial manoeuvre could ultimately prove extremely inefficient. It would seem logical, therefore, to attach some cost to any such increased diversion, to discourage use of such a manoeuvre where a more economical alternative is available. That cost could be rather more than a simple linear function of the angle of diversion.

Since later stages in a compound manoeuvre will be tested only briefly for feasibility (as opposed to selection of an optimal manoeuvre from a wide range of options, in the first stage), the possibility must be considered of accumulated error giving a false result. For example, a manoeuvre extrapolated to five stages may appear feasible in the fifth

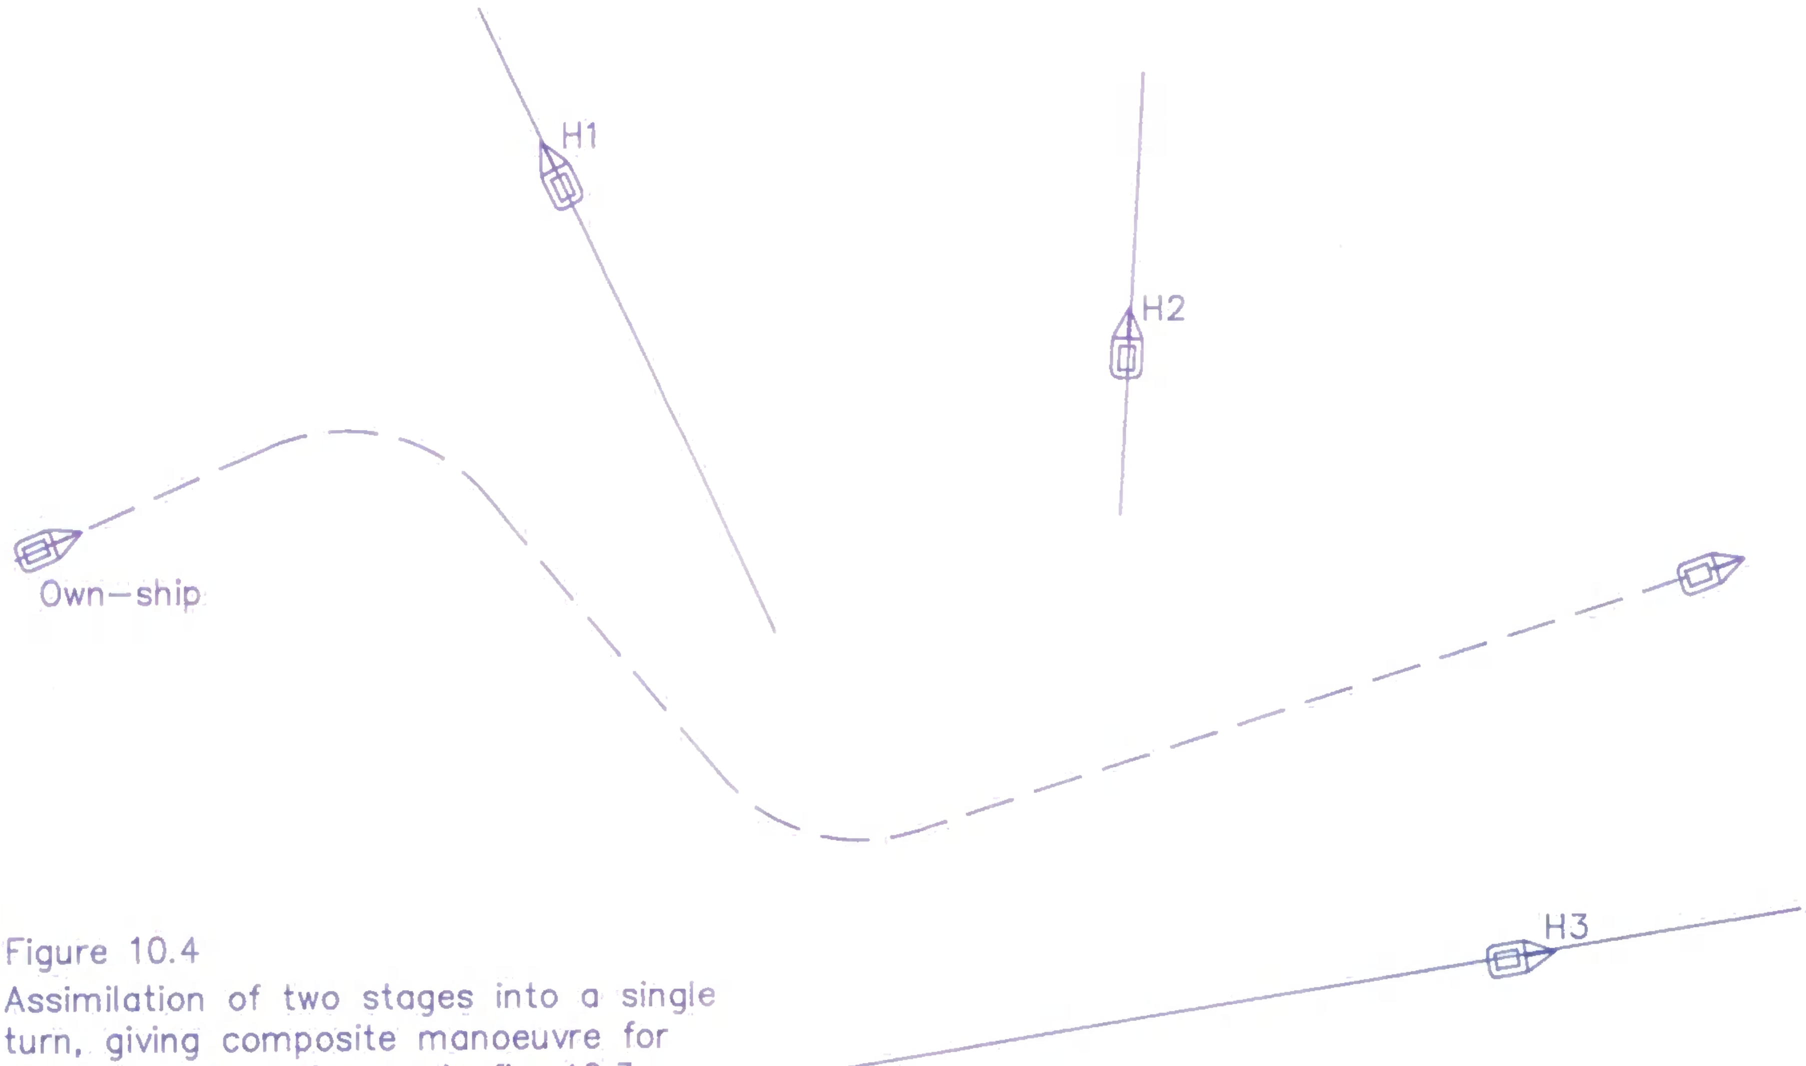

Figure 10.4
Assimilation of two stages into a single turn, giving composite manoeuvre for hazard configuration as in fig. 10.3.

stage, whereas the actual position at the end of stage four may render stage five impracticable - see figure 10.5. One way to combat this problem would be to add a percentage 'growth' to the domain size for each successive stage of extrapolation beyond the first manoeuvring stage. This added error margin should ensure that any extrapolated manoeuvring stage which appears feasible with an enlarged domain should prove to be so with a normal domain when the manoeuvre is actually carried out. This approach may, of course, disqualify a manoeuvre which would have succeeded. Since any collision avoidance strategy must offer certainty of safe completion, occasional rejection of a workable option is a price which must be paid for such certainty.

Considerations relating to multi-ship encounters are further compounded by the certain expectation that the master of any hazard vessel will not only consider our own-ship, but also any other vessels in the vicinity, when devising his own plan of action. It is therefore incumbent upon this expert system, insofar as possible, to 'second-guess' the intentions of masters of other vessels in the vicinity vis-a-vis each other, as well as in respect of our vessel. However, if taken to extremes, this can lead to recursive reasoning of the form "What does he think that I think that he is going to do?" (comparable to the deadlock at a mini-roundabout when there is a car waiting in each approach road). Common sense suggests that such logic stop at the first level, with advice from expert mariners on the likely intentions of other vessels in a variety of situations. Such expectations could be built into the rule base at some future time, possibly operating in conjunction with data on hazard vessel manoeuvring characteristics (see section 11.5).

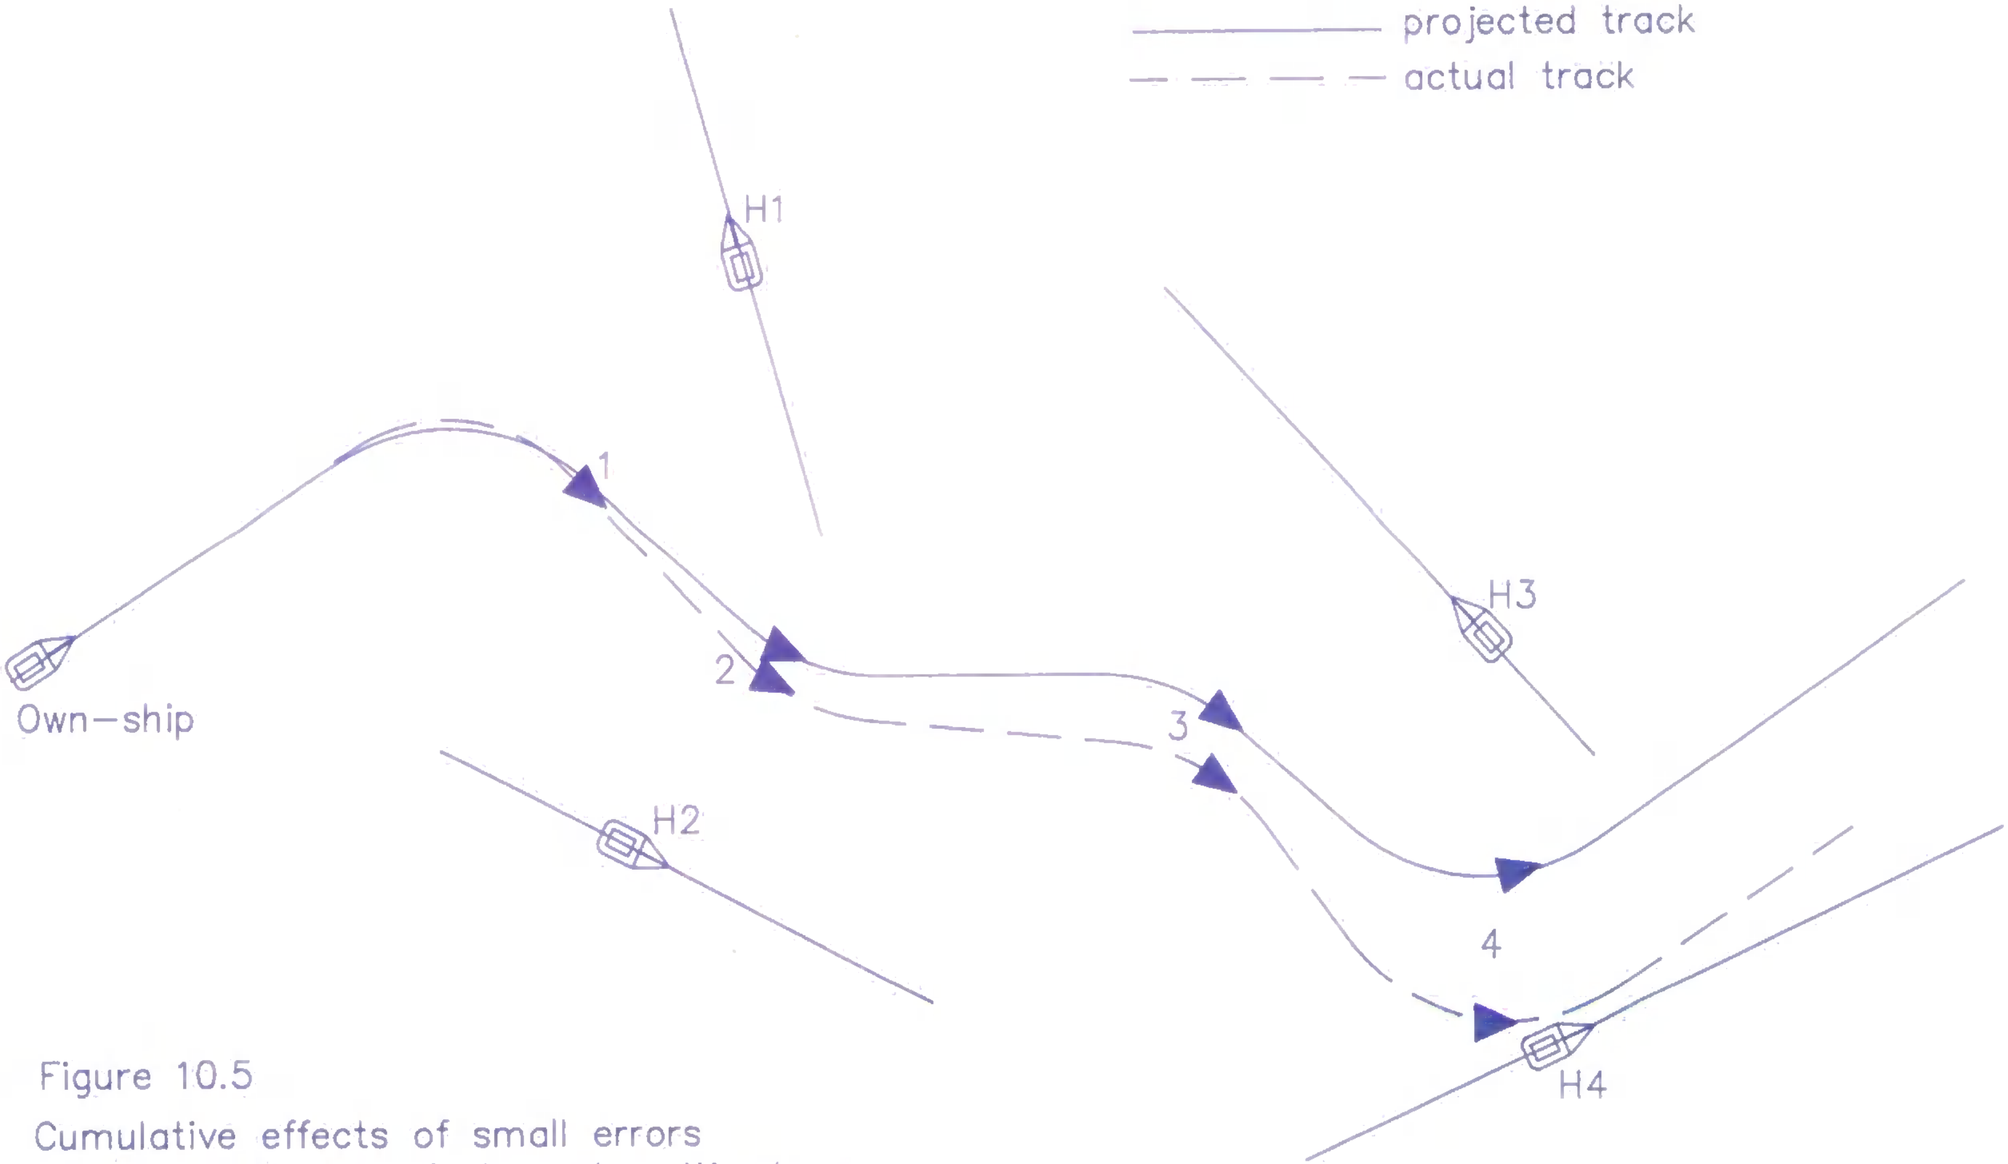

Figure 10.5

Cumulative effects of small errors in implementation of planned multi-stage manoeuvre, leading to non-viable final stage.

CHAPTER 11

CONCLUSIONS AND FUTURE DEVELOPMENTS

"The outcome of any serious research can only be to make two questions grow where only one grew before."

Thorstein Bunde Veblen (1857-1929):

The Place Of Science In Modern Civilization

11.1 Conclusions

The prime objective of this research has been to investigate the application of computer-based expert systems technology to the task of collision avoidance. This has entailed:

(a) analysis of the nature of the problem: what a collision is, different types of collision situations, circumstances under which such situations may arise, possible consequences if remedial action is not taken;

(b) definition of principles by which a situation of collision or near-miss may be recognised; specifically, forward-looking extrapolation of current circumstances can show whether there is any risk of such an event;

(c) identification of appropriate boundaries or 'safety zones' around the entity to be protected and/or other entities which constitute potential hazards; these boundaries may be distance-based or time-based; a progression of successively smaller boundaries may indicate increasing levels of immediacy in terms of required action;

(d) formulation of avoidance strategies to deal with potential collision or near-miss situations, and principles whereby the appropriate strategy may be selected for a specific situation;

(e) consideration of suitable forms of communication between the expert system, the environment in which it is operating, and the user of that system.

For reasons detailed in chapter 1, the maritime setting was chosen as a suitable field for exploration of these concepts. Consequently the main thrust of this research has been directed towards the formulation and implementation of a computer-based system for intelligent marine collision avoidance. Whilst it was not expected that this work would deliver, as an end-product, a fully viable marine collision avoidance system, it was intended that the resulting system would embody all the concepts necessary for the ultimate perfection of such a tool.

Through the duration of this research, the developing system has gone through various stages on two different computer systems. Since each stage had been anticipated, in outline, at the previous stage (and therefore, working back iteratively, in embryo at least in the initial design phase), the transition from each stage to the next has been one of enhancement, rather than complete restructuring. This concept of 'upward compatability' is a major tenet of the present system, offering ease of assimilation of new concepts - including those detailed in chapter 10 and the later sections of this chapter.

It has not been possible, or appropriate, to work methodically through stages (a) - (e) (as detailed above) as successive steps in the development of the required system. The process has been more generative in nature, all of these stages being advanced in parallel, each contributing its part to the overall concept. In fact, the element which received most attention at the outset was item (e), the user interface. It was felt that the nature of the communication between system and user was a consideration which could, and should, be planned fairly thoroughly at the outset, without regard for the details of decision-making and manoeuvring. As a consequence, the 3-window display format has remained much the same throughout; use of dialog boxes, pop-up menus and mouse-driven operations have added sophistication as the system has progressed, in accordance with the original plans.

In principle it has not been difficult to identify the problem. Shipping statistics show a consistent record, and regular newspaper reports carry a grim reminder, that collisions at sea are an actuality and there is much room for improvement in the safety record of marine transport. Having said this, the reasons for each particular disaster are often the subject of valid differences of opinion and it ill behoves a non-master-mariner to be wise after the event. However, certain underlying concerns may be recognised, and hopefully addressed in the formulation of a computer-based system. These centre on:

(a) the wide range of information at times having to be assimilated by one person on the bridge of a ship;

(b) the non-specific wording of certain key phrases in the collision avoidance regulations, necessitating subjective interpretation by mariners;

(c) the occasional use of unqualified, or otherwise incompetent, staff as duty officers on marine vessels.

Consequent on these observations, it followed logically that any computer-based system should:

(a) provide pertinent and timely information in an easily-digestible form at the point of need;

(b) apply a well-reasoned interpretation of the collision avoidance regulations, which would not lead to uncertainty in the minds of masters of other vessels or make rash assumptions regarding the intentions of those masters;

(c) give clear visual (and possibly audible) warning, in good time, of impending danger, and equally clear advice on how the situation could best be safely resolved (this could ultimately lead to the marine equivalent of the 'dead man's handle', where the system would control the vessel through a potential disaster situation if no action is taken by the Officer of the Watch).

The first of these observations has determined the nature of the information displayed in each of the three windows on the computer screen: brief, to the point, layout and use of colour carefully chosen to communicate key facts quickly and easily. The second has proved a recurring theme throughout the work, requiring a flexible yet unambiguous response to every conceivable combination of circumstances. Successive levels of sophistication of the system rule base have been primarily concerned with formalising the observed interpretation of the regulations by experienced mariners, in a manner which will ensure satisfactory resolution of any encounter situation. The third consideration is again encompassed in the form of the advice given in the 'review' window of the system display - clear, concise statements of action to be taken. At present, this information is selected as a user option; however, a variant of the display routine (created for the validation trials) shows this advice as it is generated by the expert system logic. Addition of an audible alarm, to sound shortly before action is to be taken, would be a very simple task.

One major achievement of this research project is the design and implementation of the expert system 'shell'. This composite structure, combining time-based inferential logic, an open-ended framework for a (similarly time-based) hierarchical system of 'rules' (or decision nodes), and an advanced user interface, has proved well adequate for the sample application undertaken. This shell could also prove appropriate for other collision avoidance situations, e.g. use of autonomous robotic devices in hazardous environments (fire, radiation, chemical hazards). Naturally, an appropriate set of rules would have to be defined for any such application, to be operated by this shell.

Within the operation of this shell, it has been possible to develop another radical concept, that of look-ahead simulation. The principle of simulation within the expert system, as a decision-making aid, has proved a necessary and sufficient tool for the successful resolution of encounter situations. This technique is considered to be an essential feature of any efficient collision avoidance system; as such, its development must be regarded as a substantial advance in this field.

The three phases of validation detailed in chapter 9 between them constitute a thorough examination of the extent to which the objectives of this research have been achieved. The nature of the presentation of information, and user interaction, has been deemed satisfactory by experts - who have also given constructive suggestions as to the form of wording in certain cases; over a very substantial number of randomly-generated encounter simulations, avoidance strategies implemented by the expert system have proved effective in every case; and the system has shown itself to be seaworthy. These results constitute a complete vindication of the approach taken. Given the expansion potential inherent in the structure of the system, they support the expectation that this research forms the basis of a viable fully operational expert system for collision avoidance.

The remaining sections of this chapter give consideration to possible enhancements of the system, as suggested earlier in the work.

11.2. Ongoing Graphical Display of Ships' Tracks

The orientation of the graphic display has been the subject of a great deal of thought: head-up or North-up? relative or absolute motion shown? if absolute, what may be done about the problem of own-ship ultimately disappearing off the edge of its own display?

A head-up display (i.e. top of screen = direction of travel for own-ship) brings with it attendant problems of rotation of the whole display as own-ship manoeuvres. Whilst technically feasible (if complicated), it seems likely that such rotation might lead to disorientation for the user. North-up has been selected as the preferred option (user selection of either option, as with some radars, may be a later enhancement).

A display of relative motion of hazards would enable own-ship to be shown static at the centre of the display. However, relative tracks present a seriously distorted image of the train of events, inhibiting an objective appraisal of the situation. It was considered that a 'true motion' display would aid in assessment of the relevant factors in the encounter.

The current version of the expert system shows true tracks of vessels, including own-ship tracking across (and ultimately off the side of) the 'tracks' window. This approach has proved adequate for the limited time period of one simulated encounter, but would rapidly fail - with the disappearance of own-ship off the side of the display - on a voyage of any duration. Various options have therefore been considered to overcome this limitation.

One possibility is to display true tracks of vessels, but continuously scroll the display as necessary to hold own-ship at the centre of the window at all times. This would effectively 'grow' the past track of own-ship outward from the centre of the window, but would also have the disconcerting effect of relocating every other feature displayed (hazard track(s), grid, possible future developments such as chart features) at each time-step. Such a regular blurring of the display would undoubtedly seriously reduce its capacity for clear communication of relevant detail.

A second option, frequently used in other computer applications, is that of 'full screen wraparound'. In this case, a ship track leaving the area at the top of the screen would reappear at the corresponding position at the bottom, and vice versa; likewise, a track crossing the left or right edge of the area would reappear on the other side.

Although preventing the loss of a ship's track off the edge, this method could not be seriously considered to aid in the clarity of presentation. Discontinuity of a track over an edge, coupled with the problems of a 'wrapped-around' track possibly using the same window area as a previous track section (i.e. the same window area representing two different regions of sea) render such a solution unacceptable. Even the erasure of track sections more than, say, 20 minutes old does not satisfactorily address these limitations.

The third, and most satisfactory, solution is a variant of the continuous scrolling described as the first option. In this, the 'rolling road' concept is applied at intervals to ensure that own-ship is always within the centre portion of the screen - i.e. never allowed to stray into the 25% margins of the graphics area.

Figure 11.1 illustrates how this principle may be put into effect. Initially, own-ship is centre screen. As the ship's passage proceeds, its track will naturally move towards the 'margins' of the display area (not marked on screen) which constitute the outer 25% of that area on each edge. At the same time, the movement of any other ship(s) in the vicinity will be shown in the appropriate relative positions - all track lines showing <u>true</u> motion.

As own-ship reaches a 'margin', the whole of the display will be translated as necessary to put own-ship back at the centre of the display. Thus track of own-ship, track(s) of any hazard vessel(s) and any other displayed features (e.g. detail from ECDIS) will be bodily shifted such that shape, orientation and relative position are maintained, but own-ship has substantial display area all around for further movement.

This approach will, of course, mean that the track of a hazard vessel, or some other feature, may be shifted out of the screen display area

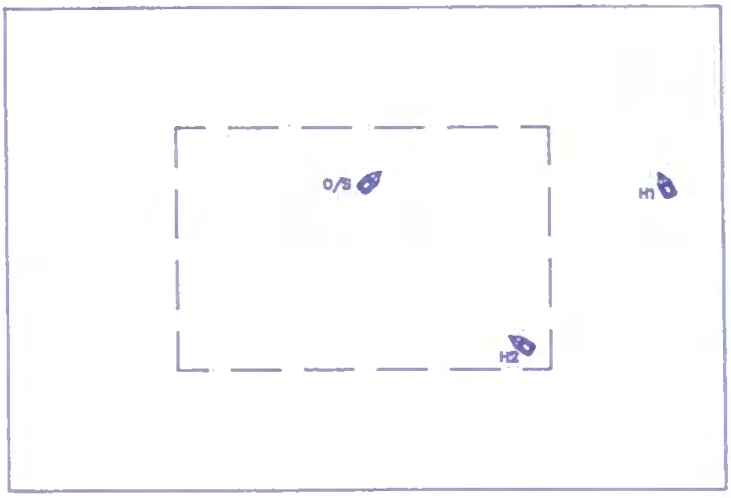

(a) Own–ship approaching top edge of boundary region (shown dotted here, but unmarked on screen). Two hazards, H1 and H2, also appear on the display.

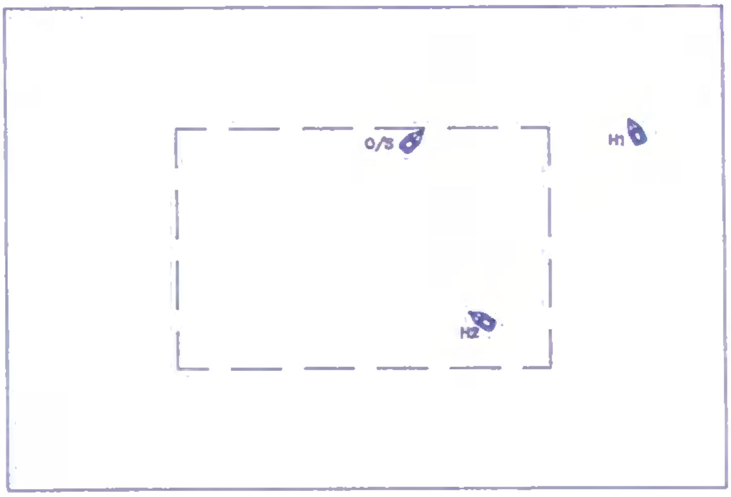

(b) Own–ship reaches upper boundary. Hazard ships have also moved on to new positions.

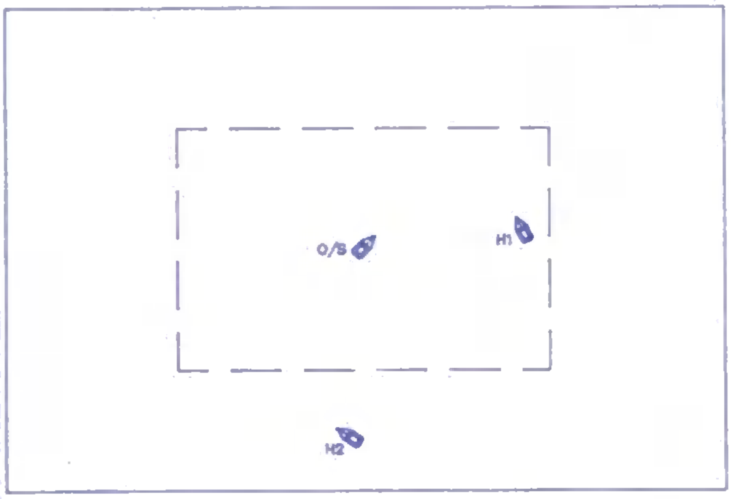

(c) On crossing the boundary, own–ship has been re–located to centre of the display area. Both hazard vessels have been moved from their new positions to maintain positioning in relation to own–ship. (Headings are unchanged).

Figure 11.1 Three stages of "rolling sea" display, showing repositioning when own–ship reaches (non–displayed) boundary.

altogether. This is in keeping with the requirements of the system, since such a policy ensures that all vessels within fairly close proximity to own-ship will be preserved on the screen display, and these are the objects of immediate concern. There are certain details which require further consideration - notably the discontinuity in track of a vessel which has 'fallen off' the display shortly before such a move, and is brought back onto screen by this process. Various possibilities suggest themselves for further investigation at the appropriate time.

The above outline of the 'rolling sea' approach to the ongoing display contains a deliberate implicit ambiguity. The 'screen display area' referred to may be taken to mean either the window area showing part of the total graphics area at any time, or the total graphics area itself; except in the reduced-scale mode, these two areas are not identical. It is probable that the 'rolling sea' approach would maintain own-ship within margins in the current window display area, rather than the total graphics area; however, this also would be given further thought at the trial implementation stage.

11.3. Extension of Review Text Information

The time-based decision process implemented by this expert system brings with it attendant complications regarding the review, or rationale of that decision process. On the one hand, the user requires a brief, clear statement of events leading to the current recommended strategy; on the other, full consideration of circumstances surrounding any encounter, set in their temporal context, will require rather more than the pithy texts appropriate to critical decision-making at times of stress.

The solution envisaged is a two-tier review facility. At the first level (already implemented) a short statement gives the essence of each decision stage, as previously described. A further refinement, yet to be included, will provide more information on each decision process. This latter text will be sent to the right-hand-side of the review window - not displayed in the normal contracted form of the window. The user will thus have the 'brief' review on display at all times (once invoked by menu selection), but will also have the option of studying the decision steps in more detail either by scrolling the review window horizontally (using the horizontal scroll-bar control) or (preferably) by expanding that window to full-screen width. This will obliterate the graphics window for the duration, but it is not anticipated that this facility would be exploited at times when the 'tracks' window is urgently needed. The 'tracks' window will, of course, reappear when the review window is contracted back to condensed form.

The space available on-screen in the 'review' window is adequate for the text generated in a single two-ship encounter. However, as the system deals with more complex encounters, or a succession of encounters, the text will automatically scroll upward, showing the text for the most recent steps at any time. Text for previous stages, or previous encounters, may be recalled by the simple expedient of scrolling upwards using the vertical scroll-bar control; scrolling downwards reverses the process.

It should be immediately apparent that the 'review window' text object holds within its memory space all of the review text generated in the course of a voyage, or simulation. By contrast, the 'rolling sea' version of the graphics display (see section 11.2) makes no pretentions to preserving all of the graphics information - any which 'drops off the edges' is lost. The alternative, conservation of all graphics for the voyage/simulation, raises a number of complications - only one of which is the amount of memory space required!

The secondary issue which arises from this is the continued existence of (letter-marked) text for which no corresponding (letter-marked) track display exists - once the process of track-window relocation has obliterated part or all of any encounter. This is not considered a problem, and the lettering should be self-explanatory. The review text essentially forms a log of the 'thought processes' of the system, with some information on the situation which engendered those thought processes. There is, of course, ample scope for reconsideration of the detail of these proposals as the intended developments are undertaken.

11.4. On-Line Updating Of The Rule Base

It is anticipated that the rule base will continue to evolve, incorporating an increasing number of decision nodes and alternative paths to accommodate a wider variety of circumstances. However, as indicated in section 3.4, updating of the rule structure is a task which requires much thought and great care. Any enhancement of the rule base must be:

(a) thoroughly planned, with due consideration of all possible consequences;

(b) carefully implemented, with strict attention to maintaining the integrity of the existing rule structure;

(c) meticulously tested in a simulated environment, before use in practice at sea.

Clearly, there is no question of updating the rule base in a 'live' situation (i.e. at sea). At present, any new rule is coded into the appropriate format and keyed into the 'rules' file by hand; any new boolean function is likewise added by hand, and links from existing rules adjusted accordingly.

Many variants exist of a simple 'learning' program, which has been around since the early days of computers, based on the binary tree structure. This program offers to identify an object, creature or person through the user answering a series of questions with 'yes/no' answers. If the program is incorrect in its identification of the object/creature /person, it follows that the system is not cognisant of that entity. It therefore requests that the user give the name of the entity they have in mind, and a 'yes/no' question which could be used to distinguish it from the incorrect response given. The system then incorporates this 'learned' information into its data file for future reference, modifying the links in the existing tree structure to incorporate the new question and additional entity.

The concept of updating a binary tree structure, stored in the form of a linked list, is not new, nor does it involve substantial difficulties in respect of programming. The technical requirement for the task envisaged is simply an extension of a well-established principle. In addition, of course, it is necessary to formulate, identify (label) and incorporate any new boolean function(s) which may be required.

It would appear totally feasible, in principle, to achieve all of this through the use of dialog boxes. One such dialog box could be used to scan through the 'rules' tree structure, looking at various nodes and identifying the break point at which a new rule should be inserted. That same box could facilitate the modification of appropriate links, signalling any inconsistencies in the proposed new arrangement. Another dialog box (or possibly the same box) could enable the input of a boolean decision function associated with this new rule; this function may refer to various standard parameters (speed, heading, relative bearing of hazard, etc) in any mathematical combination. It should be possible to adopt some form of stylised notation for such functions, to simplify their construction; mouse-selection of parameters from a list of those available would eliminate any potential errors due to mis-spelling or incorrect usage.

It might also be appropriate, as a final measure, to provide a facility to 'single-step' through the rule base. Used in conjunction with a facility to set-up any required test situation, this diagnostic aid would enable careful analysis of the consequences of any modifications. Such a diagnostic could well be supported by the provision of a printed 2-dimensional schematic of the new rule base.

11.5. Manoeuvring Models For Hazard Vessels

As indicated in section 4.4, it is anticipated that benefits may be derived from holding a 'library' of standard ship types, with their associated manoeuvring characteristics. It should then be possible, in some circumstances, for the system operator to identify a potential hazard vessel as fitting one of those types; the characteristics of that vessel could then be incorporated by the system logic into the decision-making process. Specific vessel types (such as dredgers and fishing vessels) are given special consideration under the collision avoidance regulations; the system rule base could be extended to acknowledge such priority treatment. More generally, an understanding of the manoeuvring capability of a hazard vessel would allow more informed decisions as to the timing and action required in an emergency manoeuvre.

A facility could be provided whereby a dialog box may be invoked by the operator, listing the various vessel types for which manoeuvring details are held. The user could then associate any one of those types with any hazard vessel currently under observation, allowing the system to apply the refined 'judgement' described above (this presupposes a default option - standard set of characteristics - to be used by the system in the absence of any such user intervention).

A further stage of refinement, where the system may automatically identify a vessel type from its observed behaviour, may be a possibility for the future, but is too speculative to be seriously considered as a possible development at the current time.

11.6. Coastal and Navigational Features

Section 4.3 considers in some detail the informational requirements of a computer-based collision avoidance system in respect of navigational chart data. The provision of such information, in a usable form, is outside the scope of this research project. However, inclusion of chart data in the decision-making process is a major consideration for such a system, and one which has been given much thought in this work. It should be noted that:

(a) no computer-based collision avoidance system can be fully effective in coastal waters without recourse to a well-structured electronic digital chart database;

(b) in the continuation of this work, every effort is being made to provide constructive input to the formulation of standards and structure for such a database;

(c) once such a database is available, the links are already in place within this expert system to incorporate such data into the decision process.

REFERENCES

[1] Cockcroft A N 1986 — Factors affecting shipping losses by collision and wreck. Journal of Navigation volume 39, 372-376

[2] Habberley J S, Shaddick C A and Taylor D H 1986 — A behavioural study of the collision avoidance task in bridge watchkeeping. Report compiled for Marine Directorate, Department of Transport January 1986

[3] International Marine Organisation 1981 — International regulations for preventing collisions at sea.

[4] Durham A 1985 — A dialogue that's past the small talk. Computing, October 24th 1985, 4-6

[5] Blackwell G K and Stockel C T 1989 — Simulation of ship encounters using multiple linked intelligent systems. Proceedings of the 1989 Summer Computer Simulation Conference, Society for Computer Simulation, Austin (Texas) July 1989, 766-771

[6] Blackwell G K, Rangachari J and Stockel C T 1991 — An intelligent interactive environment for a maritime real-time expert system. Proceedings of the 1991 Summer Computer Simulation Conference, Society for Computer Simulation, Baltimore (Maryland), July 1991, 358-363

[7] Blackwell G K, Rangachari J and Stockel C T 1991 — Ships that pass in the night. Intelligent management of a 'fuzzy' event stream. Proceedings of the 1991 European Simulation Multiconference, Society for Computer Simulation, Copenhagen (Denmark), June 1991, 542-547

[8] Goodwin E M 1975 — A statistical study of ship domains. Journal of Navigation volume 28, 328-336

[9] Davis P V 1981 — Computer modelling of marine traffic behaviour. PhD Thesis, CNAA (Plymouth Polytechnic), 1981

[10] Davis P V, Dove M J and Stockel C T 1980 — A computer simulation of marine traffic using domains and arenas. Journal of Navigation volume 33, 215-222

[11] Colley B A, Curtis R G and Stockel C T, 1983 — Manoeuvring times, domains and arenas. Journal of Navigation volume 36, 324-328

[12] Konyn M 1987 — Collision avoidance at sea and on a marine simulator using automatic encounter detection techniques. PhD thesis, CNAA (Polytechnic of North London), 1987

[13] Cahill R A 1983 — A fresh look at the Andrea Doria/Stockholm collision. Collisions and their causes, Fairplay Publications, London, 1983, 1-9

[14] Chudley J 1991 — An investigation into integrated navigation systems. Transfer report MPhil-PhD, Polytechnic South West (Plymouth), 1991

[15] Chudley J, Dove M J and Tapp N J 1989 — A review of ship simulator models, past, present and future. Proceedings of the Third European Simulation Congress, United Kingdom Simulation Council, Edinburgh (Scotland), 1989, 449-456

[16] Tapp N J 1989 — A non-dimensional mathematical model for use in marine simulators. MPhil thesis, CNAA (Plymouth Polytechnic), 1989

[17] Dove M J 1984 — A digital filter/estimator for the control of large ships in confined waters. PhD thesis, CNAA (Plymouth Polytechnic), 1984

[18] Penrose R 1990 — The emperor's new mind. Vintage Press, London, 1990, 14-30

[19] Schildt H 1987 — Artificial Intelligence using C. Osborne McGraw-Hill, New York, 1987

[20] Motte R and Calvert S 1988 — Operational considerations and constraints in ship-based weather routeing procedures. Journal of Navigation volume 41, 417-433

[21] Smeaton G P, Dineley W O, Fawcett S and Busby J 1990 — The electronic chart - an interactive video approach. Proceedings of the 7th Biennial International Symposium of the Hydrographic Society (Hydro 90), December 1990, Southampton, 22:1-22:11

[22] Jakobsen P A, Kyrkjeeide A and Sandvik R 1990 — The Seatrans Project. Proceedings of the 7th Biennial International Symposium of the Hydrographic Society (Hydro 90), December 1990, Southampton, 24:1-24:10

[23] International Hydrographic Bureau 1990 — Provisional specifications for chart content and display of ECDIS (special publication no. 52), Monaco, 1990

[24] Atkinson H L and Stockel C T 1992 — Charting the way ahead. Proceedings of 7th International Symposium on Vessel Traffic Services, Canadian Coastguard, Vancouver (British Colombia), June 1992

[25] Cahill R A 1983 — Interaction. Collisions and their causes, Fairplay Publications, London, 1983, 117-119

[26] Cahill R A 1983 — Converging overtaking/crossing. Collisions and their causes, Fairplay Publications, London, 1983, 78-91

[27] Colley B A 1985 — Computer simulation of marine traffic systems. PhD thesis, CNAA (Plymouth Polytechnic), 1985

[28] Acorn Computers 1989 — RISC OS programmers reference manual. Cambridge 1989

Appendix A

MATHEMATICAL CONSIDERATIONS IN AVOIDANCE MANOEUVRING

"Let no one ignorant of mathematics enter here."

Plato (429 - 347 BC)

Inscription written over the entrance to the Academy.

Biographical Encyclopedia (I.Asimov).

A.1. Time-to-go Before Projected Domain Infringement

Figure A.1 shows a target vessel at the centre of a domain, radius D. Another vessel (own-ship) is approaching on a relative heading indicated by the green line, marked C; C is the effective closing distance on that heading. This line makes an angle V with the displacement vector, R, whose length represents the current range between the two vessels.

If $V > 90$ degrees

or $D > R * \sin(V)$

then own-ship's relative heading does not intersect the domain around the target, i.e. there is no risk of domain infringement.

Otherwise, time-to-go = C / S

where S is speed of own-ship relative to hazard.

Figure A.1 Assessment of likelihood of domain infringement and calculation of time-to-go to infringement.

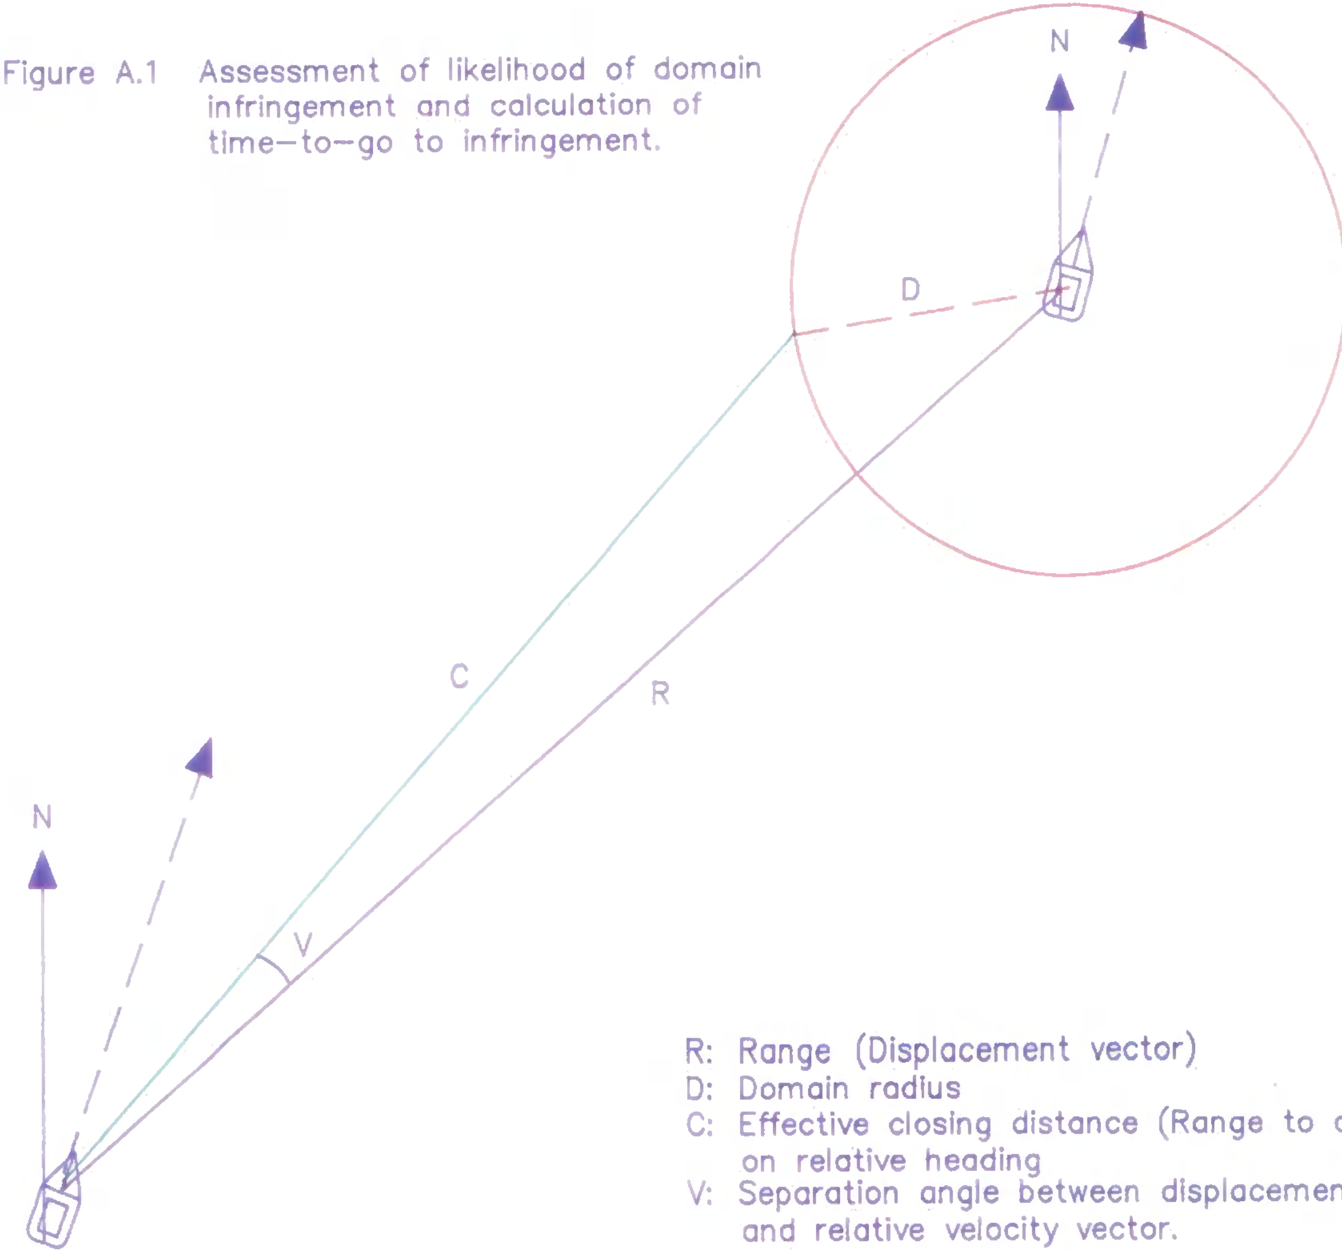

R: Range (Displacement vector)
D: Domain radius
C: Effective closing distance (Range to domain) on relative heading
V: Separation angle between displacement vector and relative velocity vector.

A.2. Overtaking From Starboard Side

Figure A.2 illustrates a situation in which own-ship is approaching a target vessel from an overtaking position on the starboard quarter. The action taken will depend on various mathematical considerations.

(i) (a) Is own-ship set to pass ahead initially?

(I.e. does V cut T2 ahead of hazard vessel?)

(b) Are the two vessels on moderately near-parallel courses?

(I.e. is $10 \text{ degrees} < \text{angle } A < 30 \text{ degrees}$?)

If either (a) or (b) is not so, alter to port and pass behind.

(ii) If own-ship is set to pass ahead on a near-parallel course, and turns to parallel-up, will it clear the domain?

I.e. is $L - d >$ domain radius?

(where $d = C * (1 - \cos(A))$)

If so, parallel-up and pass ahead.

Otherwise, alter to port and pass behind.

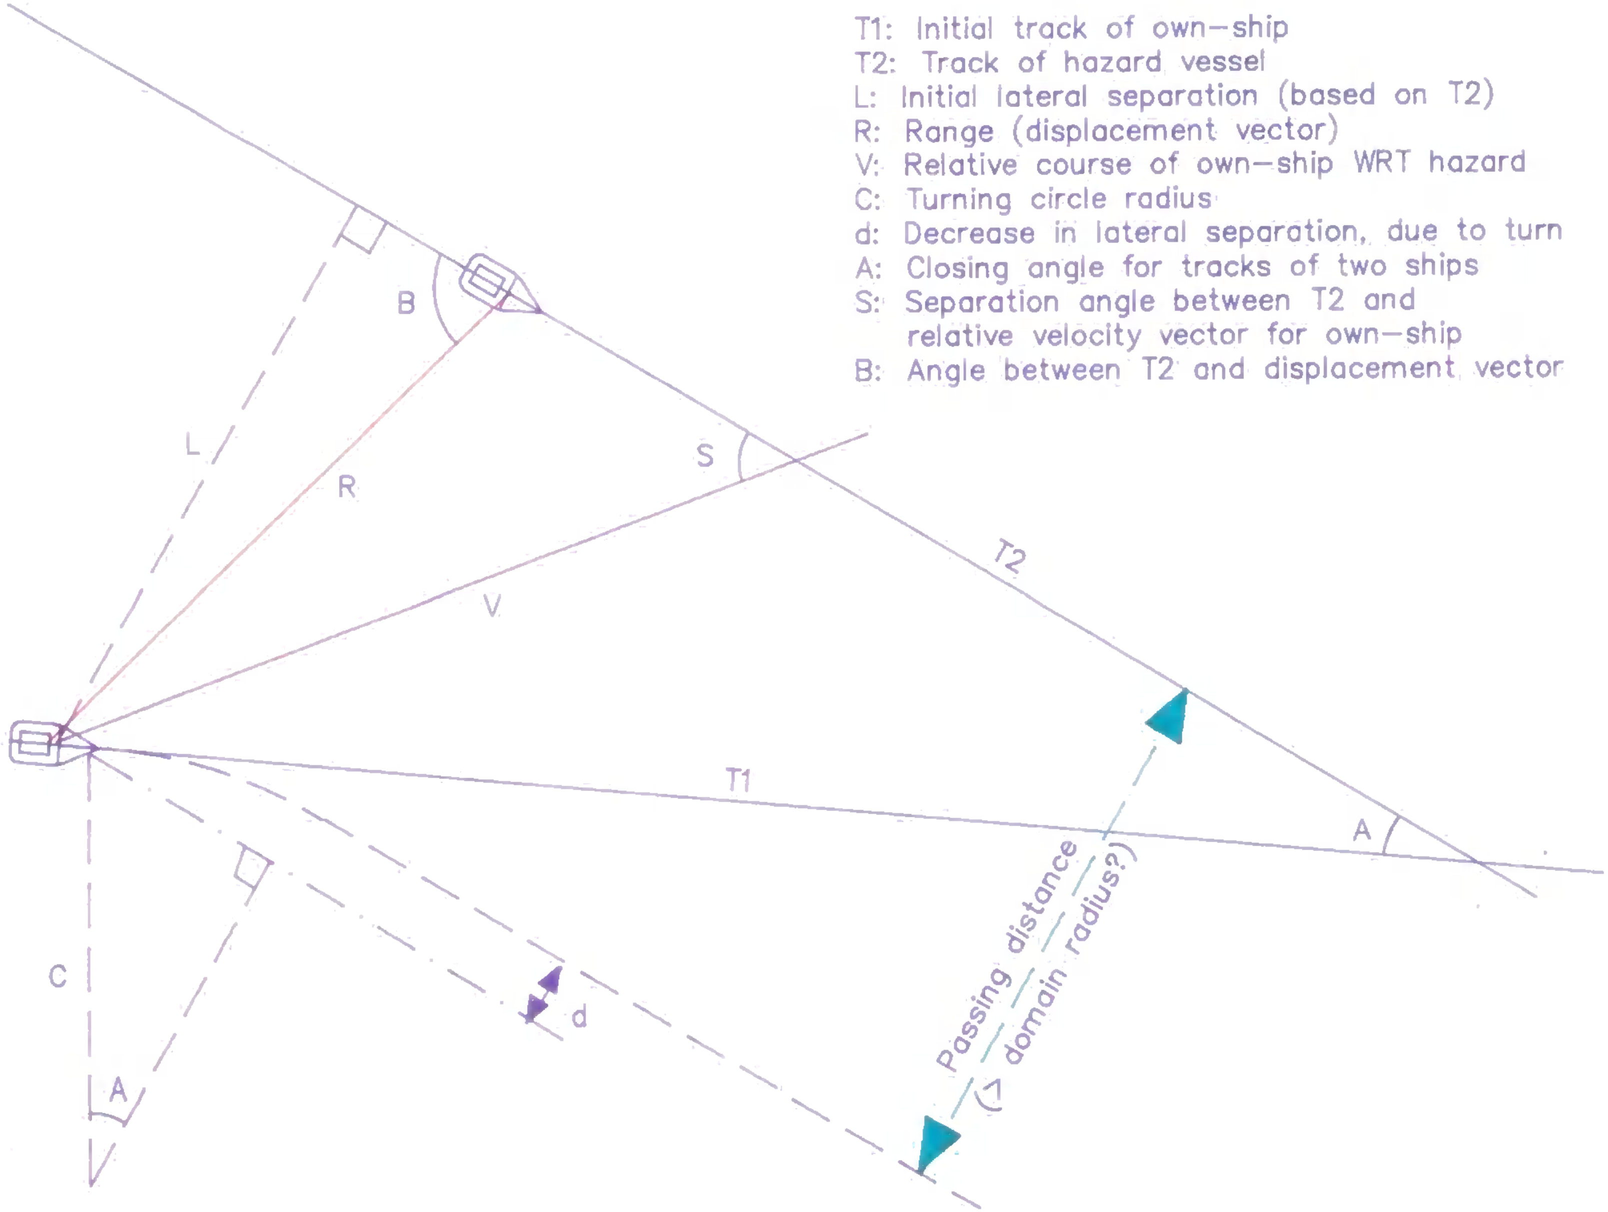

Figure A.2 Calculation of viability of parallel-up for overtaking from starboard quarter

A.3. Analytical Approach to Stages of a Manoeuvre

(for evaluation of POMT - see section 6.3)

A.3.1. Initial Turn

The initial turn to remove the threat of domain infringement must be at least the maximum of:

(i) turn required to clear domain (A + B in fig.A.3 - A being negative if own-ship is initially set to pass behind hazard), and

(ii) 'minimum turn' as required by collision avoidance regulations (subjectively interpreted by mariners, set as a parameter in this expert system).

A further extent of turn may be required by particular rules, to satisfy situation-specific requirements.

As previously stated, the domain may be considered to encompass either own-ship or the hazard vessel, whichever is more expedient mathematically - either option ensures the required separation between the two vessels. In this case (as in most) it is convenient to envisage the domain surrounding the hazard vessel.

It is then necessary to consider whether the domain circle around the hazard vessel intersects own-ship's turning curve as the two ships pass each other; if so, is own-ship on the intersected arc of the turning curve at the time of intersection or not?

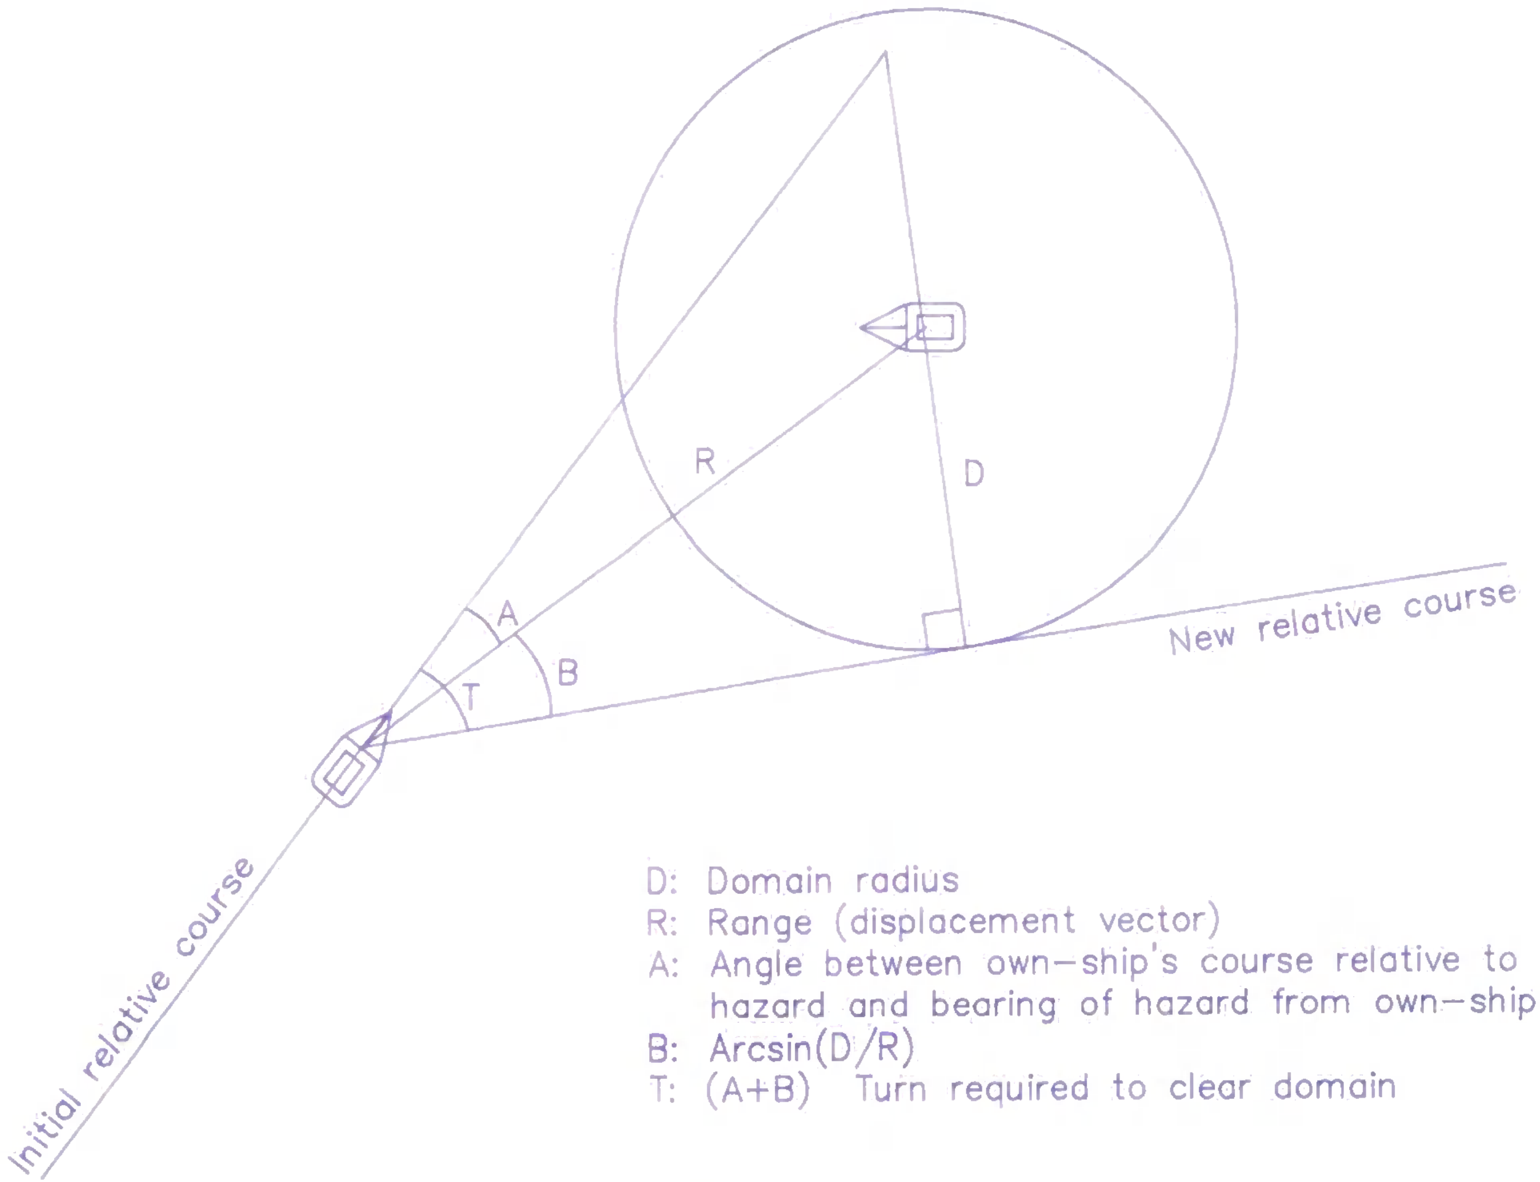

Figure A.3 Calculation of initial turn required to clear domain

This question may be restated as follows:

(i) at the commencement of the turn, does the projected turning curve intersect the current domain circle around the hazard vessel?

With reference to fig. A.4:

T = distance between hazard and centre of turning circle

If $T < C + D$

then turning curve intersects initial domain position.

(ii) (a) if the answer to (i) is YES: does that domain circle become tangential to the turning curve (i.e. clear the turning curve) at a prior time to own-ship reaching that tangential point on the turning curve? (Since the hazard's projected path is a straight line, such a tangential point is known to exist).

I.e. Time taken for M < Time taken for S ? (on fig. A.5)

If so, the domain is cleared, otherwise the domain is infringed.

(b) If the answer to (i) is NO: does the hazard's path cause the domain circle around the hazard to intersect own-ship's turning curve (two tangential positions) or miss it? (Sine rule gives two solutions for F1, F2 - or none).

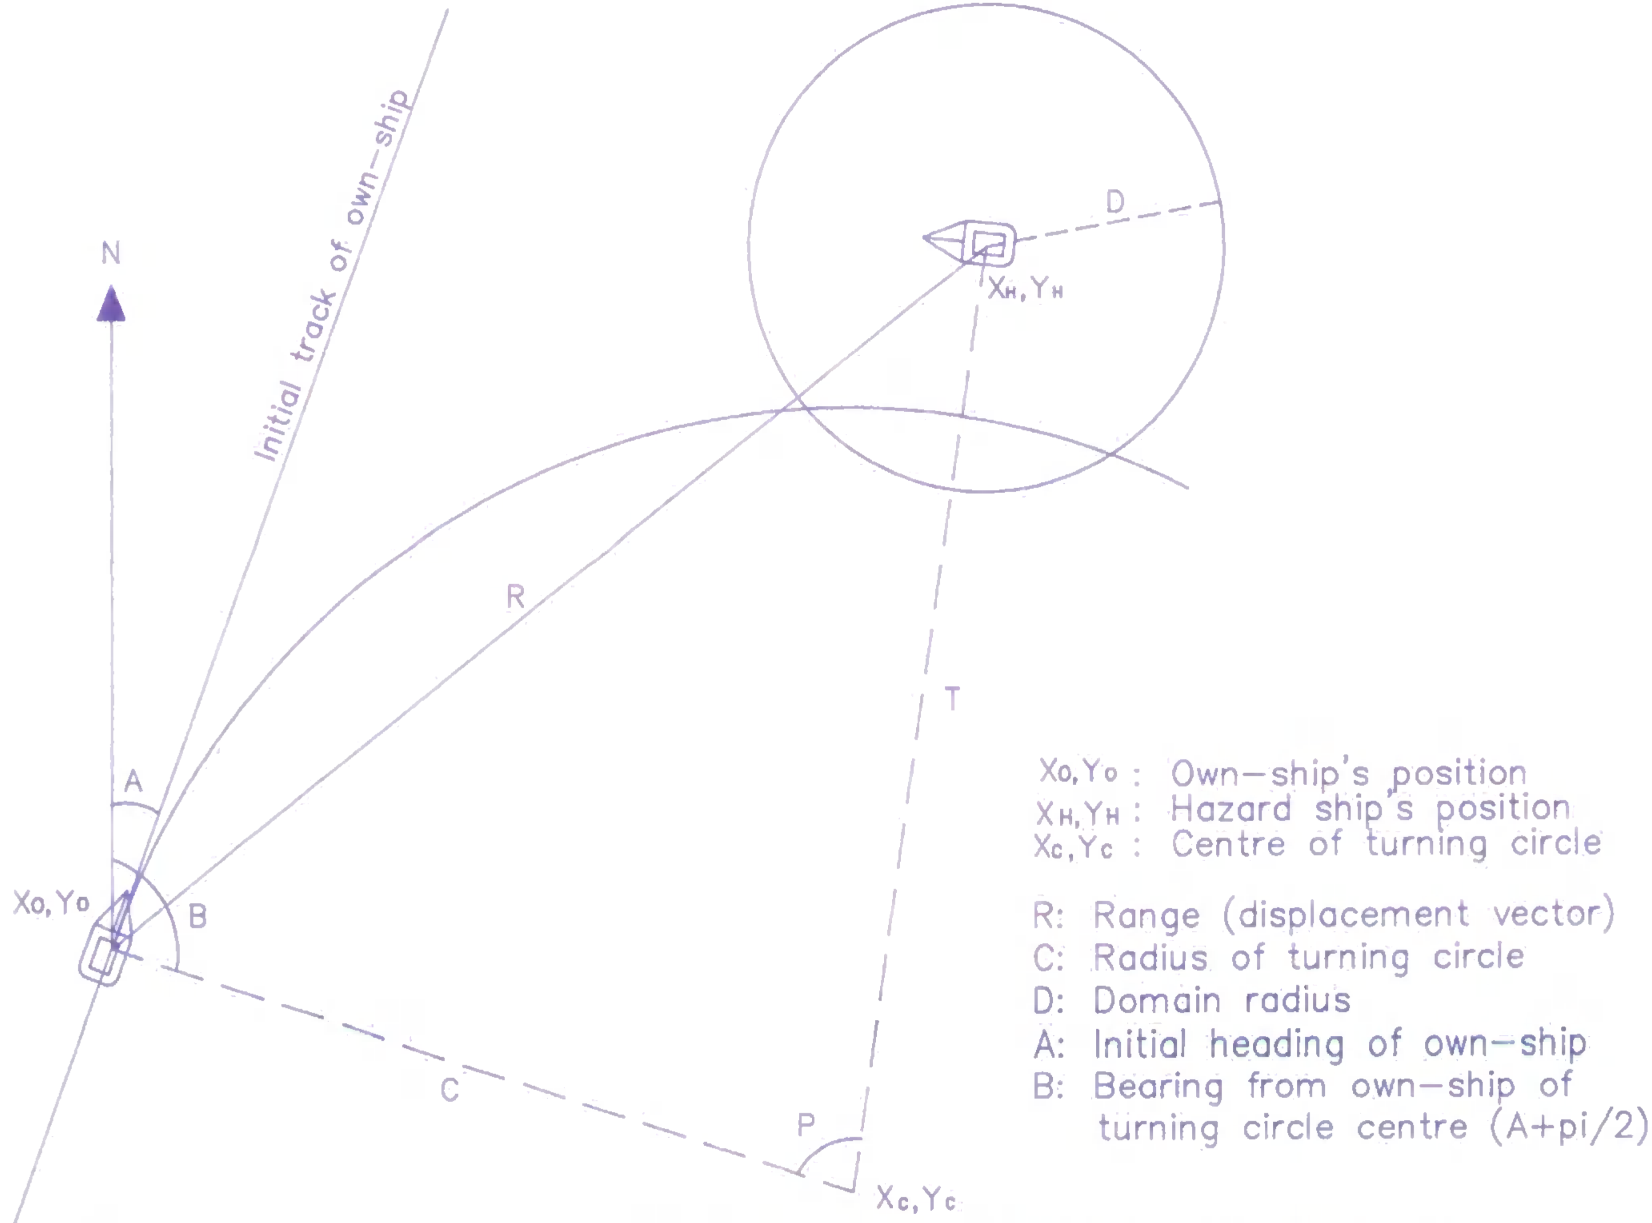

Figure A.4 Check whether initial turn intersects hazard's initial domain position.

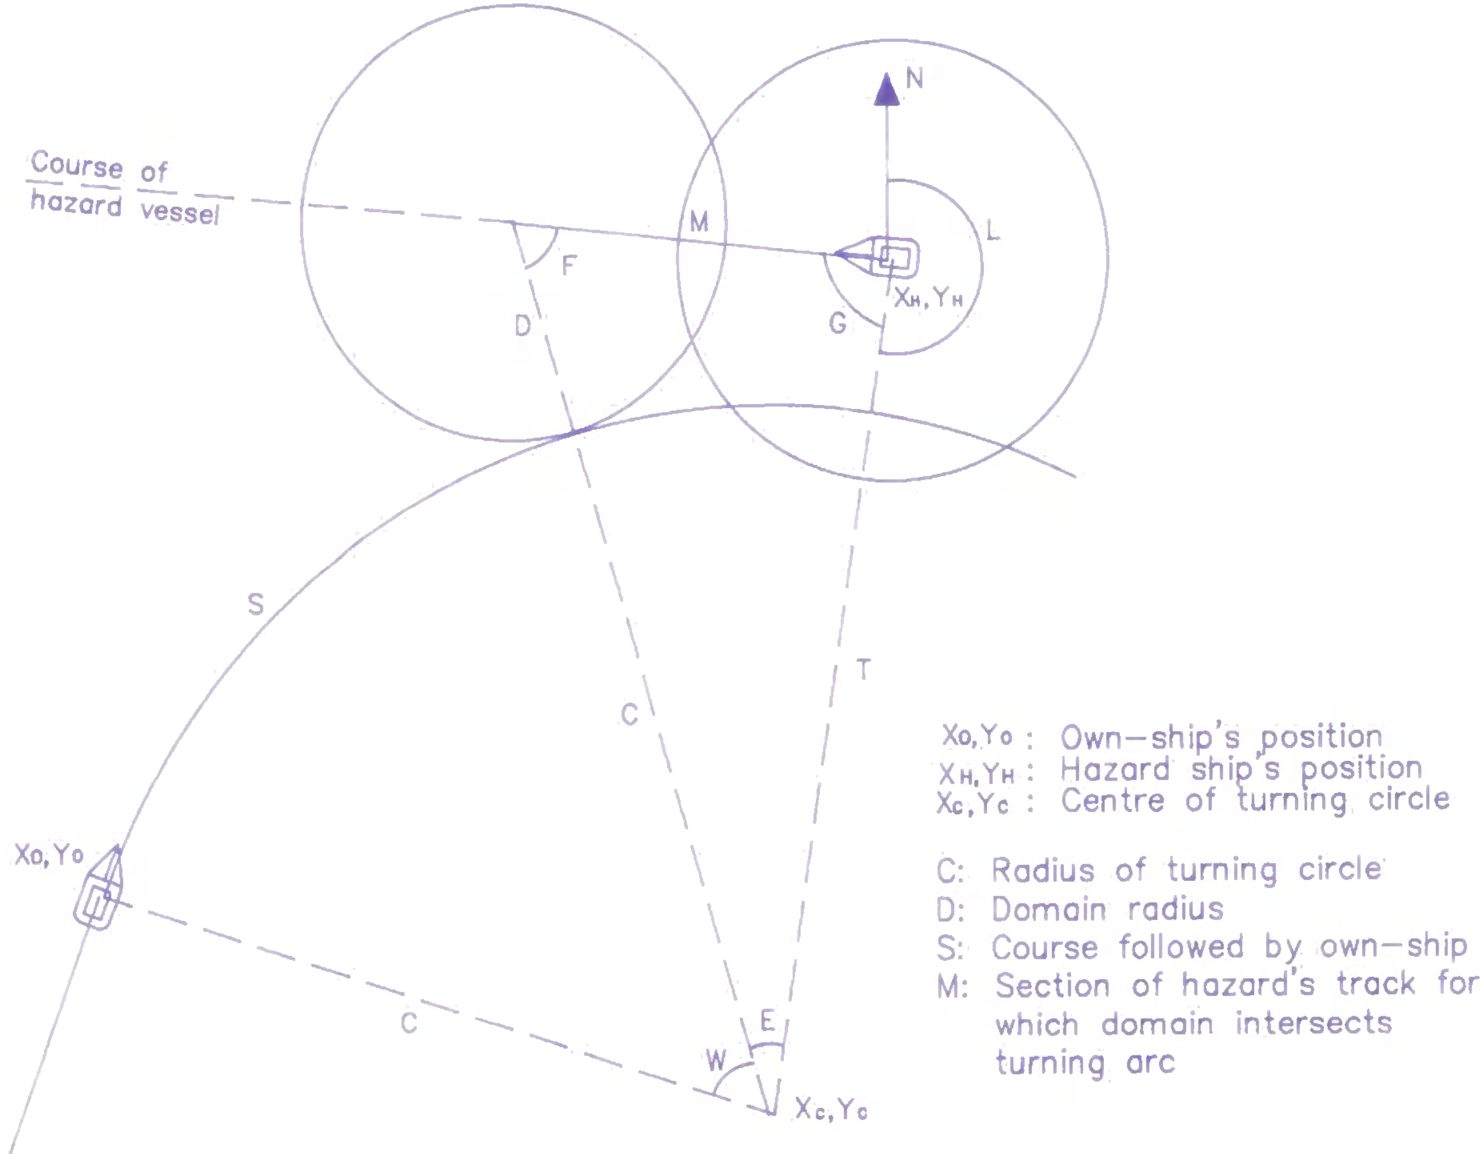

Figure A.5 Check whether hazard's domain clears turning curve before infringement by own-ship.

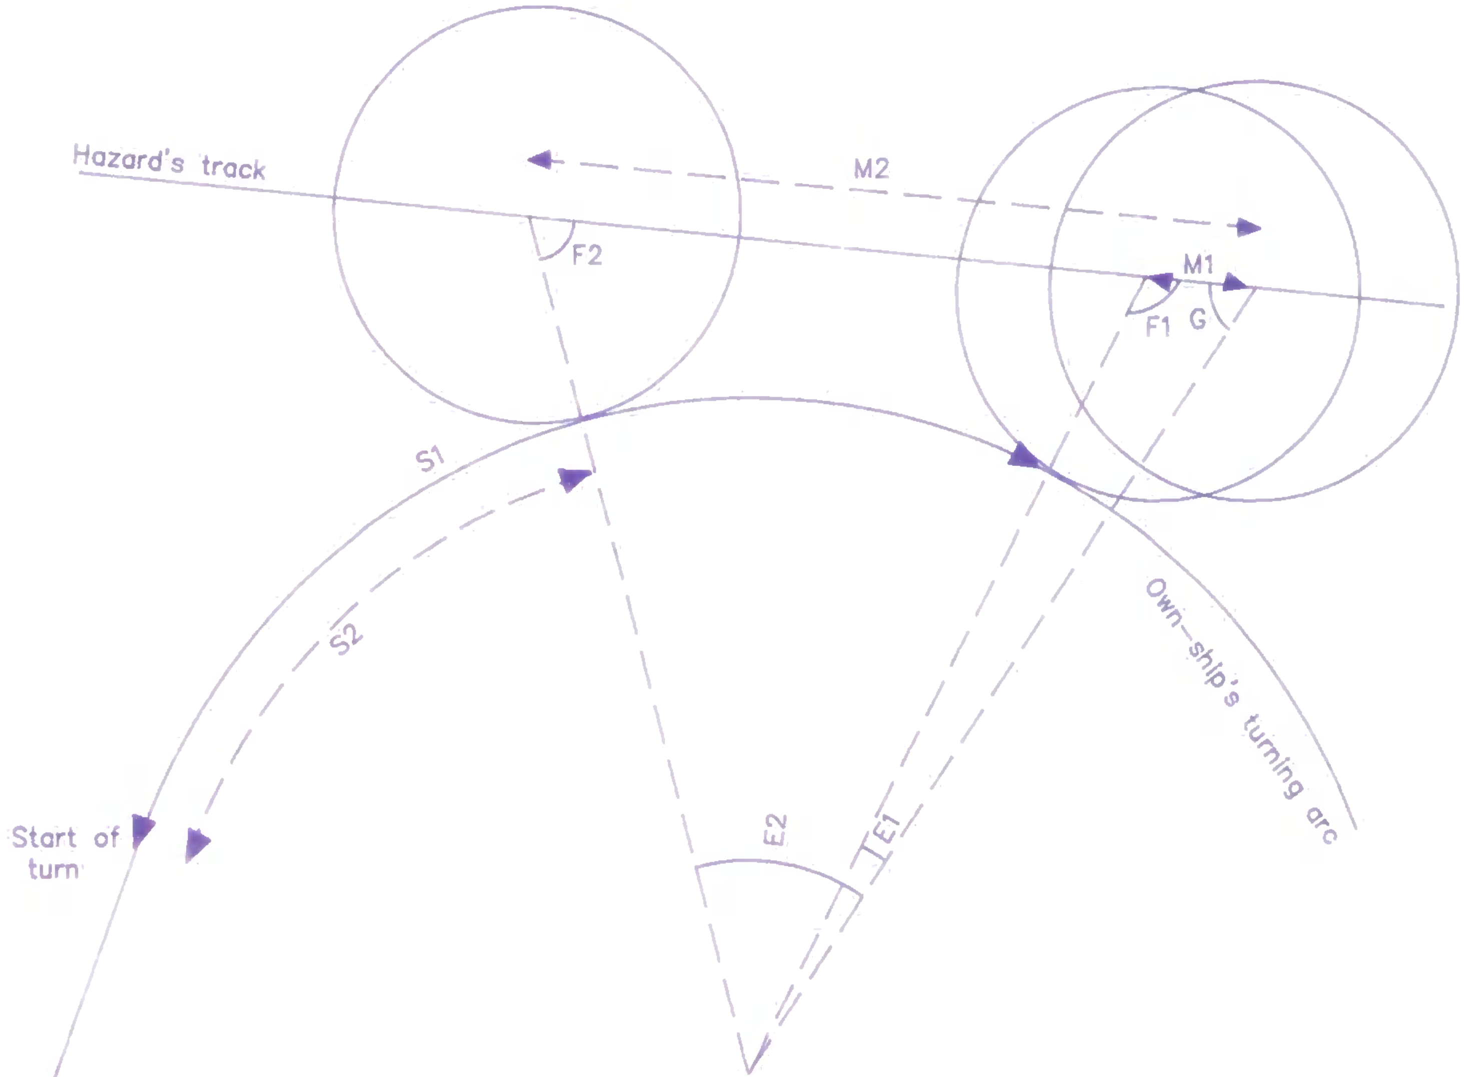

Figure A.6 Check whether own-ship infringes hazard's domain as that domain transit s own-ship's turning curve.

If the latter, no domain infringement is anticipated and own-ship is clear to go.

If the former,is own-ship on the turning curve between those two tangential points at any time whilst the domain is transiting that curve? (See fig. A.6).

I.e. is Time taken for M1 < Time taken for S1

and Time taken for M2 > Time taken for S2 ?

If so, the domain is infringed at some stage;

otherwise, no domain infringement, own-ship is clear to go.

Together, these criteria may be used to ascertain whether the projected turning curve at any time takes own-ship within the domain circle around the hazard vessel: if so, the manoeuvre must be rejected as non-viable; if not, time and vessel displacements may be calculated for the required initial turn as defined earlier. The simulation then hands control back to the rule base, in case the prevailing rule requires a greater extent of turn for any reason. The ensuing straight-line section may then be calculated, as described below.

A.3.2. Straight Line Section

Consideration of a curved track from the point of view of relative motion raises problems, due to the consequent distortion of that track; hence the use of absolute motion (and the lesser problem of two interacting motions) in A.3.1 above. By contrast, straight line motion yields easily to treatment in relative terms, and enables one of the objects under consideration to be held static for calculation purposes. Consequently the straight-line section of the simulated manoeuvre is calculated with the hazard vessel (+ surrounding domain) held static, and motion of own-ship considered relative to this.

Figure A.7 illustrates the requirement for the final (relative) course to clear the (static) domain around the hazard vessel. This in turn defines the distance to be covered on the straight line section of the altered course - without allowance, initially, for the portion covered by the turning curve taking own-ship back onto its original heading.

It is then necessary to consider the turning curve required to effect this stage of the manoeuvre, and whether domain infringement occurs during that second turn. For this purpose it is necessary to again consider absolute motion; note also that this 'minimum' straight line section referred to above will be reduced by that portion taken up by the first half of the turning curve. Only the viability of this turn is of concern here, not its consequences in terms of final vessel positions and speeds or time of completion.

If the turning curve does not intersect the initial portion of the domain around the hazard vessel (i.e. at the start of the turn back onto course) then domain infringement cannot occur, since by definition the hazard is moving away from that turn. Otherwise, the motion of the domain around the hazard must be compared (in time) with the motion of own-ship around its turning curve, as before; if own-ship reaches the

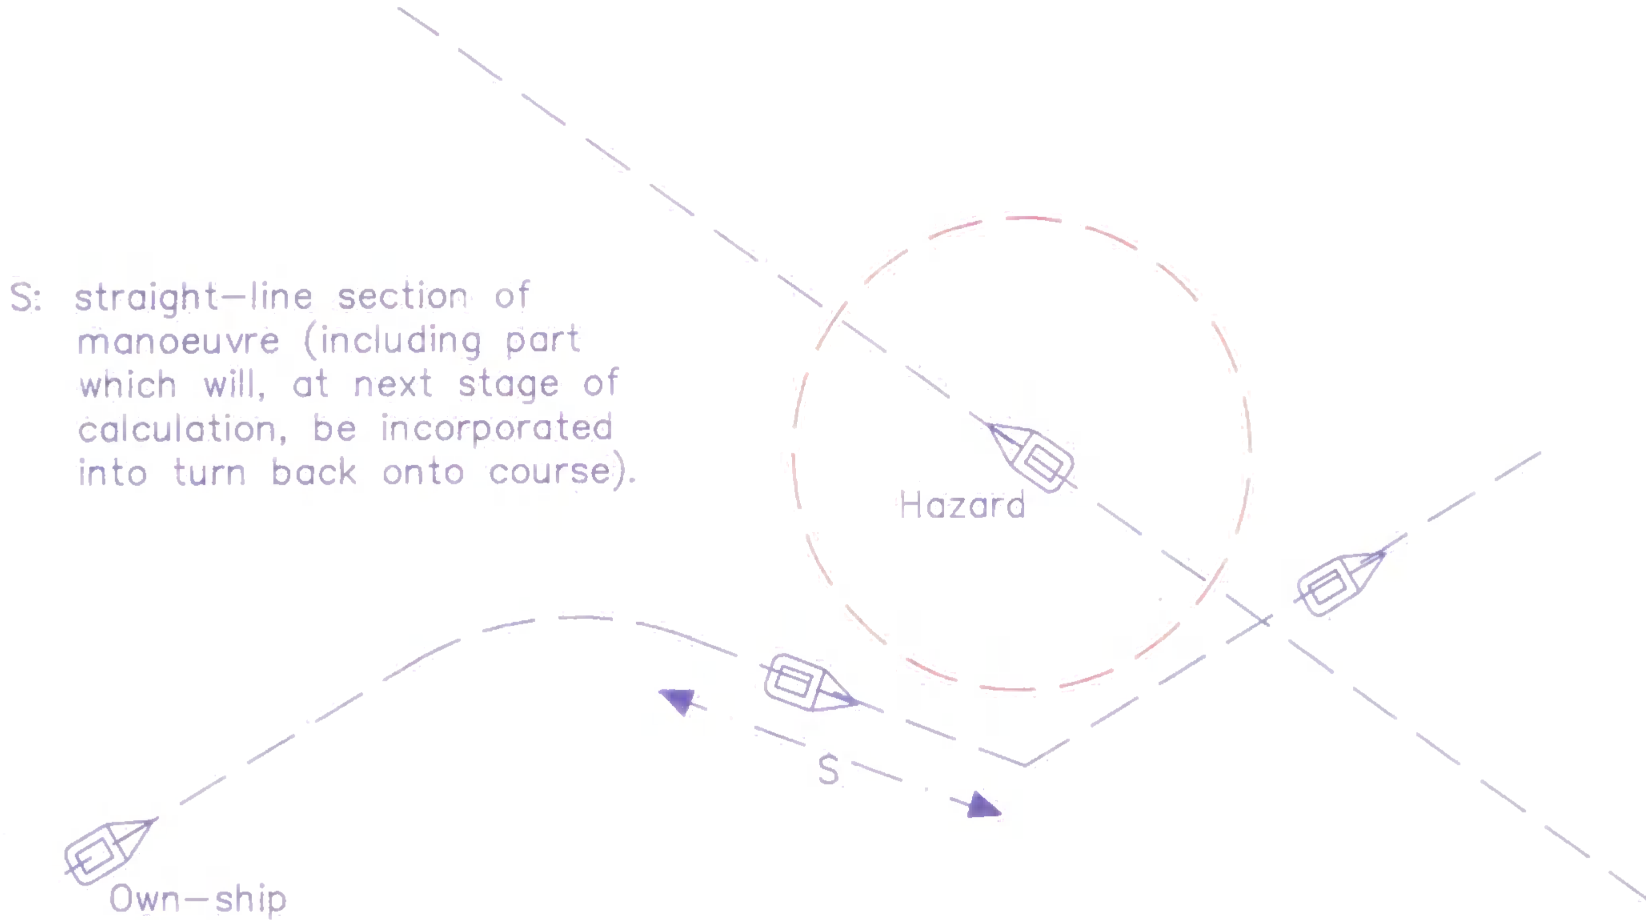

Figure A.7 Identification of appropriate point for turning back onto original course after straight-line section to clear domain around hazard vessel. For such a decision, relative velocity of own-ship may be used (considering hazard as static), since turning curve is not involved at this stage.

point of tangential contact between domain and turning curve <u>before</u> the domain clears that point of contact, then domain infringement will occur - otherwise the domain will not be infringed.

If the straight line section as originally calculated does not lead to clearance of the domain, then an increment is added and the test for viability repeated - until a successful (theoretical) return to course is achieved. This process yields the required straight line section for the manoeuvre (subject to the proviso following), and thus the time and position of both vessels at the end of that section, for the purposes of calculating the cost function. As with the initial turn, certain circumstances may require an extended straight line section; for this reason, after the above calculation, control is handed back to the rule base to allow such further extension if required by the relevant rule, before calculation of the cost function.

The above mathematical treatment of these two sections of a manoeuvre, though intricate, compares favourably with repeated application of the rule structure to 'step' own-ship through a hypothetical manoeuvre. This saving, multiplied by the number of look-ahead simulations evaluated for a single manoeuvre, makes the concept of multiple look-ahead simulation, and thus the POMT, a feasible proposition with the processing power currently available.

The first approach to the above analysis task involved application of standard trigonometry and circle geometry; this appeared to satisfy the requirements for all simulated sample encounters. However, as indicated in section 9.1, random simulation of a large number of encounters showed up limitations to this approach, the generality of the mathematics being limited by implicit assumptions in the geometry involved. As a result of that observation, this form of analysis was replaced by a variant of the iterative binary halving process. This process involves calculation of the range between the two vessels: (a) at the start of the manoeuvre; (b) at the end of the manoeuvre; and (c) at three intermediate positions equally spaced in time. The time giving the least range is identified, and the time-steps to either side of this are taken as the limits for the next stage of the iteration. This process is repeated until either the range falls within the domain radius (i.e. domain infringement) or the time-interval becomes too small for domain infringement to be possible at current speeds. This approach makes no specific assumptions regarding the geometry involved, and has proved totally effective over one thousand randomly-simulated encounters as documented in section 9.1.

APPENDIX B

RESEARCH VESSELS USED FOR SEA TRIALS

Since September 1990, the Marine Dynamics Research Group at Polytechnic South West has been engaged in collaborative research with Britannia Royal Naval College, Dartmouth. The college maintains a number of naval picket boats for training in small vessel handling, learning elementary navigation and practice in application of the rules for inshore navigation. One of these vessels, picket boat 9, has been equipped with instrumentation and sensors to facilitate the undertaking of this joint research project. BRNC has, from time to time, generously made this vessel available (with support crew) for associated work, such as the research described in this thesis. As stated in section 9.3, another such vessel (picket boat 5) has similarly been made available in circumstances where a second vessel has been required as a target.

Picket boat 9 is a twin-screw, twin rudder 13 metre vessel, with a standard hull form (making it suitable for hydrodynamic modelling in future developments of this work). It is fitted with a Kelvin-Hughes Concept HR3000A ARPA, the display for which is mounted to the starboard side of the bridge (see fig. 9.6(b)). Workspace and mains power are available adjacent to this unit, providing an ideal site for the Archimedes-based expert system (see fig. 9.6(a)). The scanner for the ARPA is positioned 1 metre above the bridge roof affixed to a steel tripod (see fig. 9.7(a)), placing it some 3.5 metres above the water-line.

Picket boat 9 is also equipped with a flux gate compass, optimally positioned to eliminate roll and pitch components of motion and record only the yaw component. A Global Positioning System (GPS) receiver is fitted to a bracket on the ceiling of the bridge, above the chart table. It is thus ideally placed for stand-alone readout, and also for possible future interfacing to a computer system.

Sensors are mounted adjacent to both drive shafts, on a bulkhead which offers a good firm basis for reliable readings. At the commencement of the joint research project, the vessel was out of the water being refitted with new engines; this also offered the ideal opportunity to fit a paddlewheel log to an appropriate position on the hull. It is anticipated that the readouts from these sensors could all be used as inputs to shipboard computer systems such as the collision avoidance system described in this thesis.

APPENDIX C

LIST OF PUBLICATIONS

Blackwell GK Burns RS and Stockel C T. 1987	An expert systems approach to collision avoidance. Proceedings of the 1987 Summer Computer Simulation Conference, Society for Computer Simulation, Montreal (Quebec), July 1987, 43-48
Burns R S Blackwell G K and Calvert S. 1988	An automatic guidance, navigation and collision avoidance system for ships at sea. Colloquium on 'Control in the marine industry'. Institution of Electrical Engineers, London, January 1988, 3/1-3/3
Blackwell G K Colley B A and Stockel C T. 1988	A real-time intelligent system for marine collision avoidance. Proceedings of Third International Conference on Applications of Artificial Intelligence in Engineering. Computational Mechanics Institute, Stanford University (California), August 1988, 119-138
Burns R S Blackwell G K and Stockel C T. 1988	Ship manoeuvring simulation and intelligent marine collision avoidance. Proceedings of USER1 working conference for users of simulation hardware, software and intelliware. Society for Computer Simulation, Ostend (Belgium) September 1988, 184-189

Blackwell G K, Colley B A and Stockel C T. 1988 — An IKBS for marine collision avoidance. Proceedings of International Conference on Maritime Communications and Control, Institute of Marine Engineers, London, October 1988, Paper 19, 1-8

Burns RS, Blackwell G K and Miller K M. 1989 — Simulation of a marine guidance and expert collision avoidance system. Proceedings of Spring Simulation Conference, Society for Computer Simulation, Tampa (Florida), March 1989, 62-67

Blackwell G K, Stockel C T and Colley B A. 1989 — Look-ahead simulation as a tool for real-time expert systems. Proceedings of 1989 European Simulation Multiconference, Society for Computer Simulation, Rome (Italy), June 1989, 157-162

Blackwell G K and Stockel C T. 1989 — Simulation of ship encounters using multiple linked intelligent systems. Proceedings of 1989 Summer Computer Simulation Conference, Society for Computer Simulation, Austin (Texas), July 1989, 766-771

Blackwell G K and Burns R S. 1990 — Onboard systems for automatic ship guidance and intelligent collision avoidance. Proceedings of International Conference on Modelling and Control of Marine Craft, Exeter University School of Engineering, Exeter, April 1990, 121-135

Blackwell G K and Stockel C T. 1990	Optimal strategies for a marine real-time expert system. Proceedings of 1990 European Simulation Multiconference, Society for Computer Simulation, Nuremberg (Bavaria), June 1990, 294-299
Blackwell G K and Stockel C T. 1990	Strategic planning for intelligent real-time management. Proceedings of 1990 Summer Computer Simulation Conference, Society for Computer Simulation, Calgary (Alberta), July 1990, 694-699
Blackwell G K Chudley J and Dove M J. 1990	The use of a mathematical model in a collision avoidance system. Proceedings of the Second International Conference on Maritime Communications and Control, Institute of Marine Engineers, London, November 1990, Paper 7, 7/1-7/6
Blackwell G K Rangachari J and Stockel C T. 1991	Ships that pass in the night. Intelligent management of a 'fuzzy' event stream. Proceedings of the 1991 European Simulation Multiconference, Society for Computer Simulation, Copenhagen (Denmark), June 1991, 542-547
Blackwell G K Rangachari J and Stockel C T. 1991	An intelligent interactive environment for a maritime real-time expert system. Proceedings of the 1991 Summer Computer Simulation Conference, Society for Computer Simulation, Baltimore (Maryland), July 1991, 358-363

PROCEEDINGS OF THE 1987 SUMMER COMPUTER SIMULATION CONFERENCE

JULY 27-30, 1987
THE QUEEN ELIZABETH HOTEL
MONTREAL, QUEBEC, CANADA

NINETEENTH ANNUAL SUMMER
COMPUTER SIMULATION CONFERENCE

Edited by
Jordan Q.B. Chou

CONFERENCE SPONSOR

SCS

THE SOCIETY FOR COMPUTER SIMULATION

Printed in the United States of America

An expert systems approach to collision avoidance

Grahame K. Blackwell, Roland S. Burns and Colin T. Stockel
The Ship Control Group
Plymouth Polytechnic
Plymouth, Devon, England

ABSTRACT

Computer simulation of marine traffic flow and collision avoidance, using traditional (procedural) programming techniques, have much to offer, both in training of mariners and in considering the effects on high-density traffic lanes of changes in such parameters as traffic volume on a particular route. However, such simulations are limited in terms of their ability to model real-world situations, due to over-simplification of the decision-making process in two-ship or multi-ship encounter situations.

A prototype expert system has been developed, incorporating the International Rules for Preventing Collisions at Sea and the control/response characteristics of own-ship - these latter parameters being varied to model different vessels. The knowledge base also contains information on own-ship's position and velocity, together with the corresponding information for the perceived hazards. From this, the system is able to evaluate the threat posed by a hazard, with reference to the concepts of domain and 'range to domain over range rate' (RDRR). The expert system then selects the best course of action and advises the operator accordingly. The situation is re-evaluated at regular intervals, and new advice given as necessary.

The expert system has been written in the 'C' programming language, on a microcomputer system capable of operating in real-time on board ship. The microcomputer used has the capacity for information input directly from the various sensor devices available.

The prime objective of this work is to provide an expert advisor for mariners, capable of receiving sensory input from all available channels - including information provided by the operator via a keyboard - and advising on the preferred course of action. Development of this system has necessitated a facility whereby the expert system may be interrogated to ascertain the reasons for its decisions; this is considered to be a useful training aid.

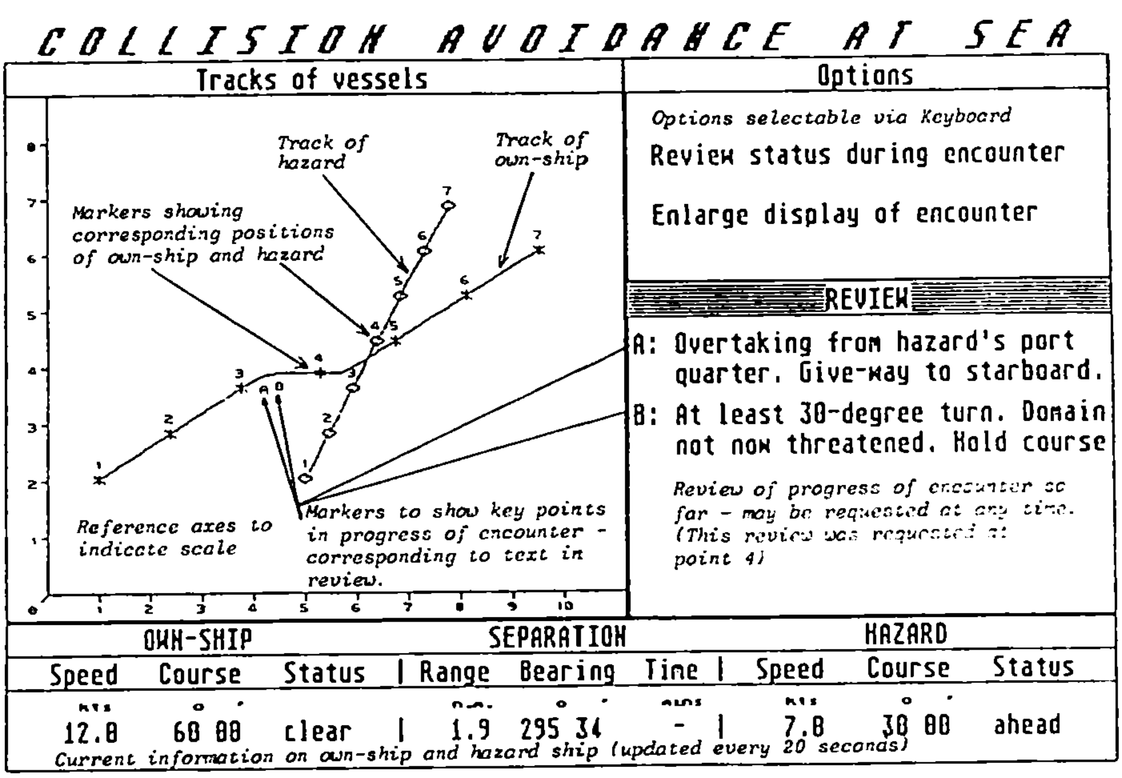

Figure 1: Annotated example of the screen display generated by the prototype system.

1. INTRODUCTION

Conventional computer simulations of marine traffic are acknowledged to be of considerable benefit in broadening the experience of trainee mariners. Such simulations are also recognised to be well-suited to modelling traffic flow along routing systems, and thus demonstrating the likely effects on such routes of changes in traffic volume, introduction of a fixed or free-moving hazard into the traffic lanes, or other variation in conditions. However, by their nature such simulations generally rely on a simplified, mechanistic approach to the decision-making process in two-ship or multi-ship encounters.

What is required is an 'intelligent' response system, capable of evaluating such an encounter between two or more vessels from the standpoint of an experienced mariner, and taking appropriate action. Such a system, on being provided with the information normally available on the bridge of a well-equipped ship, should direct the course of action of that vessel as would a competent mariner. This involves:

(a) Recognising an encounter (potential hazard) situation in good time;
(b) Identifying the type of encounter, and the status of one's own vessel (own-ship) in that encounter, according to the International Rules for Preventing Collisions at Sea - such status would normally be either 'stand-on' (having right of way) or 'give-way', according to one's position relative to the other vessel, and other relevant factors;
(c) Choosing a course of action which combines a sensible safety margin with the least practicable inconvenience (where own-ship is judged to be stand-on, this will - initially at least - mean no alteration of course);
(d) Maintaining a watching brief on the other vessel(s) in the encounter, and being ready to take avoiding action if necessary, should the situation change

Clearly implicit in this 'seaway code' is the understanding that any other vessel encountered will be assumed to be 'well-behaved' until proved otherwise, but that an identical perception of, and reaction to, the encounter by the master of another vessel cannot be taken for granted. It is also imperative always to adopt a give-way stance in relation to any hazard other than another vessel under control, for example a derelict or a rock!

An intelligent system of the sort described is currently under development, with a view to providing a microcomputer -based on - board facility which would act in the role of Expert Adviser to the man on the bridge. This Expert System would receive sensory input relating to own-ship and any vessels or other objects within range, and provide advice on the preferred course of action at each stage in an encounter; on request, the system would explain its advice in terms of the decisions involved. Development is currently at the 'early prototype' stage.

2. THE PROTOTYPE

At its current stage of development, the system deals with open-sea encounters (not in restricted channels or routing systems, no fixed obstacles) between two manned vessels moving subject to fairly idealised parameters. The rule base (decision-making logic) for the system is fixed, consisting of rules for:

(a) Identifying the presence of a potential hazard, and assessing degree of threat in terms of expected time to collision or near-miss - this time factor is central to the decision on when to take avoiding action, if necessary;
(b) Identifying the type of encounter, including several variants in some cases (the primary types of encounter being head-on, crossing, and overtaking), and fixing the status of own-ship and perceived status of hazard-ship in the encounter;
(c) Negotiating the stages of the encounter, with due regard for appropriate safety margins.

The computer program for this prototype system effectively performs three functions: first, it provides a reference for decision-making for own-ship; second, it has to provide simulated data for own-ship, data which will ultimately be acquired through sensors from instrumentation on-board ship; third, it has to simulate a hazard ship which must also give the appearance of being under some form of intelligent control, independent from the control of own-ship.

The program, then, must be seen as two separate entities:

(1) a (currently rather simple) expert system, which interrogates the knowledge base available for a specific ship - this would include radar data on position, speed and course of the hazard vessel (subject to limited accuracy) - and gives direction on the course of action to be taken, by setting status flags in that knowledge base.
(2) a simulator which, given initial data for a vessel, will calculate the course for that vessel, subject to course changes introduced by the expert system.

The interface between these two parts is the knowledge base, which holds velocities, displacements, relative bearings etc generated through the actions of the simulator for reference by the expert system, and likewise records the decisions of the expert system for action by the simulator. For convenience, the expert system logic is used to control the hazard ship as well as own-ship; it is intended, at the next stage, to provide a manual override for the hazard ship, so that it may be independently controlled to enable testing of further features of the expert system. At a later stage, it is envisaged that two computer systems should each model the action of one ship, and that these two models should interact via a communications link between the systems.

The microcomputer system in use is an Atari 520ST with 1 megabyte of main memory, operating under GEM* and programmed in 'C' for fast real-time processing involving complex mathematical calculations. GEM (Graphics Environment Manager) is a WIMPS (Window - Icon - Mouse Processing System) environment, offering (a) full window management of multiple screen windows for text and/or graphics; (b) use of icons, or representative on-screen symbols, as prompts to user action; (c) user interaction via a 'mouse', moving a screen cursor for option selection and other interactive graphics applications; (d) multiple event handling under full application control; (e) limited multitasking, comprising user application and background system functions. Some of these features are used at a fairly low level in the current prototype, but they are expected to become more significant as the system advances in sophistication.

*GEM is a registered trade mark of Digital Research Incorporated.

The simulator operates on a 20-second discrete time interval, the simulation proceeding at several steps per second, real time. At each time-step the data for each ship is updated, and each vessel is independently submitted to the expert system for decisions as to subsequent action.

Figure 1 shows an annotated example of a typical screen display, which consists of four windows under a main title:

- Window 1 provides a graphical display of the two-ship situation, updated as the situation proceeds:
- Window 2 displays the available options, namely (a) to show an enlargement of the encounter part of the situation so far, in the graphics window, and (b) to provide a review in window 3 of the key decisions in the encounter so far, with cross-references to the corresponding key points on the graphics display of own-ship's course;
- Window 4, along the bottom of the screen, shows current data on own-ship, and on hazard vessel to the extent that sensors such as radar could be expected to provide such data. In particular, the status of own-ship is determined in detail by the expert system. Assessment by the master of hazard ship as to his status can only be judged from own-ship by whether or not hazard ship has deviated to a noticeable extent from its set course, presumably for collision avoidance purposes. Such an observation gives rise to the two possible status conditions for hazard-ship: 'turning' (observable deviation from original heading) or 'ahead' (on original heading).

3. THE INFERENCE ENGINE

This term is used to describe the logic system whereby the rule structure of an expert system is brought to bear on the knowledge base in order to draw conclusions, possibly under probability constraints and subject to confidence limits. Such probabilistic considerations are envisaged in later developments of this system, but the 'inference engine' of this prototype rests primarily on previous research in the field of marine collision avoidance. In particular, two well-researched concepts dictate the nature and extent of action taken: the domain, or 'danger zone' around a vessel, into which no other vessel or obstacle must be allowed to encroach; and the arena, or 'circle of influence', a larger zone around a vessel, inside which any other vessel or obstacle must be considered a potential hazard.

The concept of the domain was advanced by Goodwin[1] as 'the effective area around a ship which a navigator would like to keep clear with respect to other ships and stationary objects'. The original concept of three distinct circle sectors of differing radii for sidelights and sternlight was later modified[2] to give a continuous circular domain boundary, with the ship off-centre in the circle so as to preserve areas subtended by sidelights and stern light. This rationalisation was based on the logical necessity for continuity of the domain boundary. (See fig. 2)

The arena concept was introduced by Davis[3] in response to the observed actions of mariners manoeuvring to prevent infringement of their domain by ships in their vicinity. This concept was refined by Colley[4] to a time factor, referred to as the RDRR (Range-to-Domain/Range-Rate) criterion - basically the time-to-go (before calculated domain infringement) at which the mariner will assess the status of his vessel and take action accordingly. For example, a mariner may choose to take action 12 minutes before potential domain infringement - a 12-minute RDRR.

Since mariners on two vessels may choose different RDRR values, and different domain boundaries, and since the Collision Avoidance Rules are based primarily on relative positions at the time of decision-taking, those two mariners may in all conscience interpret a situation at different times, and therefore possibly in a different way. The Collision Avoidance Rules specify actions in terms of the type of encounter, so different interpretations may lead to conflicting actions. Such uncertainty, as well as the possibility of 'careless drivers', must be catered for in a fully-effective expert system.

The prototype system has as its first objective the protection of the integrity of the domain. This is simplified to a circle centred at own-ship; it is anticipated that future developments will favour the 3-sector domain, based on three concentric circles, with 'fuzzy' (probability-weighted) boundaries between the sectors. The RDRR criterion

fig. 2

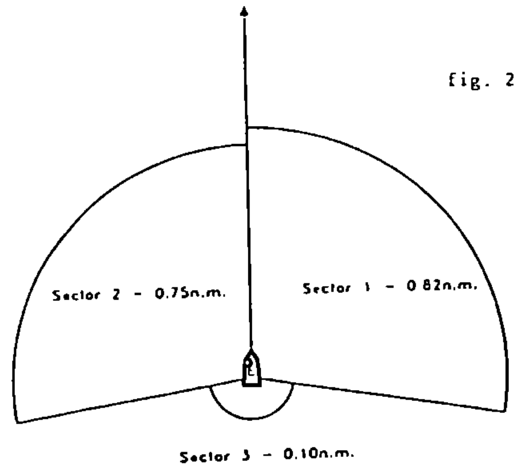

(a) Goodwin domain (example)

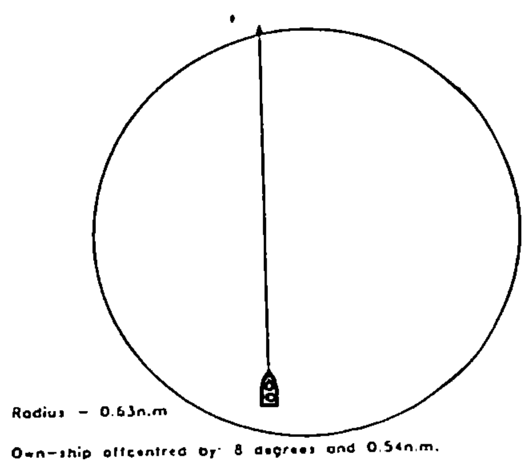

(b) Davis domain (example)

is used to determine decision-time; future developments will incorporate the RDRR+ criterion put forward by Konyn[5], which allows for 'thinking time' between situation appraisal and action. Both the timing and extent of the action are important, since a ship must not only take action but must also be seen to be taking action in good time by the master of the other vessel.

As a further simplification of the prototype, the hazard vessel is allocated the same domain radius and RDRR value as are input for own-ship. This ensures that both vessels appraise the situation at the same time, obviating for the time being a complication described earlier.

4. THE RULES

As previously stated, encounters fall into three main types:

(1) Head-on
Two vessels are approaching each other on (almost) reciprocal bearings. A vessel in such an encounter must consider itself give-way and alter course to starboard (right); it should not begin to alter back onto course until the other vessel has passed 90° on the port beam (lefthand-side);

(2) Crossing
Two vessels are approaching each other at an angle such that each can see one side of the other; more specifically, one vessel is in the sector subtended by the port or starboard light - 112½° measured from the bow (front) on each side. The ship with the other vessel to starboard should give-way, altering course to starboard and going behind. It should not alter-back until the stand-on vessel has passed 15° to port of own-ship's original heading;

(3) Overtaking
A vessel is deemed to be overtaking if it is approaching from within the stern-light sector, ie the rear 135° arc. The overtaking vessel should give-way as follows:
(a) If on near-parallel course, go behind to starboard; alter-back when safe to do so;
(b) If on port quarter of stand-on vessel, action as for (a);
(c) If on starboard quarter, and set to overtake ahead on current course, and if able to parallel courses without infringing domain, parallel-up; when well past and clear to pull across, return to original heading;
(d) Otherwise, go behind to port; alter-back when safe to do so.

The above definitions embody the International Regulations for Preventing Collisions at Sea[6], plus clarification of specific practical details as observed by Colley[7] and implemented in a previous simulation. These definitions form the essence of the rule structure of the prototype system under discussion here.

Using these rules, every collision avoidance manoeuvre is taken through four stages, as follows:

(1) When the RDRR criterion is infringed, the give-way vessel initiates a turning action (at present, this is taken as a constant-velocity arc of a circle);
(2) Once the required turn is effected (Minimum 30° - except in (3) (c) above - so as to be clearly observable), the vessel ceases its turn;
(3) When original heading may be resumed without further risk of domain infringement (and any other conditions are satisfied), alter-back is initiated;
(4) When back on original heading, manoeuvre is terminated, vessel proceeds on that heading.

It should be noted that the function of the 'inference engine' at each time-step is to take the basic parameters (speed, heading, position) provided by the simulator. From these it generates a standard set of further parameters - relative bearing, relative velocity, angle between relative bearing and relative heading; etc - and applies the rules to them to ascertain which situation pertains to the next step.

5. EXAMPLES

Fig. 1 shows own-ship overtaking hazard from hazard's port quarter.

Fig. 3(a) shows a crossing encounter with own-ship giving-way; a review has been requested at point 7. Fig. 3(b) shows enlarged detail of the same encounter, with a full review.

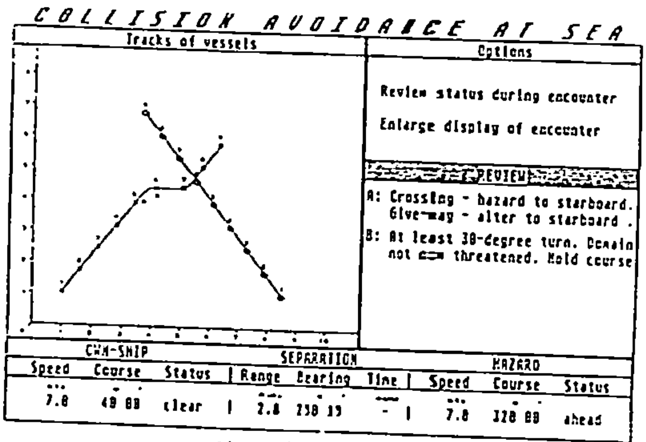

fig. 3(a)

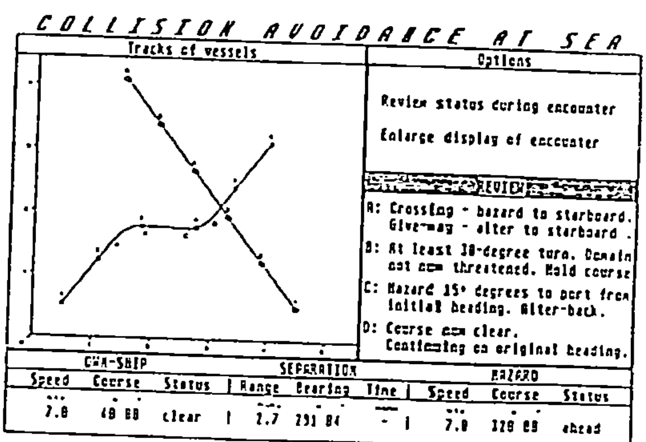

fig. 3(b)

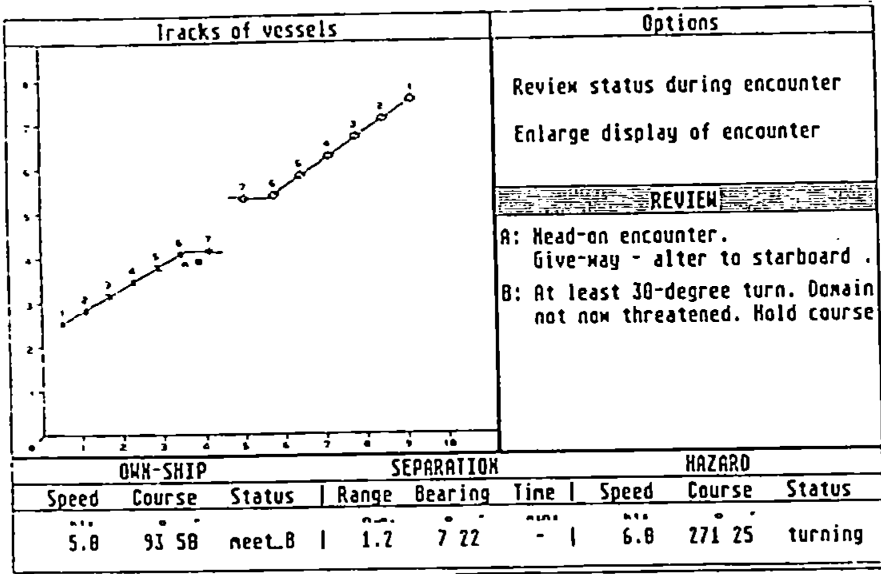

N.B. The screen header has been removed from figs. 4(a), 4(b) and 5.

Fig. 4(a) shows a head-on encounter partway through; the status information gives detail of the status of own-ship (2nd stage of a head-on encounter), but indicates only that hazard ship has diverted from original heading.

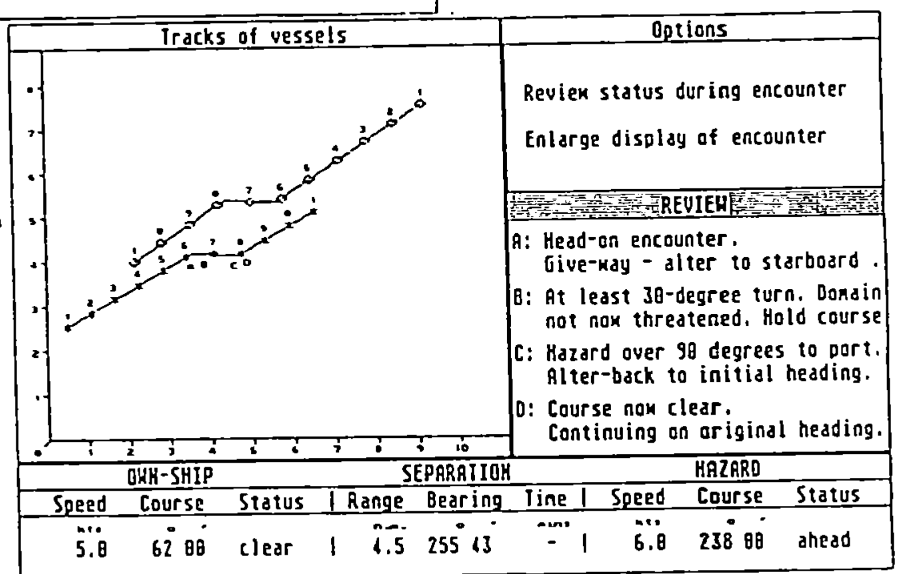

Fig. 4(b) shows the same encounter completed, with full review and both vessels back on original heading.

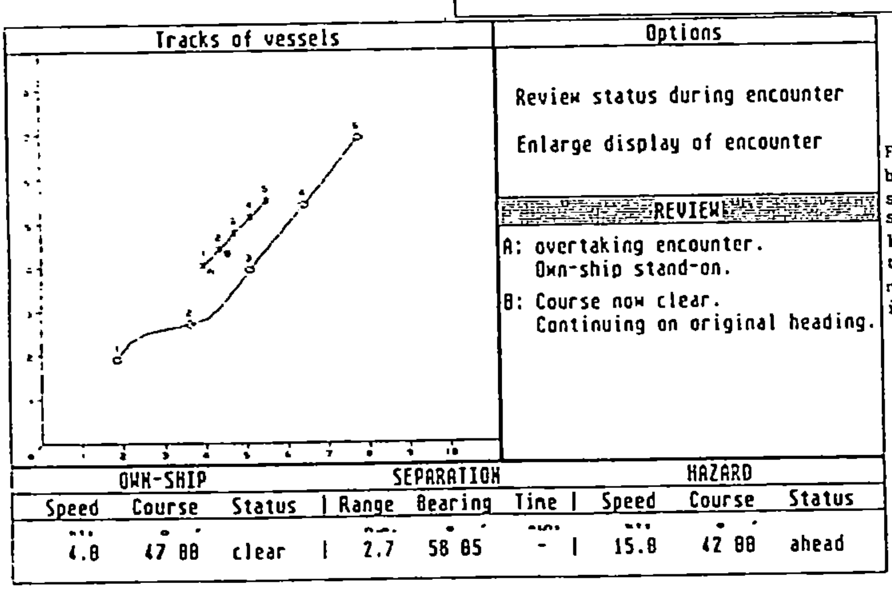

Fig. 5 shows own-ship being overtaken by hazard vessel, very soon after start of simulation run. Table 1 shows the first part of the data-logging file for this run including the initial data input. (N.B. negative time-to-go is used to indicate 'domain not threatened').

data for own-ship

Course:- (i) Degrees:	47.00
(ii) Minutes:	0.00
Speed in knots:	4.00
Turning circle (n.mls):	0.60
Position:- x:	4.00
Position:- y:	4.00
Domain radius in mls:	0.80
Arena (RD/RR) in minutes:	10.00

data for hazard ship

Course:- (i) Degrees:	42.00
(ii) Minutes:	0.00
Speed in knots:	15.00
Turning circle (n.mls):	0.40
Range in nautical miles:	3.00
Bearing:- (i) Degrees:	225.00
(ii) Minutes:	0.00

OWN-SHIP X	Y	COURSE	SPEED	TIME TO GO	HAZARD X	Y	COURSE	SPEED	RANGE
4.00	4.00	47.00	4.00 :	12.14 :	1.88	1.88	42.00	15.00 :	3.00
4.02	4.02	47.00	4.00 :	11.81 :	1.93	1.94	42.00	15.00 :	2.94
4.03	4.03	47.00	4.00 :	11.47 :	1.99	2.00	42.00	15.00 :	2.88
4.05	4.05	47.00	4.00 :	11.14 :	2.05	2.06	42.00	15.00 :	2.82
4.07	4.06	47.00	4.00 :	10.81 :	2.10	2.13	42.00	15.00 :	2.76
4.08	4.08	47.00	4.00 :	10.47 :	2.16	2.19	42.00	15.00 :	2.70
4.10	4.09	47.00	4.00 :	10.14 :	2.21	2.25	42.00	15.00 :	2.63
NEW STATUS: own-ship: standon									
NEW STATUS: hazard ship: pass0_A									
4.11	4.11	47.00	4.00 :	9.81 :	2.27	2.31	42.00	15.00 :	2.57
4.13	4.12	47.00	4.00 :	9.88 :	2.33	2.37	53.94	15.00 :	2.51
4.15	4.14	47.00	4.00 :	-60.00 :	2.40	2.41	65.87	15.00 :	2.45
4.16	4.15	47.00	4.00 :	-60.00 :	2.48	2.44	77.81	15.00 :	2.40
4.18	4.17	47.00	4.00 :	-60.00 :	2.56	2.45	77.81	15.00 :	2.36
4.20	4.18	47.00	4.00 :	-60.00 :	2.64	2.47	77.81	15.00 :	2.31
4.21	4.20	47.00	4.00 :	-60.00 :	2.73	2.49	77.81	15.00 :	2.26

Table 1

Initial part of data-logging file for overtaking run.

(see fig. (5))

6. CONCLUSIONS AND FUTURE DEVELOPMENTS

Clearly negotiation of fixed obstacles, restricted channels, rogue ships and uncontrolled mobile hazards must all be catered for, as must satisfactory resolution of multi-ship encounters. The current primitive rule base is open to considerable extension and increased sophistication in its present form. Ultimately, probability considerations ('fuzzy logic') must be incorporated to enable a realistic assessment of the 'best' course of action in complex scenarios - it is anticipated that this could be assimilated into the existing framework without major restructuring. Given the current speed of the simulation process, there is substantial scope for increased complexity whilst still operating within real-time constraints. Application of fuzzy logic would include the imposition of confidence limits on sensory data received via instrumentation, and on the consequent action.

As indicated previously, it is intended at an early stage to handle the simulations of the two vessels on separate microcomputer systems, either or both of which may be under the direction of the embryo expert system; alternatively, one of the systems may be under operator control, or subject to programmed 'aberrations'.

In respect of the visual display and user interaction, it is anticipated that the features of GEM may be used more fully to offer more comprehesive review information, option selection and graphic detail selection using the mouse, review of course data, and other facilities as the need becomes apparent (A data-logging file is already created for post-mortem purposes - see Table 1).

A long-term objective is that of an expert system 'shell' for collision avoidance. In such a structure, it would be possible for the experienced user to extend the rule base, adding new encounter situations in terms of the possible combinations of parameter values giving rise to such situations and advisable responses to such situations; such a shell would also have value in encounters of controlled moving objects other than marine vessels. Major consideration must be given in such a system to (a) maintaining real-time operational capability, and (b) safeguarding against possible catastrophe by inexpert use of such a facility.

7. REFERENCES

1. Goodwin, E.M., 1975 "A Statistical Study of Ship Domains", Journal of Navigation, 28, 328-336.
2. Davis, P.V., Dove, M.J. and Stockel, C.T., 1980 "A Computer Simulation of Marine Traffic Using Domains and Arenas", Journal of Navigation, 33, 215-222.
3. Davis, P.V., Dove, M.J. and Stockel, C.T., op. cit.
4. Colley, B.A., Curtis, R.G. and Stockel, C.T., 1983 "Manoeuvring Times, Domains and Arenas", Journal of Navigation, 36, 324-328.
5. Konyn, M., 1987 "Collision Avoidance at Sea and on a Marine Radar Simulator using Automatic Encounter Detection Techniques", PhD Thesis, London.
6. International Marine Organization, 1981 "International Regulations for Preventing Collisions at Sea".
7. Colley, B.A., 1985, PhD Thesis, Plymouth, Devon.

COMPUTING AND CONTROL DIVISION
COLLOQUIUM ON
"CONTROL IN THE MARINE INDUSTRY"

ORGANISED BY
PROFESSIONAL GROUP C9 (CONTROL TECHNIQUES AND APPLICATIONS)

ON THURSDAY, 28 JANUARY 1988

DIGEST NO 1988/19

An Automatic Guidance, Navigation and Collision Avoidance System for Ships at Sea

R S Burns, G Blackwell and S Calvert

Introduction

It has been suggested that 85% of all marine collisions and groundings are due to human error. On this evidence alone there is a case for research and development into automating the control and guidance systems which are installed in ships. Furthermore this does not preclude relative levels of automation in the navigation and guidance process, neither does it preclude sub systems which advise the mariner of the action to take. For example, Automatic Radar Plotting Aids are generally accepted as taking a great deal of tedious work away from the Officer of the Watch, whilst providing him with much more information than he could obtain manually. The work of the Ship Control Group is directed towards development of these concepts.

The Group's activities are divided into navigation, control and guidance, collision avoidance and weather routeing. By the very nature of the research and development, mathematical modelling of ship systems is an all important aspect. This paper describes current work being undertaken on behalf of the DTI to develop a voyage management system for automatic guidance of a ship through all phases of a passage. The system will automatically maintain a vessel on a predetermined track, improve the probability of arriving in port on time in poor weather conditions, and reduce the workload on the bridge staff.

The Ship Guidance Problem.

In its research programmes the Ship Control Group has used stochastic optimal control theory [1, 2]. An important part of this theory is the separation principle, which allows a given optimisation problem to be reduced to two other problems whose solutions are known, namely an optimal filter in cascade with a deterministic controller. The problem considered has thus been the design of optimal multi-variable system to control simultaneously position and velocity of the vessel. The deviation from the desired values of these parameters can only be corrected, for most vessels, by operation of the rudder(s) and/or the main engine(s). A feature of an optimal system is that it will seek to maximise or minimise a global parameter, J, called the cost function or performance index. This is based upon the summation of the weighted errors over some time interval. Examples of the time interval might be the time to complete the pilotage phase of the voyage, the time to complete the oceanic phase of the voyage, or the total time from departure to arrival ports.

The pilotage, coastal and oceanic phases of the voyage may require different or modified cost functions. Whilst safety of the vessel, and hence collision avoidance, is paramount, the weighting of track keeping, positional accuracy, course keeping, speed and other parameters will change during a passage. For example in the approaches to a port it is vitally important that a large vessel keeps to the buoyed channel and hence automatic track keeping will have maximum weighting. In this instance the track is defined by the channel in to the port.

Ship Control Group, Plymouth Polytechnic

However once in to open water the navigator has a wider choice available. There will still be limitations in coastal waters, but once clear of the coastline the minimum distance will involve a great circle track, possibly taking the vessel in to bad weather areas. Hence the shortest distance may not mean the minimum time at sea. The cost function will now change and as weather data is received the on-board computers will be required to update the best route in order to avoid the worst weather. Alternatively the cost function might now be required to minimise fuel consumption, entailing longer time at sea, but possibly reducing voyage costs. There is, after all, no point steaming at 20 knots, only to find that there is no berth available at the arrival port. So port knowledge (berths and bunkers) could provide limiting constraints which will affect the oceanic performance index.

Optimal filtering, using a Kalman-Bucy filter, is a stochastic technique which combines noise corrupted measurements of a dynamic system with other known information about the system, in order to obtain best estimates of the variables, or states, that govern the system. It should be noted, however, that the process assumes that the system is linear and the errors are gaussian. As a ship constitutes a non-linear system, when parameters such as large alterations of course and/or speed, shallow water effects, and trim are considered there must be some limitation to the technique. In the work undertaken by the Ship Control Group the problem has been overcome by assuming constant course and speed during each sample period. This is reasonable providing sample times are small when compared with such factors as ship time constants and time between way points.

The Coastal Phase

Once clear of the port the coastal phase of the voyage is undertaken using the same cost function as for the harbour phase. The vessel is automatically maintained on track using information from available navigation aids to minimise track and course error. Theoretical studies are underway to incorporate an Automatic Collision Avoidance System (ACAS). This project is being undertaken in two areas. One research programme is concerned with applications of Artificial Intelligence to the International Rule of the Road [3]. The work is at present in the early stages and computer simulations are being undertaken. In another project researchers have been developing software to automatically guide the vessel when risk of collision or grounding exists. Computer studies have shown that it is possible for the cost function to be modified when an approaching ship poses a threat to the safety of navigation. The on board computer is programmed to change the weighting to allow for a collision situation and for the vessel to take avoiding action. Trials will commence shortly to ensure that the automatic collision system can be integrated with the navigation and guidance systems.

The Oceanic Phase

Having cleared the land the navigator will now wish to follow the most economic route to the coastline at the other end of an oceanic passage. Weather routeing [4] is now an established process in reducing overall running costs. In computer simulations the researchers have shown that it is possible for the cost function to be changed so that weather data transmitted via a satellite link can be used to plan and execute the most cost effective route for the vessel to follow. The technique of dynamic programming is being used to compute optimal routes. In recent simulations of a crossing from Cork to New York by the 37000 dwt Dart Atlantic during November 1986, when large depressions were forecast for the great circle route, the calculated optimum route was very close to that given by a weather routeing agency. It must be emphasised that the route planning is advisory in the initial stages of the research. However it is intended that this information will also be fed to the optimal controller so that the vessel can automatically follow the most cost effective route during the oceanic phase of the voyage.

Conclusions

The Ship Control Group's ultimate aim is to produce a voyage management system which takes navigational data from such navigation aids as are fitted on the vessel, collision avoidance data from the ship's radars, together with meteorological data from shore stations via a satellite link [5]. After appropriate filtering, this data will be used as input to a multivariable optimal controller which will maintain the vessel on the correct track between posrts, with due consideration to safety, efficiency and economy of operation. This Group is some way towards achieving this aim, as the navigation and guidance system has been tested afloat, whilst simulation studies have demonstrated that the automatic collision avoidance and weather routeing systems can be incorporated in to the overall voyage management system. The next step is to develop a complete prototype system and undertake extensive trials in coastal and oceanic waters.

References

1. M J Dove, R S Burns and J L D Evison, "The Use of Kalman Filters in Navigation Systems - Current Status and Future Possibilities", Proc Conf Computer Aided Design, Manufacture and Operations in the Marine and Offshore Industries, Washington DC USA, (Sept 1986), Computational Mechanics Institute.

2. R.S Burns, M J Dove and C T Stöckel, "A Multivariable Controller for the Automatic Guidance of Large Ships in Confined Waters", Proc 1986 Summer Simulation Conference, Reno Nevada USA, (July 1986), Society for Computer Simulation.

3. G K Blackwell, R S Burns and C T Stockel, "An Expert Systems Approach to Collision Avoidance", Proc 1987 Summer Simulation Conference, Montreal Canada, (July 1987), Society for Computer Simulation.

4. R H Motte, R S Burns and S Calvert, "An Overview of Current Methods Used in Weather Routeing." J Nav, (1988), To be Published.

5. M J Dove, R S Burns and K M Miller, "A Voyage Management System To Improve the Economy of Operations of Ships", Proc 8th Biennial Seminar on "The Economic Ship", Singapore, (Dec 1987), Joint Conf Inst of Marine Engineers and Society of Naval Architects and Marine Engineers

AIENG88 THIRD INTERNATIONAL CONFERENCE ON APPLICATIONS OF ARTIFICIAL INTELLIGENCE IN ENGINEERING 8–11 AUGUST 1988

MEMO TO: Contributors to Third International Conference on Applications of Artificial Intelligence in Engineering

FROM: Professor John Gero
Department of Architectural Science
University of Sydney NSW 2006
Australia

Your contribution has the following reference:

Blackwell, G. K., Colley, B. A. and Stockel, C. T. (1988). A real-time intelligent system for marine collision avoidance, *Artificial Intelligence in Engineering: Diagnosis and Learning*, Computational Mechanics Publications, Southampton, pp.119-138.

Thank you for all your efforts. I am also enclosing the title page and list of contents.

26 April 1988

A REAL-TIME INTELLIGENT SYSTEM FOR MARINE COLLISION AVOIDANCE

G.K.Blackwell, B.A.Colley, C.T.Stockel

Ship Control Group, Plymouth Polytechnic, Plymouth, Devon, England

ABSTRACT
The task of collision avoidance at sea is, by definition, rule-based. Mariners are in fact constrained by a set of rules governing encounters between vessels at sea. However, these rules are not totally definitive as to details of timing, change of course, or acceptable clearances. Such details are left to the judgement of the mariner - judgement which improves over years of experience and accumulated wisdom. By the same token, interpretation of complex situations involving a number of hazards requires skill, judgement and experience, as well as a book of rules.

Such considerations identify this task as an ideal candidate for a computer-based expert system, preferably housed in a microcomputer situated on the ship's bridge. In marked contrast to most expert systems, however, such a collision avoidance system would be characterised by:
(a) Sensory input of dynamically-changing data, such as ship's speed and course, also range and velocity of any potential hazards;
(b) Substantial mathematical processing, notably evaluation of a variety of intricate trigonometrical functions;
(c) Ability to respond in real time, whilst re-evaluating the situation continually, as would the expert mariner.

These requirements suggest the need for a combination of software tools: a programming language capable of control-level operations, incorporating fast and powerful mathematical processing capability; coupled with a user environment which supports the hierarchical rule structure associated with such

systems, as well as providing a well-structured and flexible user interface. Detailed consideration of the latter requirement suggests that an object-oriented language, such as Smalltalk, is better suited than a list-processing or Prolog-type language.

Development is under way on such an expert system, using the 'C' language to handle sensory data input and mathematical function processing. Work is currently at the 'second prototype' stage, having evaluated the effectiveness of the first prototype by a multiple run of 500 simulated encounters. The current model, written completely in 'C' and operating under the GEM environment, incorporates an expert system module and a simulator module, each operating independently on each of two knowledge bases, one for own-ship, the other for one hazard vessel. Multiple windowing provides the user with current information, both graphical and numerical, on ships' status, plus the option to review the reasons for decisions taken.

The original prototype operated the rule base as a simple decision tree within the body of the expert system program module; mark II uses a content-free inference engine operating on explicitly-defined rules, with consequent actions. Current development is concerned with provision of an object-oriented 'front-end' module, responsible for decision-making and interfacing with the user. The next move should be to isolate the expert system for own-ship, with all extraneous influences (simulated, for the time being) updating the knowledge base via communication ports.

The prime objective of this work is to provide an expert adviser for mariners, capable of receiving sensory input from all available channels - including information provided by an operator via a keyboard - and advising on the preferred course of action.

1. INTRODUCTION

It is almost axiomatic to say that collision avoidance at sea requires a rule-based approach; this is true whether or not computers are involved in the process. Indeed, a superficial appraisal of the problem might elicit the conclusion that a simple procedural application of a set of rules (or conditions), leading to consequent actions, would suffice - is not just such a set of rules in constant use by mariners?

Such a procedural approach has been used very effectively in computer simulation of marine traffic flow and collision avoidance [1], with benefits in such areas as training of mariners and modelling the effects of changes (eg in traffic volume) in high-density traffic lanes. However, this approach is not a practical option in real-world situations, ignoring as it does the need for experience and commonsense in application of these rules.

The International Regulations for Preventing Collisions at Sea [2], although quite specific as far as they go, cannot be regarded in any sense as rigorous definitions of preconditions and consequent actions. Details as to timing, clearances, suitable course alterations, are left very much to the discretion of the mariner - discretion which is perfected by years of accumulated experience and wisdom. Furthermore, the variety of 'non-standard' situations - particularly those involving a number of vessels - is such that no rule book could adequately incorporate all such contingencies and still provide a useful reference.

Clearly it is possible to provide 'rules of thumb' for specific types of situation, and to add supplementary rules for variations in these situations. This rule structure would be founded on the previously-mentioned anti-collision regulations, would incorporate the accumulated wisdom of expert mariners, and would presumably also be tailored to reflect the response characteristics of the system operating that rule structure - speed of response, breadth of information available to the system, possible consequence of misjudgement (confidence limits). The information referenced by such a structure would be of two types: static and dynamic.

Static information relates to fixed characteristics of the vessel (some possibly specific to current voyage): length, beam, maximum speed, minimum turning circle, safe clearing distance, and a variety of technical data; relevant chart data for the area being navigated could also be regarded as static information, though of a rather different type. Dynamic information, to be constantly updated, would include such considerations as current speed and course, plus data on position, speed and course of any potential hazards in the vicinity. It should be noted that there is some degree of overlap between these two types of data, for example one's assessment of a safe clearing distance may depend on the nature of the hazard to be cleared.

The task as defined is clearly susceptible to handling by means of a suitable expert systems package, ideally running on a microcomputer system located on the bridge of the ship. However, two major considerations set this task apart from most conventional applications of expert systems experienced to date:

(1) The dynamic information required must, by its very nature, be input to the system directly via a range of sensors, including such instrumentation as radar; furthermore, the scope and limitations of such data must be given due regard by the expert systems logic.

(2) The decision processes involved require extensive mathematical calculations, primarily of trigonometrical functions, to be carried out in real time.

These considerations would appear to preclude any of the recognised artificial intelligence languages as a suitable vehicle for the major processing requirements of such a system. A fair compromise would seem to be a 'calculation' module, written in a language suited to control and mathematical processing, front-ended by a 'decision/user-interface' module set in an environment more suited to rule-handling and man-machine interface (MMI) considerations. This is the direction of the research presented in this paper.

2. SYSTEM SPECIFICATION

An intelligent response system of the type proposed here should, in practical terms, be capable of evaluating such an encounter between two or more vessels from the standpoint of an experienced mariner, and taking appropriate action. It should be noted at the outset that such a parallel is necessarily limited in its application, since the mariner may have access to information which is not quickly and easily available to a computer system, and vice versa. For example, change in aspect of an approaching vessel will immediately be visually apparent to a human observer, whereas computer interpretation of radar data may require several well-spaced readings before producing a corresponding conclusion; conversely, the computer could probably deduce the effect of a change in rudder angle, in terms of closest point of approach to a hazard, far more quickly and accurately than the most experienced ship's master. Such considerations should be borne in mind both in the process of knowledge elicitation and in creating/extending the knowledge base. Subject to these provisos our system, on being provided with the information normally available on the bridge of a well-equipped ship, should direct the course of action of that vessel as would a competent mariner. This involves:

(a) Recognising an encounter (potential hazard) situation in good time;
(b) Identifying the type of encounter and the status of one's own vessel (own-ship) in that encounter, according to the International Regulations for Preventing Collisions at Sea - such status would normally be either 'stand-on' (having right of way) or 'give-way', according to one's position relative to the other vessel (subject to types of vessel involved, fishing vessels and deep-draught vessels having right of way in many circumstances);
(c) Choosing a course of action which combines a sensible safety margin with the least practicable inconvenience (where own-ship is judged to be stand-on, this will - initially at least - mean no alteration of course);
(d) Maintaining a watching brief on the other vessel(s) in the encounter, and being ready to take avoiding action if necessary, should the situation change - it is not unknown for 'rogue' vessels to disregard the regulations and plough straight on when they should give way, or even to turn into danger.

Clearly implicit in this 'seaway code' is the understanding that any other vessel encountered will be assumed to be 'well-behaved' until proved otherwise, but that an identical perception of, and reaction to, the encounter by the master of another vessel cannot be taken for granted. It is also imperative always to adopt a give-way stance in relation to any hazard other than another vessel under control, for example a derelict or a rock!

3. CURRENT DEVELOPMENTS

A first prototye system, written completely in 'C', dealt with open-sea encounters (not in restricted channels or routing systems, no fixed obstacles) between two manned vessels moving subject to fairly idealised parameters [3]. The complete system comprised an 'expert system' module, a simulator module, and a knowledge base for each of the two vessels, own-ship and hazard. At 20-second intervals in real time (each simulated by approximately one second of computer time), the expert system module would assess the status of own-ship by reference to its knowledge base; indicators being set in that knowledge base for action to be taken. The knowledge base for the hazard vessel would likewise be examined and flagged for subsequent action, independently, by the same expert system module. The simulator module would then update the dynamic information for each vessel in turn by reference to relevant parameters, at the same time actioning any status flags set for that vessel by the expert system. Since under normal conditions certain information on each vessel would be available to the master of the other ship via radar and other instrumentation, such information was interchanged between the two knowledge bases by the simulator module. The central requirement of the expert system was to ensure that no hazard violated an area of proximity around the vessel, referred to as the 'domain' - more information on this and other associated concepts is given later in this paper.

Fig. 1 shows an annotated example of the multi-window screen display generated by the prototype system.

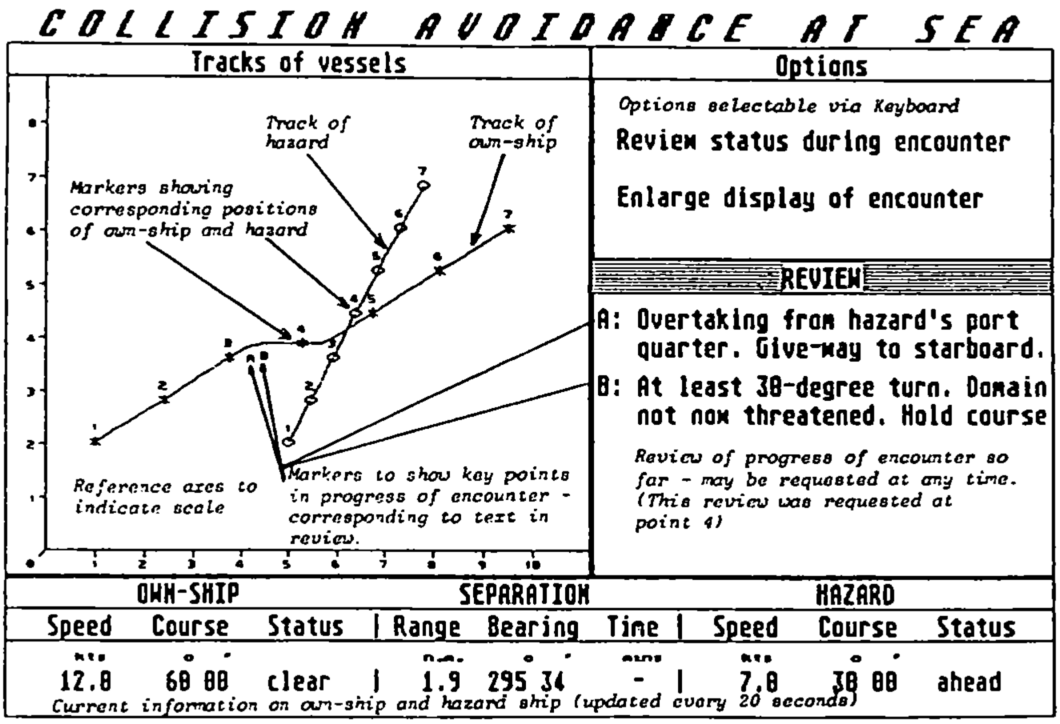

Figure 1 : Annotated example of the screen display generated by the prototype system.

This display format is substantially the same for the second (current) prototype. Further examples of screen displays for a variety of encounter situations are given later in this paper.

The effectiveness of this first prototype was evaluated by a multiple run of 500 simulated encounters, each generated as follows: two vessels, each with random course, speed and turning circle (the latter two parameters being within specified bounds), were placed in close proximity to the extent that each was very likely to consider the other a hazard, according to their terms of reference. The two vessels, one designated own-ship, the other hazard-ship, were then effectively taken back in time ten minutes along their tracks, and the simulation allowed to proceed from that point - if an encounter situation already existed by that time, the time-reversal was extended to a point just before that encounter before letting the simulation proceed. The separation at closest point of approach (CPA) between the two vessels was calculated (i) with no collision avoidance manoeuvre attempted, (ii) subject to the collision avoidance manoeuvre brought into play by the expert system module. The results of this exercise are shown in the bar chart in Fig. 2.

COMPARISON OF CPA's FOR A RANDOM SAMPLE OF 500 ENCOUNTERS

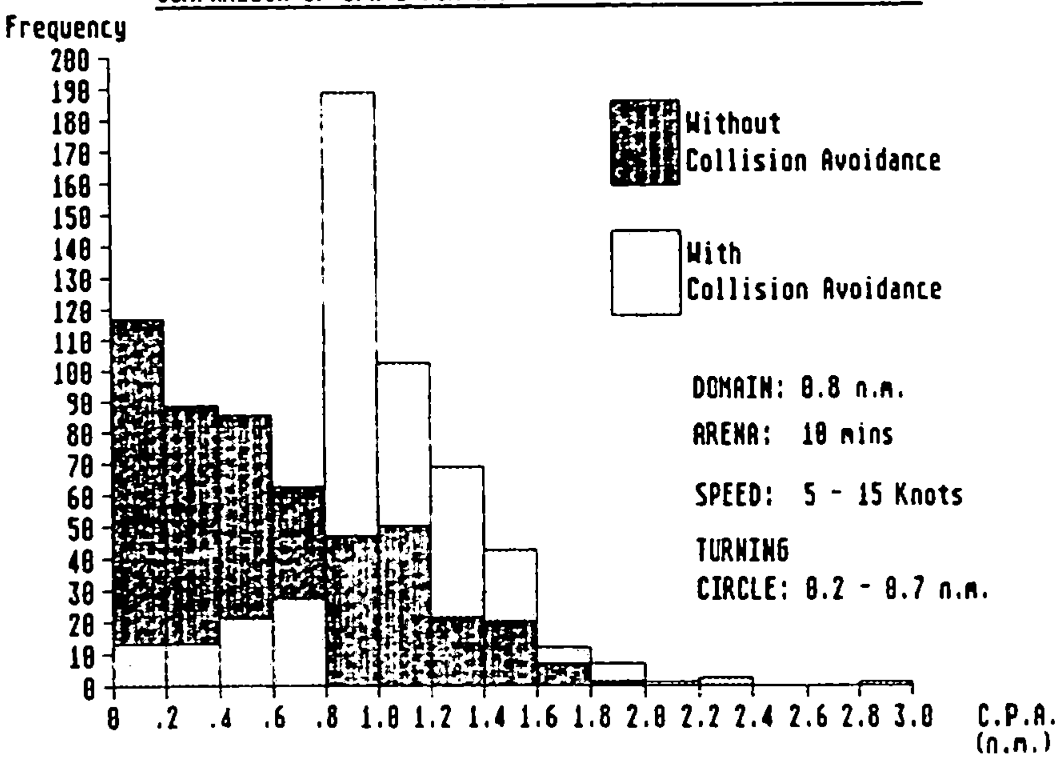

Figure 2

The results of this multiple simulation illustrate that, in the large majority of cases, potential domain violations were avoided by the expert system invoking appropriate collision avoidance strategies; moreover, the manoeuvres involved did not, in the main, involve excessive course alterations - CPA separations for these manoeuvres are clearly bunched just outside the domain boundary. However, a small number of the encounters still show CPAs within the domain boundary - this is clearly unacceptable. The reasons for this are considered to be threefold:

(1) The rule base in the initial expert system module was shown to have certain 'blind spots', in which each vessel assessed the other as the give-way vessel, and the required avoidance action was not taken;

(2) In using a fairly simple model for generation of simulated vessels, some realism was lost - vessels with unlikely combinations of characteristics were involved in a number of the offending encounters;

(3) A small number of these simulations highlighted the need for an increased level of sophistication in the strategy for recognising and handling encounter situations. Specifically, the fixed look-ahead, used to assess the time to initiate avoidance manoeuvres, appeared inadequate in certain cases; an ongoing look-ahead, extrapolating at regular intervals from current status, seemed indicated to ensure action at the optimum time. This is considered to be the most important development to arise from the multiple simulation exercise.

All of these points have been taken on board (!) in developing the latest, second, prototype, also operating completely in 'C' under GEM, also driving two vessels through a simulator module operating under the direction of an expert systems module. The greatest change, however, is in the nature of the expert system. In the first prototype, the rule base was fixed, encoded in the program logic. In this later version, the rule base has been separated out as a hierarchical set of conditions with associated actions - with the potential for extension and/or modification, given a suitable user interface. A content-free inference engine applies those rules to the information for each vessel in turn.

It should be made clear that this prototype is intended purely as a vehicle to test certain concepts, and in this respect has proved most useful. It has made possible development and testing of a fundamental rule structure, and enhancement of that structure in the light of experience. In this context, it has provided a test-bed for the intricate trigonometrical functions essential to the evaluation of those rules, and a preliminary check on the viability of real-time operation of such a system. Later in this paper, consideration is given to the question of setting these initial developments in a more appropriate environment for decision-making and user interaction.

4. PRINCIPLES OF OPERATION

As previously indicated, the total system drives two simulated vessels, each with its own operational parameters, and each (independently) subject to the same expert system module. It follows that both vessels should behave in an orderly manner, as prescribed by the anti-collision regulations, and that neither vessel will have to take emergency evasive action at a late stage to combat non-cooperation by the other vessel - i.e 'rogue' ships are not a consideration at this stage of development. It is intended, in the near future, to model the two ships on separate computers, with interaction via a communications link; it is further intended to provide a manual override on one of the ships - the hazard - to facilitate testing of an extended rule base aimed at handling such 'rogue' behaviour.

At the present phase of development, the rule base consists of rules for:

(a) Identifying the presence of a potential hazard, and assessing degree of threat in terms of expected time to infringement of the domain - this time factor is central to the decision on when to take avoiding action, if necessary;
(b) Identifying the type of encounter, including several variants in some cases (the primary types of encounter being head-on, crossing and overtaking), and fixing the status of own-ship and the perceived status of hazard-ship in the encounter;
(c) Negotiating the stages of the encounter, with due regard for appropriate safety margins.

The rules identifying and handling the three types of encounter are as follows:

(1) Head-on
Two vessels are approaching each other on (almost) reciprocal bearings. A vessel in such an encounter must consider itself give-way and alter course to starboard (right); it should not begin to alter back onto course until the other vessel has passed 90 degrees on the port beam (lefthandside);

(2) Crossing
Two vessels are approaching each other at an angle such that each can see one side of the other; more specifically, one vessel is in the sector subtended by the port or starboard light - 112.5 degrees measured from the bow (front) on each side. The ship with the other vessel to starboard should give way, altering course to starboard and going behind. It should not alter back until the stand-on vessel has passed 15 degrees to port of own-ship's original heading;

(3) Overtaking
A vessel is deemed to be overtaking if it is approaching from within the stern-light sector, ie the rear 135-degree arc. The overtaking vessel should give way as follows:
(a) If on near-parallel course, go behind to starboard; alter-back when safe to do so;
(b) If on port quarter of stand-on vessel, action as for (a);
(c) If on starboard quarter, and set to overtake ahead on current course, **and** if able to parallel courses without infringing domain, parallel-up; when well past and clear to pull across, return to original heading;
(d) Otherwise, go behind to port, alter-back when safe to do so.

The inference engine for the current system rests primarily on previous research in the field of marine collision avoidance. In particular, two well-researched concepts dictate the nature and extent of action taken: the domain, or 'danger zone' around a vessel, into which no other vessel or obstacle must be allowed to encroach; and the arena, or 'circle of influence', a larger zone around a vessel, inside which any other vessel or obstacle must be considered a potential hazard.

The concept of the domain was advanced by Goodwin

[4] as 'the effective area around a ship which a navigator would like to keep clear with respect to other ships and stationary objects'. The original concept of three distinct circle sectors of differing radii for sidelights and sternlight was later modified [5] to give a continuous circular domain boundary, with the ship off-centre in the circle so as to preserve areas subtended by sidelights and stern light. This rationalisation was based on the logical necessity for continuity of the domain boundary. (see fig. 3). The expert system (current and anticipated) has as its first objective the protection of the integrity of the domain. This is currently simplified to a circle centred at own-ship; thoughts for the future incline towards the 3-sector domain, with 'fuzzy' (probability-weighted) boundaries between the sectors.

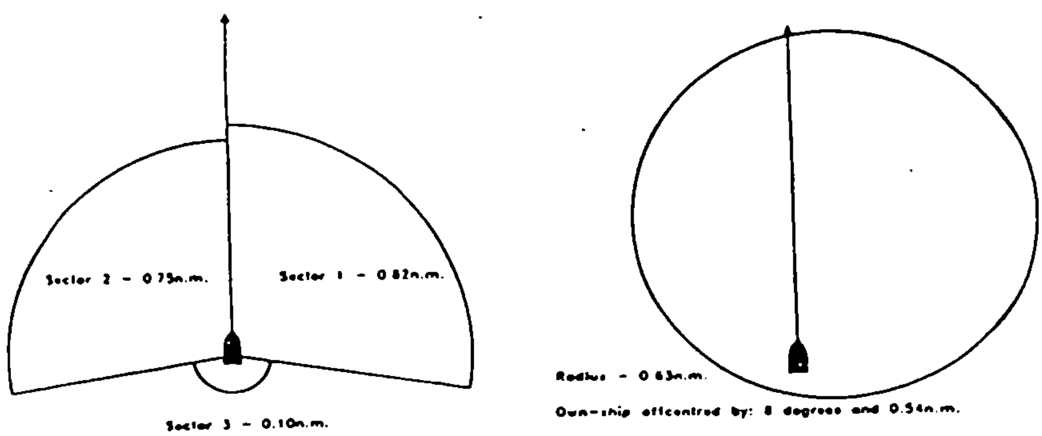

(a) Goodwin domain (example) (b) Davis domain (example)

Figure 3.

The arena concept was introduced by Davis [5] in response to the observed actions of mariners manoeuvring to prevent infringement of their domain by ships in their vicinity. This concept was refined by Colley [6] to a time factor, referred to as the RDRR (Range-to-Domain/Range-Rate) criterion - basically the time-to-go (before calculated domain infringement) at which the mariner will assess the status of his vessel and take action accordingly. For example, on the basis of his ship's handling characteristics a mariner may choose to take action 15 minutes before potential domain infringement - a 15-minute RDRR. A modified RDRR criterion, proposed by Konyn [7], gives implicit recognition to the fact that mariners' responses vary according to the nature of the anticipated encounter; this allows for reaction times required to recognise and respond to different types of encounter.

The RDRR criterion, as defined by Colley, is used in the model to determine decision-time. It had been planned to extend this to Konyn's RDRR+ criterion, however this plan has been rather overtaken by events. As indicated earlier, the RDRR look-ahead appears inadequate for certain juxtapositions of vessels, and developments are now in hand to provide an ongoing look-ahead. This is effectively a repeated forward-looking simulation (within the simulation provided by the model) - which should yield a variable RDRR value to suit the specific situation. This will, of course, add to the overheads in terms of processing time, emphasising the need for fast mathematical processing to meet the essential requirement of real-time operation.

One important aspects of collision avoidance manoeuvres is that they must be clearly visible and identifiable as such from another vessel - no other form of communication between vessels can be assumed. Consequently, any manoeuvre undertaken at the direction of the expert system must be of sufficient magnitude (a 'clear turn') and in good time to be seen from another vessel; such considerations must be set against the temptation to use the power of the computer to 'optimise' course changes purely on the basis of 'safe clearance' criteria. The current model uses a minimum turn of 30 degrees, and a minimum RDRR criterion as defined as a parameter for each vessel - typically 10 minutes, but varied to suit the handling characteristics of the vessel being simulated.

The rules as implemented embody the International Regulations for Preventing Collisions at Sea, plus clarification of specific practical details as observed by Colley [8] and used most effectively in a previous simulation. Using these rules, every collision avoidance manoeuvre is taken through four stages, as follows:

(1) When the RDRR criterion is infringed, the give-way vessel initiates a turning action;
(2) Once the required turn is effected (minimum 30 degrees), the vessel ceases its turn;
(3) When the original heading may be resumed without further risk of domain infringement (and any other conditions are satisfied), alter-back is initiated, necessitating a turn in the opposite direction;
(4) When back on the original heading, the manoeuvre is terminated, the vessel proceeds on that heading.

It should be noted that the function of the inference engine at each time-step is to take the basic parameters (speed, heading, position) provided by the simulator. From these it generates a standard set of further parameters - relative bearing, relative velocity, angle between relative bearing and relative heading, etc - and applies the rules to them to ascertain which situation pertains to the next step.

In the exended look-ahead mode under development, for use with a variable RDRR criterion, the system will, at regular intervals, extrapolate from current data through the anticipated manoeuvre. Hence the optimum time-to-go for initiation of the appropriate collision avoidance strategy is identified.

5. EXAMPLES

Figure 1 shows an overtaking manoeuvre with an approach from the port quarter. Own-ship has given-way to pass behind the hazard ship before correcting-back.

Figure 4(a) shows the first stages of a head-on encounter. Note that the status of own-ship is known in detail, but hazard ship is known only to be changing course. Figure 4(b) shows the completed manoeuvre, with both vessels back on course.

Figures 5(a) and 5(b) show the two possible manoeuvres for overtaking from starboard side. In 5(a) own-ship is able to parallel-up and pass the hazard vessel, to overtake ahead without infringing the domain; in 5(b), circumstances such as speed and turning circle rule this out, so own-ship passes behind hazard vessel.

Figure 6 shows hazard vessel overtaking from own-ship's port quarter, by altering to starboard and passing behind before correcting-back. Table 1 is a printout of the first few entries in the datalogging file for this encounter, which could be used to reconstruct or analyse the event.

Figure 7 shows a crossing manoeuvre, in which own-ship has given-way to hazard vessel.

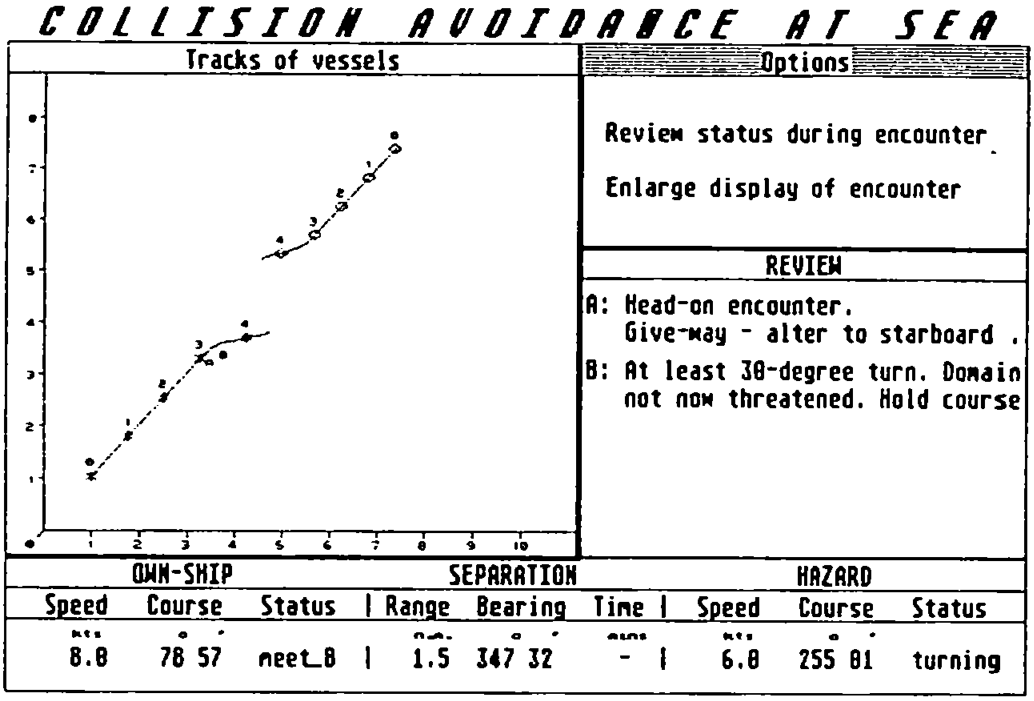

Figure 4(a). Head-on Encounter - Initial stages.

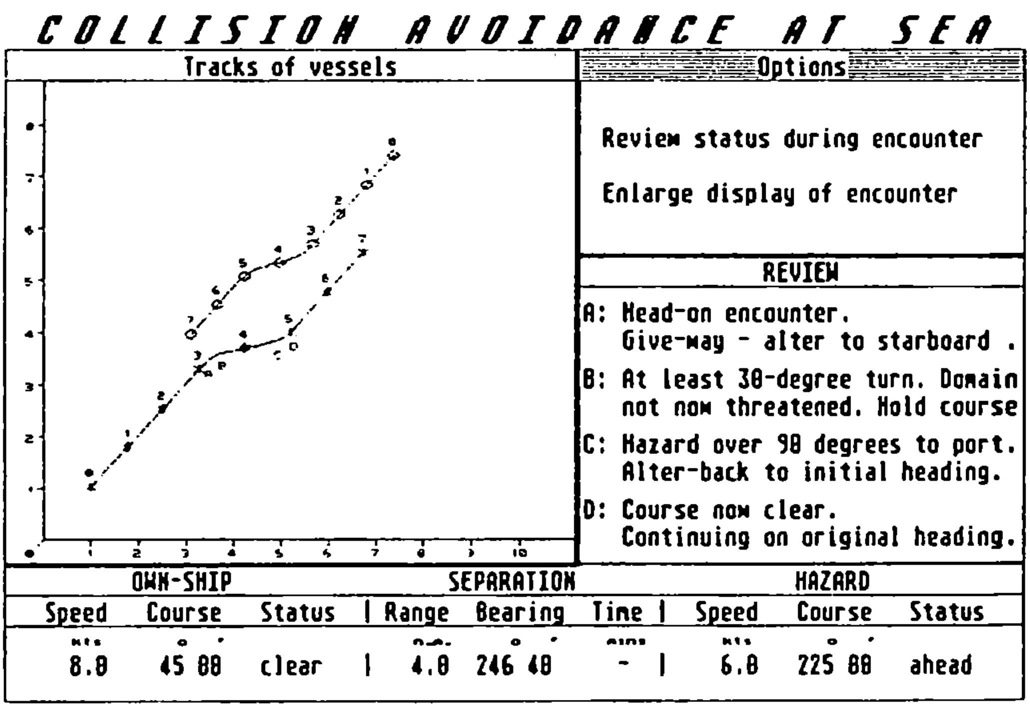

Figure 4(b). Completed Head-on Encounter.

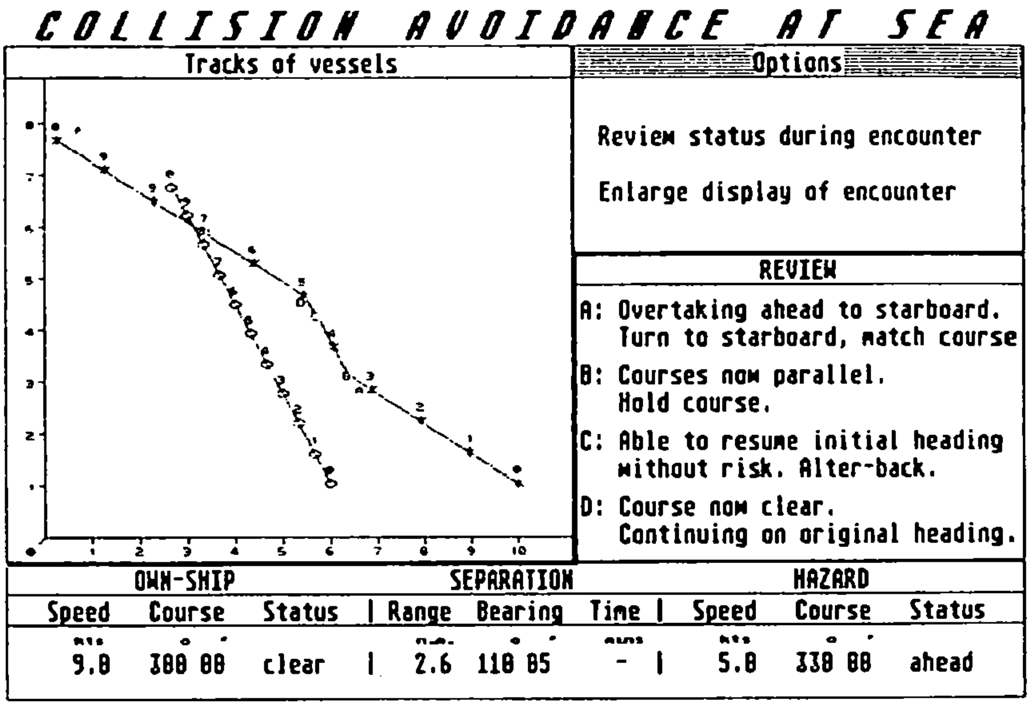

Figure 5(a). Overtaking from starboard side (i).

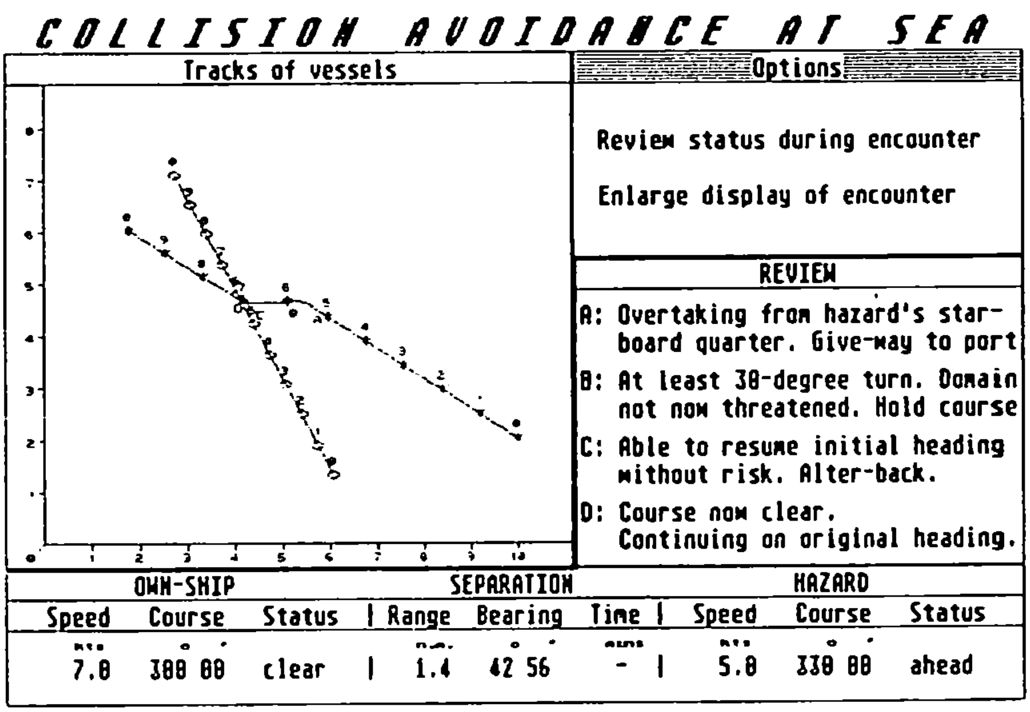

Figure 5(b). Overtaking from starboard side (ii).

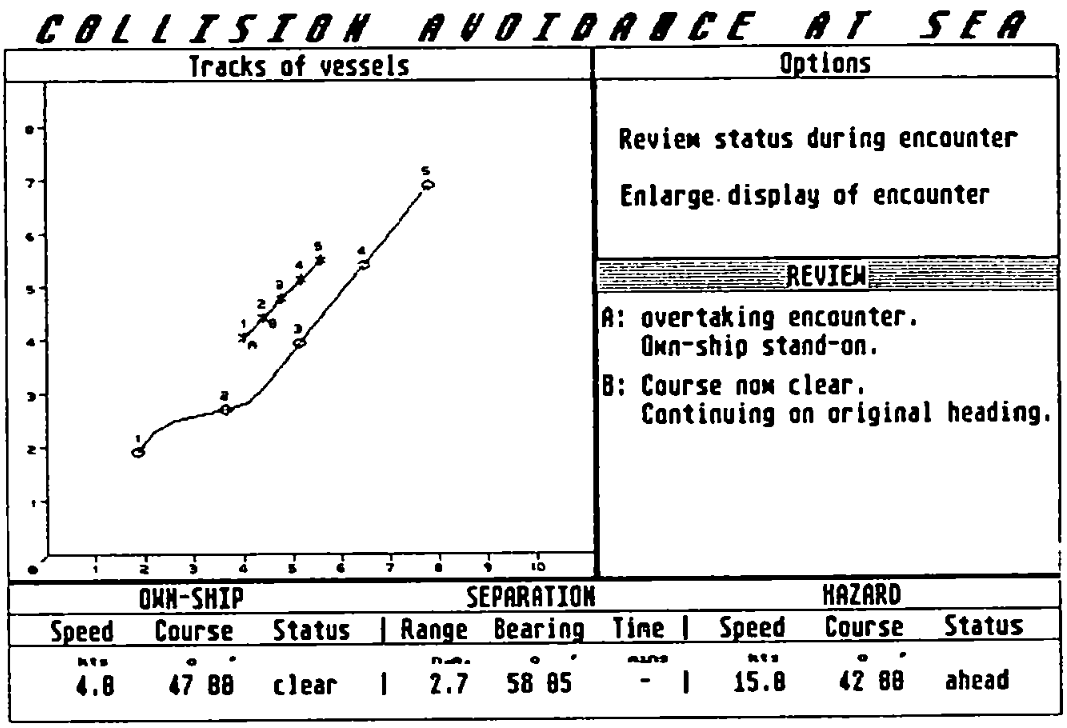

Figure 6. Hazard overtaking from port quarter.

data for own-ship

Course:-	(i) Degrees:	47.00
	(ii) Minutes:	0.00
	Speed in knots:	4.00
	Turning circle (n.mls):	0.60
	Position:- x:	4.00
	Position:- y:	4.00
	Domain radius in mls:	0.80
	Arena (RD/RR) in minutes:	10.00

data for hazard ship

Course:-	(i) Degrees:	42.00
	(ii) Minutes:	0.00
	Speed in knots:	15.00
	Turning circle (n.mls):	0.40
	Range in nautical miles:	3.00
Bearing:-	(i) Degrees:	225.00
	(ii) Minutes:	0.00

OWN-SHIP X	Y	COURSE	SPEED	TIME TO GO	HAZARD X	Y	COURSE	SPEED	RANGE
4.00	4.00	47.00	4.00	12.14	1.88	1.88	42.00	15.00	3.00
4.02	4.02	47.00	4.00	11.81	1.93	1.94	42.00	15.00	2.94
4.03	4.03	47.00	4.00	11.47	1.99	2.00	42.00	15.00	2.88
4.05	4.05	47.00	4.00	11.14	2.05	2.06	42.00	15.00	2.82
4.07	4.06	47.00	4.00	10.81	2.10	2.13	42.00	15.00	2.76
4.08	4.08	47.00	4.00	10.47	2.16	2.19	42.00	15.00	2.70
4.10	4.09	47.00	4.00	10.14	2.21	2.25	42.00	15.00	2.63

NEW STATUS: own-ship: standon

NEW STATUS: hazard ship: pass0_A

X	Y	COURSE	SPEED	TIME TO GO	X	Y	COURSE	SPEED	RANGE
4.11	4.11	47.00	4.00	9.81	2.27	2.31	42.00	15.00	2.57
4.13	4.12	47.00	4.00	9.88	2.33	2.37	53.94	15.00	2.51
4.15	4.14	47.00	4.00	-60.00	2.40	2.41	65.87	15.00	2.45
4.16	4.15	47.00	4.00	-60.00	2.48	2.44	77.81	15.00	2.40
4.18	4.17	47.00	4.00	-60.00	2.56	2.45	77.81	15.00	2.36
4.20	4.18	47.00	4.00	-60.00	2.64	2.47	77.81	15.00	2.31
4.21	4.20	47.00	4.00	-60.00	2.73	2.49	77.81	15.00	2.26

Table 1. Start of datalog file for above encounter.

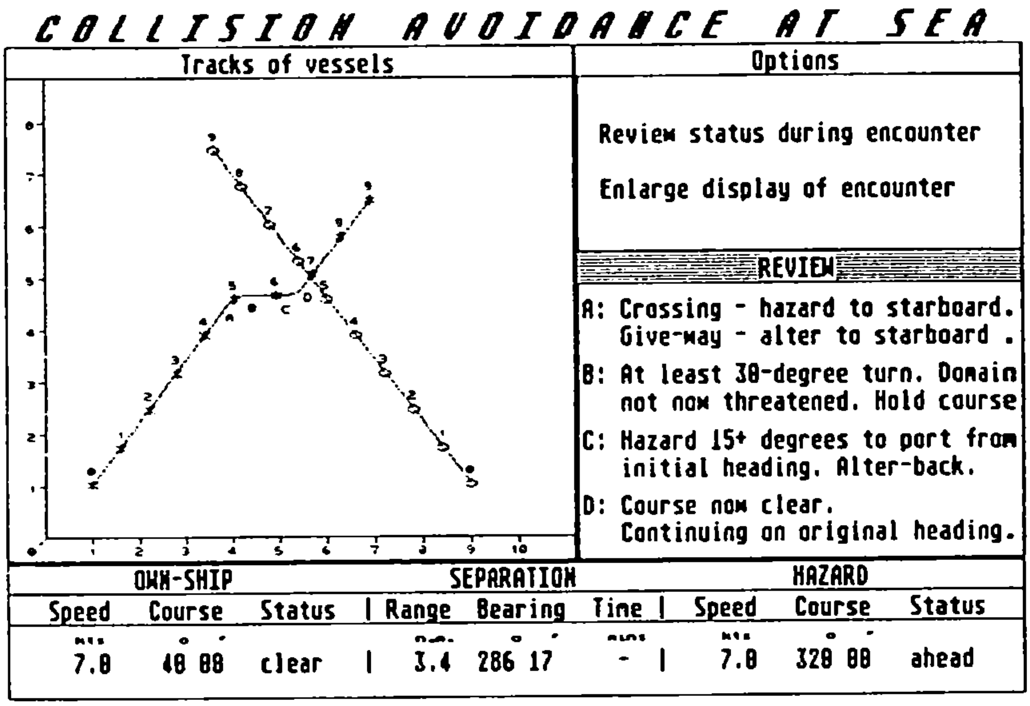

Figure 7. Crossing manoeuvre, own-ship giving way.

6. PROJECTED DEVELOPMENTS

The ultimate objective of this work, a bridge-based real-time expert adviser for marine collision avoidance,is clearly some way off, though current work suggests that it is achievable. It is perhaps worth identifying the major hurdles still to be cleared, and how these are viewed in the light of the work so far.

(1) Rogue vessels
The rule base can be extended, within the present structure, to incorporate avoiding action in cases of rogue vessels, vessels not under command, and fixed hazards.

(2) Multi-ship encounters
In manoeuvres involving more than one hazard, a two-part strategy is envisaged. First, the system would attempt to identify at an early stage a course alteration which would minimise, or possibly remove, the confliction between own-ship and the multiple hazard. This is a frequently-observed ploy of masters on cross-channel ferries. Secondly, the look-ahead technique would be applied to assess the consequences of various avoidance manoeuvres,

considering a sequence of single encounters. It is anticipated that this would be a generative process, building up a projected sequence of steps using GPS (General Problem Solver) techniques. Probabilistic constraints are also likely to be applied, to assess the 'best' course of action.

(3) User Interface
In its present form, the system can handle a small, fixed set of user requests, giving standard, pre-programmed responses. There is no facility at present for the user to extend or amend the rule base, nor is this considered appropriate for this 'test-bed' version. Serious consideration is currently being given to a suitable user environment for the final system, enabling proper access to the rule base and having the capacity for a flexible response. The nature of the entities and the tasks involved tends to indicate an object-oriented approach, rather than list-processing or so-called logic programming. There is no doubt, however, that speed is of the essence in the mathematical processing, if real-time operation is to be achieved; such languages do not tend to put speed of execution high on the priority list. For this reason, it is planned to retain 'C' as the 'number-cruncher', whilst Smalltalk is under consideration for the user environment. Smalltalk provides a very flexible and 'user-friendly' environment, and would appear to provide the potential for interfacing 'C' routines. An alternative, also now under consideration, is 'Objective C', which at first sight appears to combine the best of both worlds - a full appraisal has yet to be undertaken.

(4) Isolation of the Expert System
The development system has of necessity included a simulated environment for the expert system, namely computer-generated data pertaining to own-ship and hazard vessel. The final expert system will consist of a stand-alone unit, receiving data via a variety of communication channels, in addition to the user's console. As an intermediate stage, the internal data links to the current expert system module are to be replaced by communication links to one or more other computers, which will simulate dynamic data for own-ship and radar/communication data for the hazard vessel. Ultimately, of course, the system will be tested, and final development undertaken,

in a real operational environment.

(5) Validation of the final system
Three phases of validation are planned:

(a) The advanced simulation described in (4) above is to be used by a number of experienced mariners. These mariners will be able to compare the system's handling of each encounter with their own views on how such a situation should be dealt with. As well as providing system validation, such input of experience should facilitate enhancement of the system's rule structure.

(b) It is hoped to interface the expert system to Plymouth Polytechnic's Racal MRNS 9000 navigation simulator, via RS232 communication links currently being set up. This simulator allows up to 4 bridge teams each to control a simulated vessel. Experienced mariners, each at the 'helm' of a simulated vessel, will be able to judge the expert system's handling of various encounter situations, viewing the system effectively as another helmsman.

(c) The system is to be installed in the Polytechnic's own research vessel for comprehensive sea trials. It may also be possible to instal the system in other vessels, to give a wider range of trial situations.

7. CONCLUSIONS

At its current stage of development, the anti-collision system has indicated the viability of the proposed product. Tests of the system have enabled enhancement of the rule base, and thus system operation, as well as pointing the way for future development. A great deal still remains to be done, but the feasibility of the project - from the computing standpoint - seems assured.

One area not tackled in this paper is the nature and quality of data provided by radar and other sensory inputs. The reliability of the inferences drawn by the system is, of course, totally dependent on such data. At a later stage, it will be necessary to set confidence limits on such inferences, having first ensured the optimum usage of such sensory data as is available. Cost of the overall system is a major factor in such considerations, since various levels of sophistication are available in such instrumentation - at a price. The operator will, of course, have the option to input via the keyboard information not otherwise directly available to the system.

8. REFERENCES

1. Colley,B.A., Curtis,R.G., and Stockel, C.T., 1984. On Marine Traffic Flow and Collision Avoidance Computer Simulation. Journal of Navigation. Volume 37, pp232-250.

2. International Marine Organisation, 1981. "International Regulations for Preventing Collisions at Sea".

3. Blackwell,G.K., Burns,R.S., and Stockel, C.T., 1987. An Expert Systems Approach to Collision Avoidance. Proceedings of Summer Computer Simulation Conference, Montreal. Society for Computer Simulation, July 1987, pp43-48.

4. Goodwin,E.M., 1975. A Statistical Study of Ship Domains. Journal of Navigation. Volume 28, pp328-344.

5. Davis,P.V., Dove,M.J., and Stockel, C.T., 1980. A Computer Simulation of Marine Traffic Using Domains and Arenas. Journal of Navigation. Volume 33, pp215-222.

6. Colley,B.A., Curtis,R.G., and Stockel,C.T., 1983. Manoeuvring Times, Domains and Arenas. Journal of Navigation. Volume 36, pp324-328.

7. Konyn,M., 1987. "Collision Avoidance at Sea and on a Marine Radar Simulator using Automatic Encounter Detection Techniques", PhD Thesis, London.

8. Colley,B.A., 1985. Computer Simulation of Marine Traffic Systems. PhD Thesis.

A WORKING CONFERENCE FOR USERS OF SIMULATION HARDWARE, SOFTWARE AND INTELLIWARE

PROCEEDINGS
OF THE
USER 1 WORKING CONFERENCE

DE KRAAL
KATHOLIEKE INDUSTRIELE HOGESCHOOL
OSTEND, BELGIUM
SEPTEMBER 6-8, 1988

EDITED BY
E.J.H. KERCKHOFFS
H. KOPPELAAR
R.C. VAN DE PERRE
G.C. VANSTEENKISTE

ORGANIZED BY
SCS INTERNATIONAL

IN COOPERATION WITH
ASIM, DBSS, ISCS, SIMS, UKSC, JSST AND CSSC

AND SPONSORED BY
DEUTECH INTERNATIONAL BV, SAR INVESTMENT PROPORTIES,
SIERENBERG & DE GANS BV, TECHNICAL MARKETING CONSULTING BV, TURBOSIM-IV

Printed in Belgium

Ship Manœuvring Simulation and Intelligent Marine Collision Avoidance

R.S. Burns, G.K. Blackwell, C.T. Stockel
Ship Control Group, Plymouth Polytechnic
Plymouth, Devon, United Kingdom

ABSTRACT

The paper describes two aspects of work currently undertaken by the Ship Control Group at Plymouth Polytechnic. The first is the use of multivariable system theory to construct a mathematical model that accurately describes, in real time, the response of both large and small ships to control inputs (demanded rudder and engines) and disturbance inputs (wind and tide). The model has three degrees of freedom, surge, sway and yaw and is particularly effective when simulating tight manoeuvres, as may be experienced in the approaches to a port.

The second aspect concerns the development of an expert system for collision avoidance at sea. The system, which ultimately will be housed in a micro-computer situated on the bridge of a ship, is now at the 'second' prototype' stage. It incorporates an expert system module and a simulator module, each operating independently on two knowledge bases, one for own-ship and one for hazard vessel. Current development is concerned with provision of an object-oriented 'front-end' module, responsible for decision-making and interfacing with the user.

SIMULATION OF SHIP MANOEUVRES

The ability to simulate, in real time, the manoeuvring characteristics of a ship is very useful, and can be used for a number of purposes. The ship designer can quickly evaluate effects such as hull geometry, rudder area, propeller thrust for example on the handling characteristics. The control engineer can use the model for autopilot analysis. Harbour design may be facilitated using multi-ship modelling systems. Realistic marine simulators can be constructed for the training of bridge personnel. Real-time models can be used on-board ship as an aid to navigation, in a Kalman filter for example. The discrete mathematical model described in this paper is both non-linear and time-varying, and has been successfully employed to simulate a range of ship types, including, fast cargo ships, car ferries, survey vessels and training vessels ranging from 8 tonnes to 17,000 tonnes gross displacement.

SHIP MATHEMATICAL MODEL

All moving rigid bodies contain six degrees of freedom. In this paper the ship is considered to be a rigid body with three degrees of freedom, namely surge, sway and yaw. Ship motions in heave, pitch and roll are considered small enough to be neglected.

From Newtons laws of motion the equations for surge, sway and yaw can be written as:

$$m\dot{u} - mrv = X \qquad (1)$$

$$m\dot{v} + mur = Y \qquad (2)$$

$$I_z\dot{r} = N \qquad (3)$$

The usual method employed to obtain the X and Y forces and yaw moment (1) is to consider the forces and moments acting on the vessel as functions of:

(i) Properties of the ship e.g. length and hull geometry.

(ii) Properties of motion e.g. velocity

(iii) Properties of the fluid e.g. density of sea water

This function can be reduced to useful mathematical form by the use of Taylor's expansion for a function of several variables. It has been shown (2) that taking only the linear terms from the expansion is insufficient to define the ship accurately. Non-linear terms that give rise to more than 10% of the global force or moment were considered of major importance, and included in the equation set (3).

The state variables chosen to represent the vessel are rudder angle, engine revolutions, forward displacement, forward velocity, lateral displacement, lateral velocity, heading and yaw rate, so the state vector is define as:

$$x^T = (\delta_A, n_A, x, u, y, v, \psi, r) \qquad (4)$$

This state is affected by the forcing vector

$$u^T = (\delta_D, n_D, u_c, v_c, u_a, v_a) \qquad (5)$$

From these eight states a set of first-order differential equations can be used to define the ship.

$$\dot{x}(t) = F(t)x(t) + G(t)u(t) \qquad (6)$$

It is convenient to partition G matrix into the control forcing functions G_c and the disturbance forcing functions G_D.

$$\dot{x}(t) = F(t)x(t) + G_c(t)u(t) + G_D(t)w(t) \qquad (7)$$

The corresponding discrete time solution is

$$x(k+1) = A(k,k+1)x(k) + B(k,k+1)u(k) + C(k,k+1)w(k) \qquad (8)$$

In the difference equation (8), if the state vector x(k), the control vector u(k) and the disturbance vector w(k)are known at the present time (kT), then the new state vector x(k+1) may be predicted at some future time (k+1)T.

Item	Description
Length between perpendiculars	160.9 m
Beam	23.17 m
Draught	9.07 m
Displacement	17,102 tonnes
Block coefficient	0.6
Number of propellers	one
Forward speed at 75 rev/min (prop. speed)	15.0 knots

Table 1 Particulars of Fast Cargo Ship

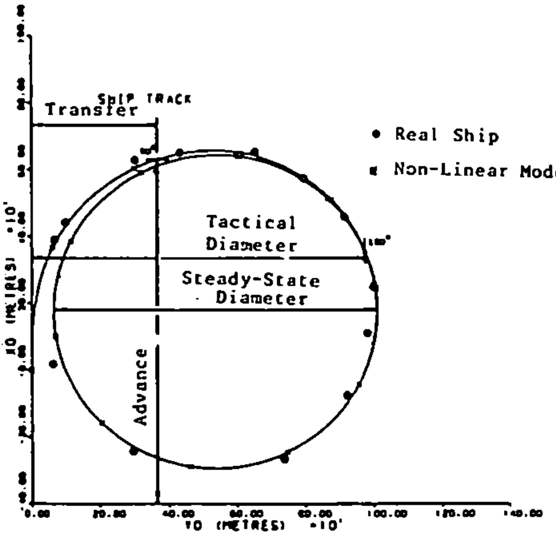

FIGURE 1

Turning Circle - Fast Cargo Ship

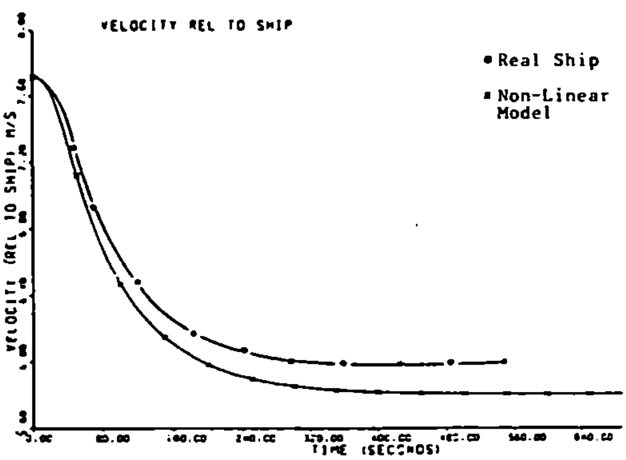

FIGURE 2

Reduction in Forward Speed - Fast Cargo Ship

Figure 1 shows a turning circle simulation for the vessel described in Table 1, with an approach speed of 15 knots and 20 degrees starboard rudder. Figure 2 compares simulated and actual reduction in forward speed during the turn. Figure 3 illustrates the vessel response in the surge direction to a step change in demanded engine speed.

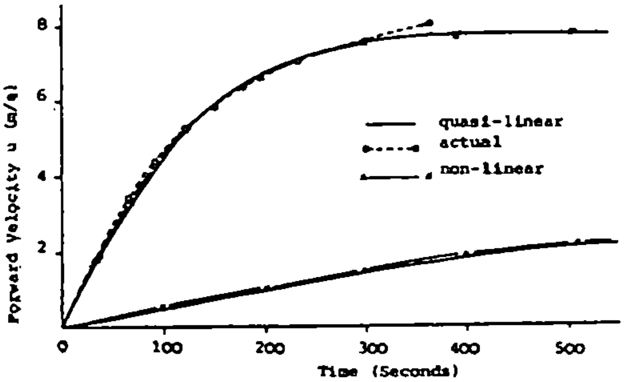

FIGURE 3

Step Change in Demanded Engine Speed - Fast Cargo Ship

SIMULATION PROGRAM

Input data to the program includes the usual geometric and mass properties, such as overall length, gross displacement and moment of inertia about the z - axis. In addition, 45 non-dimensional linear and non-linear hydrodynamic coefficients for the vessel are required. These can be provided by the Ship Control Group for particular ships under a consultancy arrangement.

A sampling period must be specified, usually between one and five seconds depending upon the size of the vessel, together with a set of initial conditions (starting position, heading, forward speed for example).

When the program is running the hydrodynamic coefficients are re-dimensionalised during each sampling period thus updating the mathematical model. This ensures that during a particular manoeuvre the instantaneous effects of yaw-rate, forward and lateral velocity are taken into account during the dimensionalising process.

This constant re-evaluation of the mathematical model means that a high degree of accuracy is maintained during manoeuvring simulations.

INSTRUMENT TRAINING VESSEL

Item	Description
Length Overall	27.43 m
Beam	5.49 m
Draught	2.09 m
Displacement	116 tonnes
Block Coefficient	0.48
Number of Propellers	2
Forward speed at 1200 rev/min (engine)	11 knots

Table 2 Particulars of Instrument Training Vessel

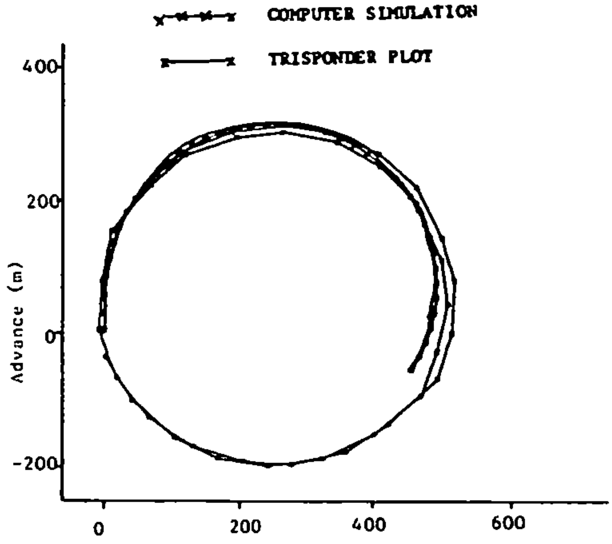

FIGURE 4

Turning Circle - Instrument Training Vessel

Figure 4 is a computer simulation of a turning circle (15 degrees starboard rudder, 11 knot approach speed) for the instrument training vessel superimposed upon the turning circle measured by the Ship Control Group using trisponder. This is a microwave position fixing system which measures range from the on-board transmitter to two or three shore-based transmitters (remotes). The system accuracy is within three metres and is widely used in hydrographic survey work.

Unlike the previous vessel, where the hydrodynamic coefficients were measured in a towing tank environment, coefficients for the instrument training vessel were evaluated theoretically, and fine-tuned by comparing simulation results with trisponder results.

In this application, the mathematical model was used in real-time by an on-board computer to accurately position the vessel on an electronic chart. Measurements of position, heading and speed were compared with predictions from the model, and a best estimate arrived at using Kalman filter techniques.

CATAMARAN

Item	Description
Number of hulls	2
Length overall	11.17 m
Width overall	4.3 m
Hull separation	3.7 m
Displacement	8.5 tonnes
Number of propellers	2
Forward speed at 2200 rev/min (engine)	8.9 knots

Table 3 Particulars of Catamaran

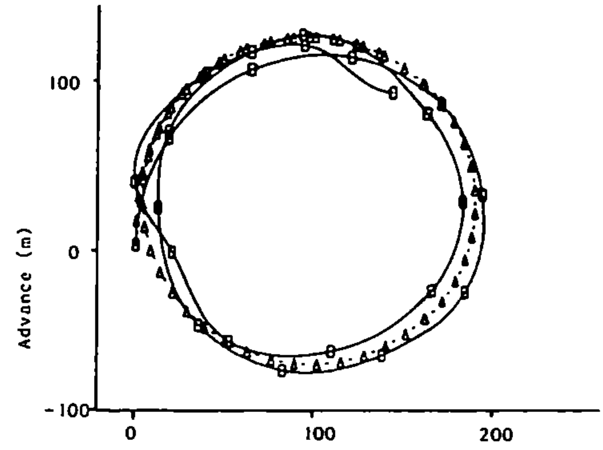

FIGURE 5

Turning Circle - Catamaran

The catamaran, with its twin hulls, was more difficult to model than mono-hull vessels. Figure 5 shows simulated and measured (trisponder) turning circles for 5 degree starboard rudder, approach speed 9 knots.

HARDWARE

The original simulation package was developed on a PRIME 9950 mainframe using FORTRAN 77. More recently the routines have been translated into C and can run on a variety of computers including IBM PC's and compatible machines.

THE COLLISION AVOIDANCE TASK

The task of collision avoidance at sea is generally recognised to be a rule-based process, as evidenced by The International Regulations for Preventing Collisions at Sea [4], a document universally accepted and respected by responsible mariners. This document appears at first sight to give a definitive sequence of operations covering every form of encounter between two marine vessels, thus making the task manageable via a straight forward procedural approach. However, a detailed study of these regulations, and (even more telling) an observation of their interpretation by mariners, reveals a fair degree of latitude in their application. The regulations are in fact a set of clear guidelines, leaving the mariner to make decisions on detail such as timing, clearances, suitable course alterations. The wide variety of 'non-standard' situations which may occur are also left to interpretation by the experienced mariner, in the light of these regulations.

Such a situation is clearly susceptible to an expert systems approach. 'Rules of thumb' for specific types of situation may be supplemented by rules for variations in these situations. This rule structure would be founded on the previously-mentioned anti-collision regulations, would incorporate the accumulated wisdom of expert mariners, and would presumably also be tailored to reflect the response characteristics of the system operating that rule structure - speed of response, breadth of information available to the system, possible consequence of misjudgement (confidence limits).

An expert system as described would require access to two types of information, static and dynamic. Static information relates to fixed characteristics of the vessel - length, beam, maximum speed, minimum turning circle, safe clearing distance, and a variety of technical data (some possibly specific to current voyage). Dynamic information, to be constantly updated, would include such considerations as current speed and course, plus data on position, speed and course of any potential hazards in the vicinity. Such information must, by its nature, be input to the system directly via a range of sensors, including such instrumentation as radar.

Such an intelligent response system should be capable of evaluating an encounter between two or more vessels from the standpoint of an experienced mariner, and taking (or advising) appropriate action. On being provided with the information normally available on the bridge of a well-equipped ship, this system should be capable of:

(a) Recognising an encounter (potential hazard) situation in good time;

(b) Identifying the type of encounter and the status of one's own vessel (own-ship) in that encounter, according to the International Regulations for Preventing Collisions at Sea - such status would normally be either 'stand-on' (having right of way) or 'give-way', according to one's position relative to the other vessel (subject to types of vessel involved, fishing vessels and deep-draught vessels having right of way in many circumstances);

(c) Choosing a course of action which combines a sensible safety margin with the least practicable inconvenience (where own-ship is judged to be stand-on, this will - initially at least - mean no alteration of course);

(d) Maintaining a watching brief on the other vessel(s) in the encounter, and being ready to take avoiding action if necessary, should the situation change - it is not unknown for 'rogue' vessels to disregard the regulations and plough straight on when they should give way, or even to turn into danger.

Clearly implicit in this 'seaway code' is the understanding that any other vessel encountered will be assumed to be 'well-behaved' until proved otherwise, but that an identical perception of, and reaction to, the encounter by the master of another vessel cannot be taken for granted. It is also imperative always to adopt a give-way stance in relation to any hazard other than another vessel under control, for example a derelict or a rock!

CURRENT DEVELOPMENTS

Development is currently at the 'second prototype' stage. Figure 6 shows the nature of the display produced by this package. The package comprises a simulator module, driving each of two vessels (own-ship and hazard vessel) independently on the knowledge bases for each of the two vessels. At 20 second intervals, ship time (simulated by 1/2 second intervals in real time) the following sequence is carried out:

(a) Certain minimal information, as would be available via instrumentation (speed, course, relative position) is communicated from the knowledge base of each ship to the knowledge base of the other;

(b) The expert system module considers the current state of each vessel independently, setting status indicators for consequent action, to be carried out by the simulator module;

(c) The simulator module updates the position, speed, and course information for each vessel in turn, with due regard for the 'recommendations' of the expert system with respect to that vessel;

(d) The screen display is updated to show the new current situation.

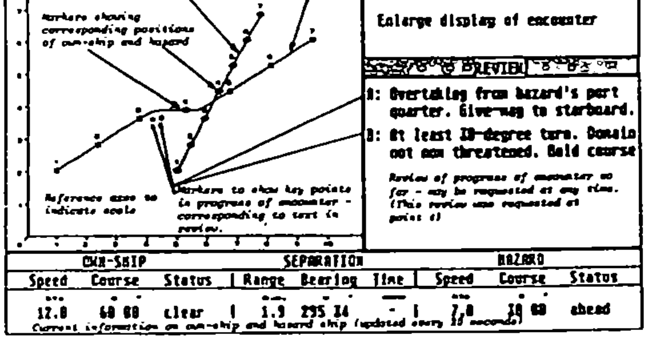

Figure 6 Annotated Sample Screen Display

The first prototype [5] contained the decision-making logic for the expert system module in such a form that the rules were 'hard-wired' into the program code. Whilst adequate for an initial test-bed, this approach was unsuited to a flexible, expanding rule base as envisaged at the outset of this paper. The second prototype incorporates a content-free 'inference engine' actioning a set of rules which form a 'decision network' - see figure 7. These rules are held as a set of nodes, or objects, each comprising a simple boolean decision function, pointers to two other nodes, and si other relevant parameters defining consequent action. The most significant of these are the two flags, one for each of the 'child nodes', indicating whether that node is to be actioned immediately or at the next time-step (solid and dotted lines respectively in figure 7). This form of structure allows for the inference/decision network to be expanded and enriched to any degree of sophistication. A simple extension will also permit 'backtracking' to explain the decision process at any stage, and parallel processing is a clear possibility without major changes.

There are several qualitative differences in this package, as compared to most other expert system packages currently in circulation. This system, as envisaged in its final form, exhibits the following advances on current 'standard' expert systems:

(a) mainly sensory input, supported by keyboard input from the user;

(b) real-time response;

(c) substantial (i.e. non-trivial) mathematical processing, mainly of multi-variable trigonometrical functions;

(d) Graphical display of current status.

For those reasons, no suitable expert systems shell has been found to fit the requirements. Indeed, one objective of this work is to provide a general-purpose 'collision avoidance' shell, with wider application than simply the marine environment.

The first two prototypes have been written completely in the 'C' programming language, and it is intended that this will also form the basis of the final system. The nodal network is currently implemented by a combination of structures and boolean functions in 'C'. Object-oriented languages (Smalltalk, Objective C, C++) are being considered for 'front-end' processing - Ie user interface and top-level decision-making. Such an approach is likely to be more flexible and 'user-friendly' than the current C-coded GEM environment (on the Atari ST microcomputer).

PRINCIPLES OF OPERATION

The operation of this package centres on two underlying principles identified in earlier research. The first is that of the domain [6], [7], an area around a vessel which the master would wish to keep free of any other ship or potential hazard. The second is that of the arena [7], the wider area around one's vessel within which any potential hazard should be carefully considered for degree of threat and appropriate remedial action. This latter concept has been refined from a distance-based to a time-based criterion, RDRR (Range-to-Domain/Range-Rate) [8], i.e. projected time to domain infringement. It is in this latter, time-vectored form that this principle is applied in the expert systems package; the time in question is shown in the line of data at the bottom of the screen display.

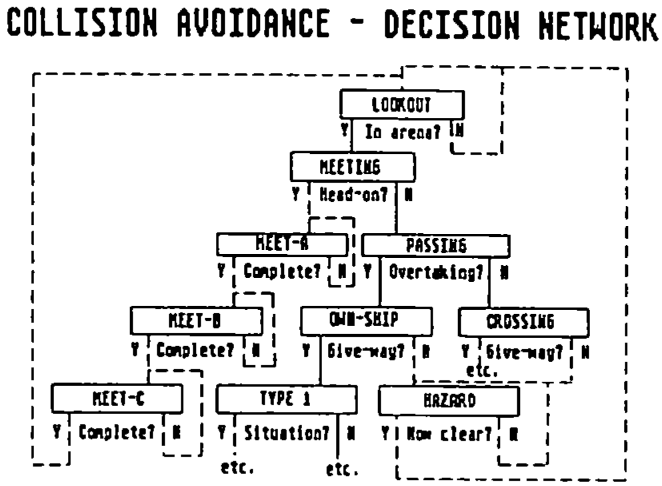

Figure 7

One important aspect of collision avoidance manoeuvres is that they must be clearly visible and identifiable as such from another vessel - no other form of communication between vessels can be assumed. Consequently, any manoeuvre undertaken at the direction of the expert system must be of sufficient magnitude (a 'clear turn') and in good time to be seen from another vessel; such considerations must be set against the temptation to use the power of the computer to 'optimise' course changes purely on the basis of 'safe clearance' criteria. The current model uses a minimum turn of 30 degrees, and a minimum RDRR criterion as defined as a parameter for each vessel typically 10 minutes, but varied to suit the handling characteristics of the vessel being simulated.

The rules as implemented embody the International Regulations for Preventing Collisions at Sea, plus clarification of specific practical details as observed by Colley [9] and used most effectively in a previous simulation. Using these rules, every collision avoidance manoeuvre is taken through four stages, as follows:

(1) When the RDRR criterion is infringed, the give-way vessel initiates a turning action;

(2) Once the required turn is effected (minimum 30 degrees), the vessel ceases its turn;

(3) When the original heading may be resumed without further risk of domain infringement (and any other conditions are satisfied), alter-back is initiated, necessitating a turn in the opposite direction;

(4) When back on the original heading, the manoeuvre is terminated, the vessel proceeds on that heading.

It should be noted that the function of the inference engine at each time-step is to take the basic parameters (speed, heading, position) provided by the simulator. From these it generates a standard set of further parameters - relative bearing, relative velocity, angle between relative bearing and relative heading, etc - and applies the rules to them to ascertain which situation pertains to the next step.

EFFICACY OF INITIAL RULE BASE

The effectiveness of the present rule structure has been evaluated by a multiple run of 500 simulated encounters. The results of this exercise are shown in the bar chart in Fig. 8.

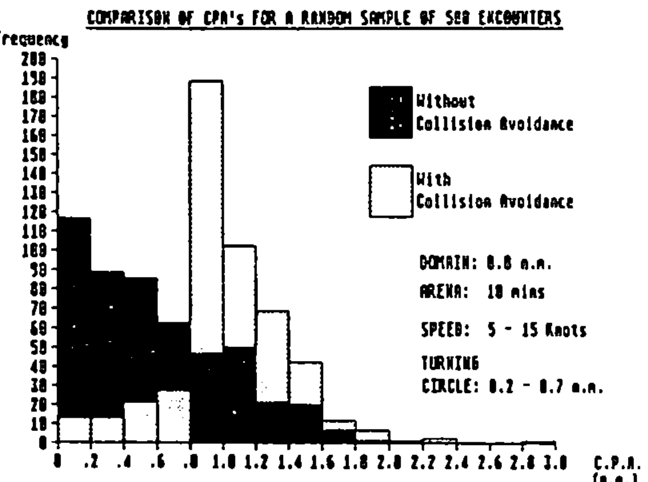

Figure 8

The results of this multiple simulation illustrate that, in the large majority of cases, potential domain violations were avoided by the expert system invoking appropriate collision avoidance strategies; moreover, the manoeuvres involved did not, in the main, involve excessive course alterations - CPA separations for these manoeuvres are clearly bunched just outside the domain boundary. However, a small number of the encounters still show CPAs within the domain boundary - this is clearly unacceptable. The reasons for this are considered to be threefold:

(1) The rule base in the initial expert system module was shown to have certain 'blind spots', in which each vessel assessed the other as the give-way vessel, and the required avoidance action was not taken;

(2) In using a fairly simple model for generation of simulated vessels, some realism was lost - vessels with unlikely combinations of characteristics were involved in a number of the offending encounters;

(3) A small number of these simulations highlighted the need for an increased level of sophistication in the strategy for recognising and handling encounter situations. Specifically, the fixed look-ahead, used to assess the time to initiate avoidance manoeuvres, appeared inadequate in certain cases; an ongoing look-ahead, extrapolating at regular intervals from current status, seemed indicated to ensure action at the optimum time. This is considered to be the most important development to arise from the multiple simulation exercise, and will form a central feature of future developments.

NOMENCLATURE

A) Matrices and Vectors

A	Discrete State Transition Matrix
B	Discrete Control Matrix
C	Discrete Disturbance Matrix
F	Continuous Time System Matrix
$G, G_C G_D$	Continouous Time Forcing Matrices
u	Control Vector
w	Disturbance Vector
x	Ship Related State Vector

B) Scalar Symbols

I_z	Ship Moment of Inertia about z axis
k	Integer Counter
m	Ship Mass
$n_A n_D$	Actual and Demanded Engine Speeds
r	Yaw rate
t	Continuous Time
u, u_c, u_a	Forward velocity of Ship, Current and Air
v, v_c, v_a	Lateral Velocity of Ship, Current and Air
X,Y,Z	Total Forces and Moments in Surge, Sway and Yaw.
δ_A, δ_D	Actual and Demanded Rudder Angles
ψ	Ship Heading

REFERENCES

1. ABKOWITZ, M. 'Lectures on Ship Hydrodynamics, Steering and Manoeuvrability.' Hy-A Report, Hy, 5., Denmark, 1964

2. BURNS, R. S. 'The Automatic Control of Large Ships in Confined Waters.' Ph.D. Thesis, Plymouth Polytechnic, September 1984.

3. BURNS, R. S., DOVE, M. J., BOUNCER, T. H., STOCKEL, C. T. 'A Discrete, Time Varying, Non-linear Mathematical Model for the Simulation of Ship Manoeuvres.' First International Maritime Simulation Symposium, Munich, June 1985.

4. International Marine Organisation, 1981. "International Regulations for Preventing Collisions at Sea".

5. BLACKWELL, G. K., BURNS, R. S., and STOCKEL, C. T., 1987. 'An Expert Systems Approach to Collision Avoidance.' Proceedings of Summer Computer Simulation Conference, Montreal. Society for Computer Simulation, July 1987, pp43-48.

6. GOODWIN, E. M., 1975. 'A statistical Study of Ship Domains.' Journal of Navigation. Volume 28, pp328-344.

7. DAVIS, P. V., DOVE, M. J. and STOCKEL, C. T., 1983 'A Computer Simulation of Marine Traffic Using Domains and Arenas.' Journal of Navigation. Volume 33, pp215-222.

8. COLLEY, B. A., CURTIS, R. G., and STOCKEL, C. T., 1983. 'Manoeuvring Times, Domains and Arenas.' Journal of Navigation. Volume 36, pp324-328.

9. COLLEY, B. A., 1985. 'Computer Simulation of Marine Traffic Systems.' PhD Thesis.

Maritime Communications
and Control
26 - 28 October 1988

Paper 19

An Intelligent Knowledge Based System for Marine Collision Avoidance

G K Blackwell, BTech, PGCE, B A Colley, PhD and C T Stockel, PhD
Plymouth Polytechnic

SYNOPSIS

A prototype expert system for collision avoidance has been developed, incorporating the International Rules for Preventing Collisions at Sea. This system simulated an encounter between two vessels whose control/response characteristics were read into the system. A built-in rule structure evaluated the situation independently for each vessel at each time-step, advising the simulator module, which then calculated the relevant parameters and updated the knowledge base for each vessel; each vessel had access only to such information about the other as would be available through ships' instrumentation. A multi-window screen display was used to show: tracks of the two vessels; current speed, direction and status information; options available; and a review of progress of the encounter, if requested, indicating reasons for decisions made.

A theoretical evaluation, based on a sample of 500 randomly-generated encounters, has shown this prototype to be highly efficient at handling most encounter situations; it has also identified areas for improvement in the decision making logic.

A second prototype, incorporating various changes, differs in one major respect from the first: the rule base has been separated out from the program logic, leaving a content-free inference engine. Each rule is now defined explicitly in terms of description of current (new) status, antecedents and consequent action (generally a sequence, involving other rules). Planned extensions to the user interface will allow modification and extension of the rule base by competent persons.

The expert system has been written in the 'C' programming language, on a microcomputer system capable of operating in real-time on board ship. The prime objective of this work is to provide an expert advisor for mariners, capable of acting on sensory input from all available channels, including information provided by an operator via a keyboard, and advising on the preferred course of action.

INTRODUCTION

A rule-based approach is an essential prerequisite to the task of collision avoidance at sea, whether or not computers are involved in the process. At first sight, it might seem that a simple procedural application of a standard set of rules, with consequent actions, would cover all requirements; indeed, all mariners are expected to abide by just such a set of rules.

Such a procedural approach has been used most effectively in computer simulation of marine traffic flow and collision avoidance.[1] The effects of changes (eg in traffic volume) in high-density traffic lanes have been highlighted by such a model, and the technique has also been used to good effect in training mariners. However, this approach takes no account of the need for experience and common sense in applying these rules and, as such, is not a practical option in real-world situations.

The International Regulations for Preventing Collisions at Sea,[2] although quite specific as far as they go, do not in themselves give rigorous definitions of preconditions and consequent actions. It is left to the mariner to decide such details

Grahame Blackwell gained a BTech degree in maths, statistics and computing, and a PGCE from Brunel University. He has since been engaged in a variety of computer applications, notably in the scientific and real-time fields. Having later moved into education, and played an active part in the development of educational computing, he is now a Senior Lecturer in computing at Plymouth Polytechnic. He is currently registered for a part-time PhD, the subject of which is 'An expert systems approach to collision avoidance'.

Dr B A Colley gained a BSc degree in nautical studies followed by a PhD degree for work on the computer simulation of marine traffic systems in 1985. Since then he has worked as a software engineer for Plessey Naval Systems and is now with British Aerospace.

Dr C T Stockel gained BSc and PhD degrees in physics from the University of London. He has since worked on the applications of computers to a variety of scientific and technological problems, being elected to Fellowship of the British Computer Society in 1985. He is at present Principal Lecturer in computing at Plymouth Polytechnic.

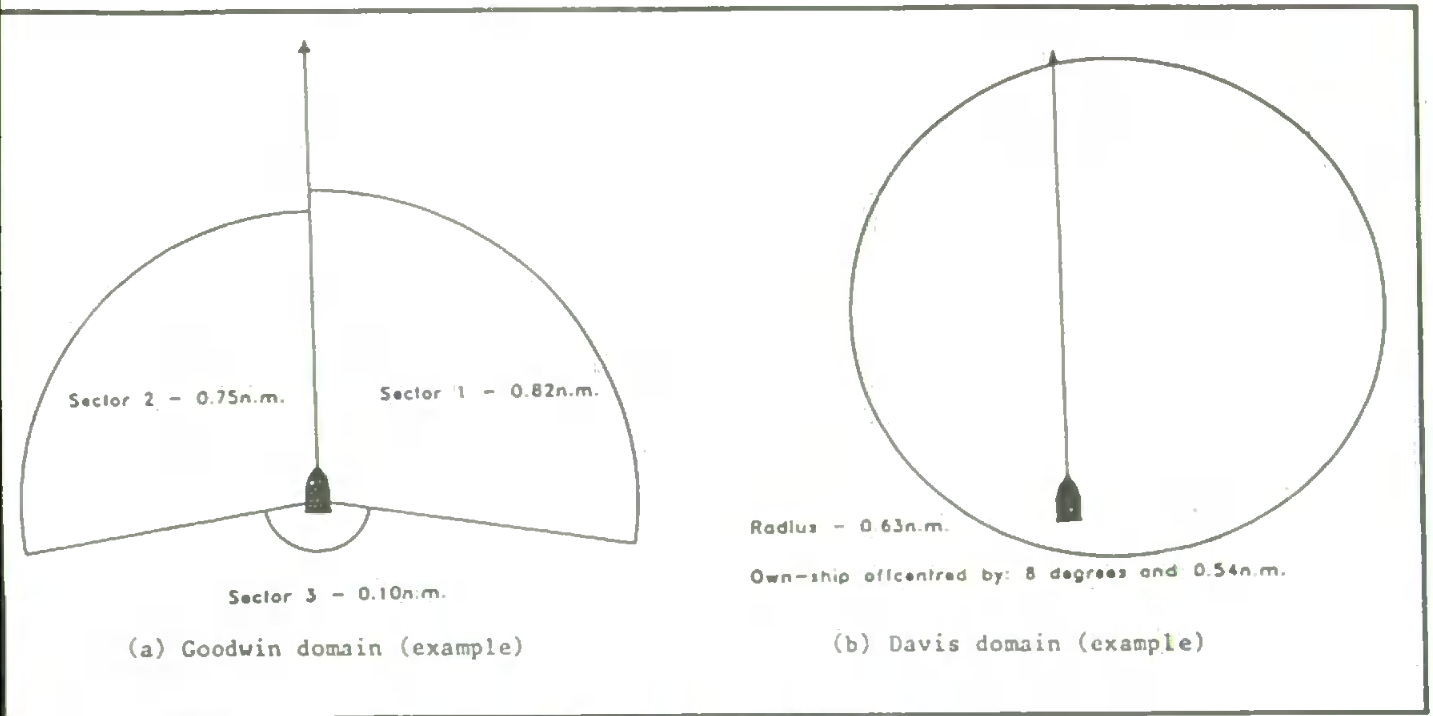

Fig. 3: Two examples of domain

2. Identifying the type of encounter and the status of one's own vessel (own-ship) in that encounter, according to the International Regulations for Preventing Collisions at Sea. Such status would normally be 'stand-on' (having right of way) or 'give-way', according to one's position relative to the other vessel (subject to the types of vessel involved, fishing vessels and deep-draught vessels having right of way in many circumstances).
3. Choosing a course of action which combines a sensible safety margin with the least practicable inconvenience (where own-ship is judged to be stand-on this will – initially at least – mean no alteration of course).
4. Maintaining a watching brief on the other vessel(s) in the encounter and being ready to take avoiding action, if necessary, should the situation change – it is not unknown for 'rogue' vessels to disregard the regulations and plough straight on when they should give way, or even to turn into danger.

A clear implication of such a code of conduct is the understanding that any other vessel encountered will be assumed to be similarly 'well-behaved' until and unless proved otherwise. It should be noted, however, that the masters of two vessels involved in an encounter will not always have the same perception of, or response to, that encounter. It is also imperative at all times to adopt a give-way stance in relation to any hazard other than another vessel under control, for example, a derelict or a rock!

CURRENT DEVELOPMENTS

An annotated example of the multi-window screen display generated by the prototype system is shown in Fig. 1. This display is substantially the same for both the current (second) prototype and the earlier version. This screen format is produced under the GEM environment, which allows the use of windows and icons for display, and a mouse for user input. Such facilities will be utilised more fully in later developments. The system is written completely in the 'C' programming language and, as yet, deals only with encounters between two manned vessels moving subject to fairly idealised parameters in open sea, ie no restricted channels or routing systems and no fixed obstacles.

The first prototype comprised an 'expert system' module, a simulator module and a knowledge base for each of the two vessels, own-ship and hazard.[3] The package operated on time-steps of 20 seconds in real time, each simulated by approximately one second of computer time. At each step, the expert system module would assess the status of own-ship by reference to its knowledge base, and set indicators in that knowledge base for action to be taken. The knowledge base for the hazard vessel would similarly be examined and flagged for subsequent action, independently, by the same expert system module. The dynamic information for each vessel in turn would then be updated by the simulator module, referencing the relevant parameters and also actioning any status flags set by the expert system in each case. As with a real situation, certain information on each vessel would be available to the master of the other ship via radar and other instrumentation, so such information would be interchanged between the two knowledge bases by the simulator module. The essential requirement of the expert system was to ensure that no hazard violated an area of proximity around the vessel, referred to as the 'domain' – the domain, and other associated concepts, are considered in detail later in this paper (Principles of Operation).

The effectiveness of this first prototype has been evaluated by a multiple run of 500 simulated encounters, each generated as follows: two vessels, each with a random course, speed and turning circle (the latter two parameters being within specified bounds) were placed in close proximity to the extent that each was very likely to consider the other a hazard, according to their terms of reference. The two vessels, one designated own-ship, the other hazard-ship, were then effectively taken back ten minutes in time along their tracks, and the simulation was allowed to proceed from that point – if an encounter situation already existed by that time, the time-reversal was extended to ffective area around a ship which a navigator would like to eep clear with respect to other ships and stationary objects'.[4] The original concept of three distinct circle sectors of differing adii for sidelights and sternlight was later modified to give a ontinuous circular domain boundary, with the ship off-centre n the circle so as to preserve areas subtended by sidelights and tern light.[5] This rationalisation was based on the logical ecessity for continuity of domain boundary (see Fig. 3). The irst objective of the current expert system, and its projected levelopments, is the protection of the integrity of the domain. n the prototype, this takes the simplified form of a circle entred at own-ship. It is likely that future developments will ncorporate the 3-sector domain, with 'fuzzy' (ie probability-veighted) boundaries between the sectors.

Based on observations of mariners manoeuvring to prevent nfringement of their domain by other vessels in the vicinity, Davis introduced a wider manoeuvring area, referred to as the arena.[5] Colley refined this concept to a time factor, considering he time-to-go (before anticipated domain infringement) at which the mariner will assess the status of his vessel and take appropriate action.[6] This time factor was labelled the Range-to-Domain/Range-Rate (RDRR). For example, on the basis of a ship's handling characteristics a mariner may choose to take action 15 minutes before potential domain infringement, ie a 15-minute RDRR. Konyn, in proposing a modified RDRR criterion (RDRR+) makes allowance for reaction times required to recognise and respond to different types of encounter.[7]

The computer model uses the RDRR criterion as defined by Colley to determine decision-time. Plans to switch to the RDRR+ criterion have been superseded by the perceived need for an ongoing look-ahead, a repeated forward-looking simulation (within the simulation provided by the model) which should provide a variable RDRR value to suit the specific situation. This additional burden on processing time serves to emphasise the need for fast mathematical processing to ensure real-time operation.

An essential requirement of any collision avoidance manoeuvre is that it must be clearly identifiable as such from another vessel. Although the computer could theoretically be used to 'optimise' (ie minimise) course changes purely on the basis of 'safe clearance' criteria, this would not, in general, meet the above requirement. Any manoeuvre undertaken at the direction of the expert system must be of sufficient magnitude ('a clear turn') and in good time to be seen from another vessel. The current model uses a minimum turn of 30°, and a minimum RDRR criterion as defined as a parameter for each vessel – typically 10 minutes, but varied to suit the handling characteristics of the vessel being simulated.

The rules, as implemented, embody the International Regulations for Preventing Collisions at Sea, plus clarification of specific practical details as observed by Colley and used most effectively in a previous simulation.[8] Using these rules, every collision avoidance manoeuvre is taken through four stages, as follows:

1. When the RDRR criterion is infringed, the give-way vessel initiates a turning action.
2. Once the required turn is effected (minimum 30°), the vessel ceases its turn.
3. When the original heading may be resumed without further risk of domain infringement (and any other conditions are satisfied), alter-back is initiated, necessitating a turn in the opposite direction.
4. When back on the original heading, the manoeuvre is terminated and the vessel proceeds on that heading.

It should be noted that the function of the inference engine at each time-step is to take the basic parameters (speed, heading, position) provided by the simulator. From these it generates a standard set of further parameters – relative bearing, relative velocity, angle between relative bearing and relative heading etc – and applies the rules to them to ascertain which situation pertains to the next step.

In the extended look-ahead mode, under development for use with a variable RDRR criterion, the system will, at regular intervals, extrapolate from current data through the anticipated manoeuvre. Hence, the optimum time-to-go for initiation of the appropriate collision avoidance strategy is identified.

EXAMPLES

Figure 1 shows own-ship overtaking hazard from hazard's port quarter.

Figure 4 shows a crossing encounter with own-ship giving-way; a full review of the encounter has been requested.

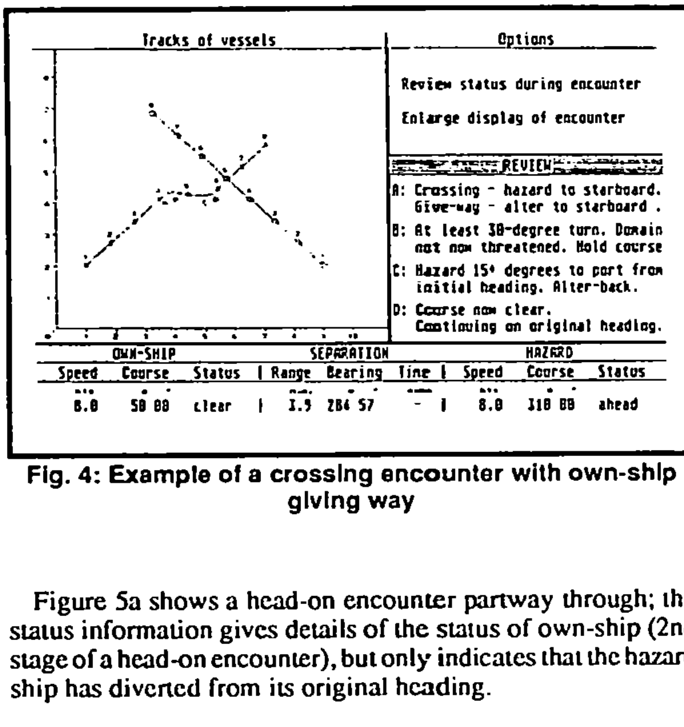

Fig. 4: Example of a crossing encounter with own-ship giving way

Figure 5a shows a head-on encounter partway through; the status information gives details of the status of own-ship (2nd stage of a head-on encounter), but only indicates that the hazard ship has diverted from its original heading.

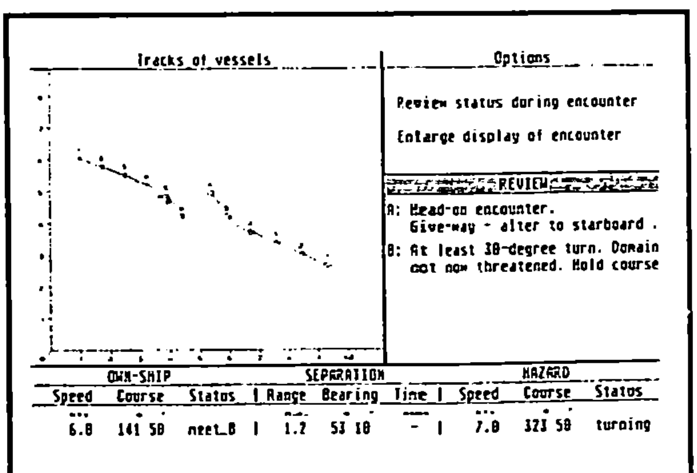

Fig. 5a: Example of a head-on encounter, partway through

Table I: First part of data logging file for the run shown in Fig. 6, including initial data input

```
data for own-ship                                   data for hazard ship

Course:-    (i)   Degrees:        47.00             Course:-   (i)   Degrees:           42.00
           (ii)   Minutes:         0.00                       (ii)   Minutes:            0.00
          Speed in knots:          4.00                       Speed in knots:           15.00
  Turning circle (n.mls):          0.60               Turning circle (n.mls):            0.40
         Position:-    x:         4.00              Range in nautical miles:             3.00
         Position:-    y:         4.00              Bearing:-  (i)   Degrees:          225.00
     Domain radius in mls:         0.80                        (ii)   Minutes:            0.00
Arena (RD/RR) in minutes:         10.00
```

OWN-SHIP X	Y	COURSE	SPEED	TIME TO GO	HAZARD X	Y	COURSE	SPEED	RANGE
4.00	4.00	47.00	4.00	12.14	1.88	1.88	42.00	15.00	3.00
4.02	4.02	47.00	4.00	11.81	1.93	1.94	42.00	15.00	2.94
4.03	4.03	47.00	4.00	11.47	1.99	2.00	42.00	15.00	2.88
4.05	4.05	47.00	4.00	11.14	2.05	2.06	42.00	15.00	2.82
4.07	4.06	47.00	4.00	10.81	2.10	2.13	42.00	15.00	2.76
4.08	4.08	47.00	4.00	10.47	2.16	2.19	42.00	15.00	2.70
4.10	4.09	47.00	4.00	10.14	2.21	2.25	42.00	15.00	2.63
NEW STATUS: own-ship: standon									
NEW STATUS: hazard ship: pass0_A									
4.11	4.11	47.00	4.00	9.81	2.27	2.31	42.00	15.00	2.57
4.13	4.12	47.00	4.00	9.88	2.33	2.37	53.94	15.00	2.51
4.15	4.14	47.00	4.00	-60.00	2.40	2.41	65.87	15.00	2.45
4.16	4.15	47.00	4.00	-60.00	2.48	2.44	77.81	15.00	2.40
4.18	4.17	47.00	4.00	-60.00	2.56	2.45	77.81	15.00	2.36
4.20	4.18	47.00	4.00	-60.00	2.64	2.47	77.81	15.00	2.31
4.21	4.20	47.00	4.00	-60.00	2.73	2.49	77.81	15.00	2.26

via RS232 communication links which are currently being set up. This simulator allows up to 4 bridge teams each to control a simulated vessel. Experienced mariners, each at the 'helm' of a simulated vessel, will be able to judge the expert system's handling of various encounter situations, viewing the system effectively as another helmsman.

3. The system is to be installed in the Polytechnic's own research vessel for comprehensive sea trials. It may also be possible to install the system in other vessels, to give a wider range of trial situations.

CONCLUSIONS

On the basis of developments so far, the feasibility of the project seems assured from the computing standpoint. The anti-collision model has shown the viability of the proposed product; the rule base has been enhanced on the basis of tests of the system; and clear pointers have been given regarding directions for future development.

This paper has not given consideration to the nature and quality of data provided by sensory inputs, such as radar. Any inferences drawn by the system will, of course, be totally dependent on such data in terms of reliability. Confidence limits will have to be set on such inferences, having first ensured optimum usage of such sensory data as is available. This could well involve cost factors, since the level of sophistication of instrumentation used will be largely determined by target price. Information not thus directly available to the system may still, of course, be input by the operator via the keyboard.

REFERENCES

1. B A Colley, R S Curtis, and C T Stockel, 'On Marine Traffic Flow and Collision Avoidance Computer Simulation', Journal of Navigation, Vol. 37, pp 232-250 (1984).
2. International Maritime Organisation, 'International Regulations for Preventing Collisions at Sea' (1981).
3. G K Blackwell, R S Burns and C T Stockel, 'An Expert Systems Approach to Collision Avoidance', Proceedings of Summer Computer Simulation Conference, Montreal, Society for Computer Simulation, pp 43-48 (July 1987).
4. E M Goodwin, 'A Statistical Study of Ship Domains', Journal of Navigation, Vol. 28, pp 328-344 (1975).
5. P V Davis, M J Dove, and C T Stockel, 'A Computer Simulation of Marine Traffic Using Domains and Arenas', Journal of Navigation, Vol. 33, pp 215-222 (1980).
6. B A Colley, R G Curtis, and C T Stockel, 'Manoeuvring Times, Domains and Arenas', Journal of Navigation, Vol. 36, pp 324-328 (1983).
7. M Konyn, 'Collision Avoidance at Sea and on a Marine Radar Simulator using Automatic Encounter Detection Techniques', PhD Thesis, London (1987).
8. B A Colley, 'Computer Simulation of Marine Traffic Systems', PhD Thesis (1985).

nulaton VI
'1989 By The Society for Computer
nulation International
3N 0-911801-51-0

Florida, March 89

Simulation of a marine guidance and expert collision avoidance system

R. S. Burns, G. K. Blackwell and K. M. Miller
Ship Control Group, Plymouth Polytechnic,
Plymouth, Devon, U.K.

BSTRACT

One of the major problems that stands in the way of the fully automated ship is that of safe guidance in and out of port. This paper addresses the problem of automatic guidance and collision avoidance, and draws upon the work undertaken over a number of years by the Ship Control Group at Plymouth Polytechnic, U.K., in terms of simulation, prototype model testing and full-scale system implementation.

Techniques to model and simulate the dynamic behaviour of a surface vessel are presented, together with simulated and actual results of an optimal ship guidance system. The problem of collision avoidance in port approaches is tackled using an expert system. Simulaltion results from a prototype system are given. Finally, the paper considers the need for an expert controller to provide integrated automatic guidance and collision avoidance.

INTRODUCTION

Although overall standards of safety at sea are very good, the approaches to a port, where traffic density is intense, may be considered a high risk area. It has been shown for example (Coldwell 1981) in the Humber Seaway, where there are 100 traffic movements per day, that there is at least one collision or grounding per week. In addition, there is the human element to consider. It has been highlighted (Panel on Human Error 1976) that 85 percent of all marine collisions and groundings are due to human error.

This suggests that there is a need for automatic guidance and collision avoidiance systems for marine vehicles in confined waters. As electronic navigation aids become more sophisticated and on-board computer systems become more available, the concept of the fully automatic ship becomes a tangible reality, and it is predicted that above 50 percent of the world's shipping will be fully automated by the turn of the century.

The work presented in this paper is concerned with the simulation of a marine guidance and collision avoidance system. A ship is considered to be a multivariable system that may be controlled by formulating an optimal policy that seeks to maximise the return from the system for a minimum cost. The optimal guidance system therefore controls simultaneously the position, heading and speed of the vessel. In addition, a rule-based expert system is simulated to handle sensory data input, and to make collision avoidance decisions based upon the International Regulations for preventing Collisions at Sea (1981).

SHIP MATHEMATICAL MODEL

All moving rigid bodies contain six degrees of freedom. In this paper the ship is considered to be a rigid body with three degrees of freedom, namely surge, sway and yaw. Ship motions in heave, pitch and roll are considered small enough to be neglected. It is convenient to describe the motion in terms of a moving system of axes coincident with the mass centre of the hull, which gives rise to an Eulerian set of equations of motion.

It is necessary to obtain the hydrodynamic surge and sway forces together with the yaw moments acting on the hull. These may be considered as functions of:

(i) Properties of the ship e.g. length and hull geometry.

(ii) Properties of motion e.g. velocity.

(iii) Properties of the fluid e.g. density of sea water.

These may be reduced to a useful mathematical form by the use of Taylor's expansion for a function of several variables (Abkowitz 1964). It has been shown (Burns et al. 1985) that taking only the linear terms from the expansion is insufficient to define the ship accurately. Non-linear terms that give rise to more than 10% of the global force or moment were considered of major importance, and included in the equation set.

The state vector $x(t)$ chosen to represent the vessel contains the following state variables:

1. Actual rudder angle
2. Actual engine speed
3. Forward position
4. Forward velocity
5. Lateral position
6. Lateral velocity
7. Heading
8. Yaw rate

The forcing vector $u(t)$ contains the following variables:

a. Demanded rudder angle
b. Demanded engine speed
c. Component of current speed in surge direction
d. Component of current speed in sway direction
e. Component of wind speed in surge direction
f. Component of wind speed in sway direction

From these eight states a set of first-order differential equations can be used to define the ship.

$$\dot{x}(t) = F(t)x(t) + G(t)u(t) \quad (1)$$

It is convenient to partition G matrix into the control forcing functions G_c and the disturbance forcing functions G_D.

$$\dot{x}(t) = F(t)x(t) + G_c(t)u(t) + G_D(t)w(t) \quad (2)$$

The corresponding discrete time solution is:

$$x(k+1) = A(k,k+1)x(k) + B(k,k+1)u(k) + C(k,k+1)w(k) \quad (3)$$

Equation (3) may be used in a recursive manner to simulate the passage of a vessel when the control vector u(k) and the disturbance vector w(k) are known.

Figure 1 shows a turning circle simulation of a fast cargo ship of displacement 17,100 tonnes. In the simulation the speed of approach is 15 knots and 20 degrees of starboard rudder has been applied.

function or performance index. This is based upon the summation of the weighted errors over some time interval, say, the time to complete the pilotage phase of the voyage. In addition to minimizing the errors in the output parameters, the optimal controller attempts to minimize also the control effort, i.e. to keep to a minimum the rudder and engine activity. The cost function is usually stated in the following quadratic terms:

$$J = \int_{t_0}^{t_1} \{(x-r)^T Q(x-r)+u^T Ru\}\, dt \quad (4)$$

Q and R are usually diagonal matrices and the values of the individual elements reflect the importance of the parameter being controlled. For example, in a track changing manoeuvre, large course errors will be incurred. If the track-keeping elements in the Q matrix are weighted more heavily than the course-keeping terms, then the majority of the control effort will be expended in reducing the track error, at the expense of the course deviation.

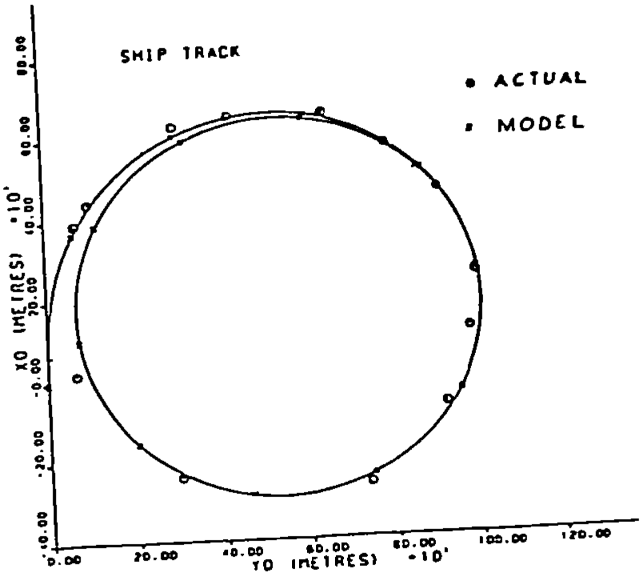

Figure 1

Turning Circle - Fast Cargo Ship

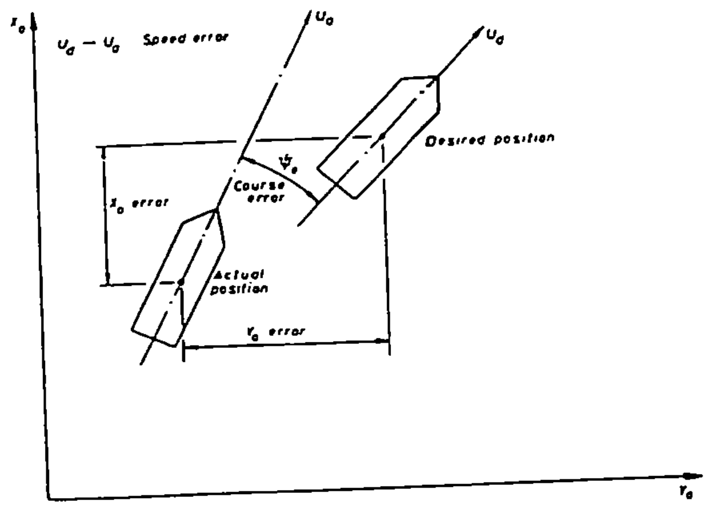

Figure 2

The Multivariable Control Problem

THE SHIP GUIDANCE PROBLEM

Multivariable Control

The guidance of a ship into the approaches to a port is a multivariable control problem.

The deviation from the desired position, course and speed, as shown in Figure 2 must be corrected for by operation of the rudder(s) and main engine(s). A feature of an optimal system is that it will seek to minimize a global parameter J, called the cost

Figure 3 shows a simulation of a full-size car ferry in the approaches to Plymouth. Here, the track-keeping is most heavily weighted, followed by the heading and forward speed. If the reduction in forward speed during a tight manoeuvre can be tolerated, the speed control loop may be dispensed with altogether.

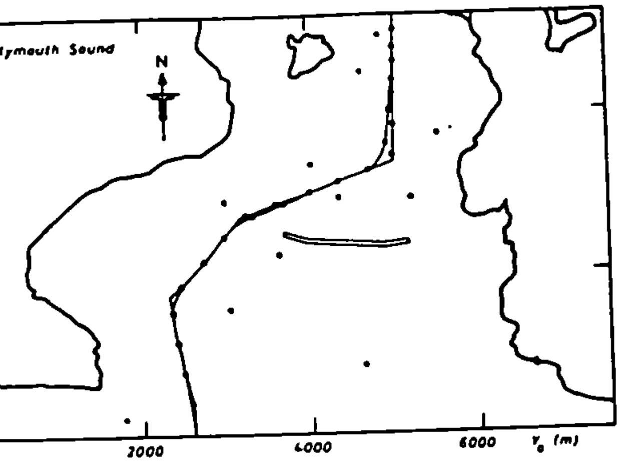

Figure 3

Computer Simulation - Approaches to Plymouth

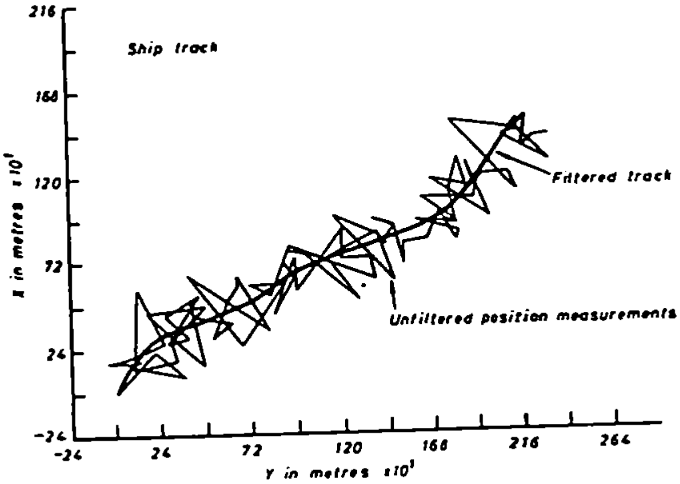

Figure 4

Filtered Track at Night

easurement and State Estimation

In Figure 3 it is assumed that all the state variables can be measured with complete accuracy. All navigational instruments contain measurement errors and a best estimate can be obtained by incorporating a minimum variance or Kalman-Bucy filter. These have been developed extensively for aerospace, and latterly marine navigation (Dove et al. 1985)

The Kalman filter is a recursive computational algorithm which works in a predictor-corrector manner. The current best estimate of the state vector x(k) is used to drive the mathematical model of the ship in real time to predict the state of the vessel at time (k+1). The predictions are compared with the measurements and multiplied by the Kalman gain matrix to obtain the best estimate x(k+1).

In determining the value of the gain matrix, consideration has to be given to measurement errors. These are assumed to be random with a Gaussian distribution, and are stated in terms of a co-variance matrix.

Figure 4 shows the simulation of a typical passage into the port under night-time conditions. The simulation assumes that positional data is being received from a Decca Navigator using a standard deviation of 200 m. It is seen that the true and filtered tracks are almost coincident.

THE COLLISION AVOIDANCE PROBLEM

A rule-based approach is an essential pre-requisite to the task of collision avoidance at sea, whether or not computers are involved in the process. At first sight, it might seem that a simple procedural application of a standard set of rules, with consequent actions, would cover all requirements; indeed, all mariners are expected to abide by just such a set of rules.

Such a procedural approach has been used most effectively in computer simulation of marine traffic flow and collision avoidance. (Colley et al. 1984) The effects of changes (e.g. in traffic volume) in high-density traffic lanes have been highlighted by such a model, and the technique has also been used to good effect in training mariners. However, this approach takes no account of the need for experience and common sense in aplying these rules and, as such, is not a practical option in real-world situations.

The International Regulations for Preventing Collisions at Sea, although quite specific as far as they go, do not in themselves give rigorous definitions of preconditions and consequent actions. It is left to the mariner to decide such details as timing, clearances, and suitable course alterations; such discretion comes with years of accumulated experience and wisdom. One must also consider `non-standard` situations, particularly those involving more than two vessels which could not all be adequately covered by any reasonable reference guide.

Clearly, it is possible to provide `rules of thumb` for specific types of situation, and to add supplementary rules for variations in these situations. Such a rule structure would be founded on the previously mentioned anti-collision regulations, would incorporate the accumulated wisdom of expert mariners and would, presumably, also be tailored to reflect the response characteristics of the system operating that rule structure - speed of response, breadth of information available to the system and possible consequence of misjudgement (confidence limits).

An expert system as described would require access to two types of information, static and dynamic. Static information relates to fixed characteristics of the vessel - length, beam, maximum speed, minimum turning circle, safe clearing distance, and a variety of technical data (some possibly specific to current voyage). Dynamic information, to be constantly updated, would include such considerations as current speed and course, plus data on position, speed and course of any potential hazards in the vicinity. Such information must, by its nature, be input to the system directly via a range of sensors, including such instrumentation as radar.

Such an intelligent response system should be capable of evaluating an encounter between two or more vessels from the standpoint of an experienced mariner, and taking (or advising) appropriate action. On being provided with the information normally available on the bridge of a well-equipped ship, this system should be capable of:

(a) Recognising an encounter (potential hazard) situation in good time;

(b) Identifying the type of encounter and the status of one's own vessel (own-ship) in that encounter, according to the International Regulations for Preventing Collisions at Sea - such status would normally be either `stand-on` (having right of way) or `give-way`, according to one's position relative to the other vessel (subject to types of vessel involved, fishing vessels and deep-draught vessels having right of way in many circumstances);

(c) Choosing a course of action which combines a sensible safety margin with the least practicable inconvenience (where own-ship is judged to be stand-on, this will - initially at least - mean no alteration of course);

(d) Maintaining a watching brief on the other vessel(s) in the encounter, and being ready to take avoiding action if necessary, should the situation change - it is not unknown for `rogue` vessels to disregard the regulations and plough straight on when they should give way, or even to turn into danger.

COLLISION AVOIDANCE SIMULATION

The simulation package in current use is at the `second prototype` stage. The package comprises a simulator module, driving each of two vessels (own-ship and hazard vessel) independently on the knowledge bases for each of the two vessels. At 20 second intervals, ship time (simulated by 1/2 second intervals in real time) the following sequence is carried out:

(a) Certain minimal information, as would be available via instrumentation (speed, course, relative position) is communicated from the knowledge base of each ship to the knowledge base of the other;

(b) The expert system module considers the current state of each vessel independently, setting status indicators for consequent action, to be carried out by the simulator module;

(c) The simulator module updates the position, speed, and course information for each vessel in turn, with due regard for the `recommendations` of the expert system with respect to that vessel;

(d) The screen display is updated to show the new current situation.

The first prototype (Blackwell et al. 1987) contained the decision-making logic for the expert system module in such a form that the rules were `hard-wired` into the program code. Whilst adequate for an initial test-bed, this approach was unsuited to a flexible, expanding rule base as envisaged at the outset of this paper. The second prototype incorporates a content-free `inference engine` actioning a set of rules which form a `decision network` - see Figure 5. These rules are held as a set of nodes, or objects, each comprising a simple boolean decision function, pointers to two other nodes, and six other relevant parameters defining consequent action. The most significant of these are the two flags, one for each of the `child nodes`, indicating whether that node is to be actioned immediately or at the next time-step (solid and dotted lines respectively in Figure 5). This form of structure allows for the inference/decision network to be expanded and enriched to any degree of sophistication. A simple extension will also permit `backtracking` to explain the decision process at any stage, and parallel processing is a clear possibility without major changes.

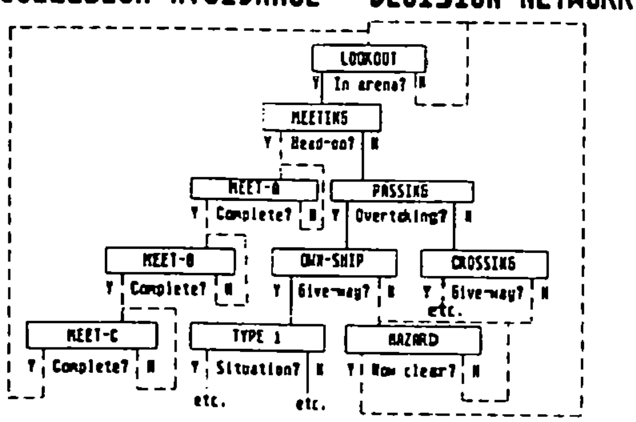

Figure 5

COLLISION AVOIDANCE - DECISION NETWORK

OPERATION OF EXPERT SYSTEM MODULES

As outlined in the previous section, the prototype package drives two simulated vessels, each with its own operational parameters, and each (independently) subject to the same expert system module. As a consequence, both vessels act in accordance with the anti-collision regulations, and there is no likelihood of one vessel having to take evasive action at a late stage to deal with non-co-operation by the other. In the near future, it is intended to model the two ships on separate computers, interacting via a communications link; an additional feature of such a system would be the provision of a manual override on one of the ships - the hazard - to enable testing of an extension of the rule base to handle `rogue` behaviour.

At the present phase of development, the rule base consists of rules for:

1. Identifying the presence of a potential hazard, and assessing the degree of threat in terms of expected time to infringement of the domain - this time factor is central to the decision on when to take avoiding action, if necessary.

2. Identifying the type of encounter, including several variants in some cases (the primary types of encounter being head-on, crossing and overtaking), and fixing the status of own-ship and the perceived status of hazard-ship in the encounter.

3. Negotiating the stages of the encounter, with due regard for appropriate safety margins.

SIMULATION RESULTS

Figure 6 shows a completed head-on encounter. Note that the status of own-ship is known in detail, but hazard ship is known only to be changing course.

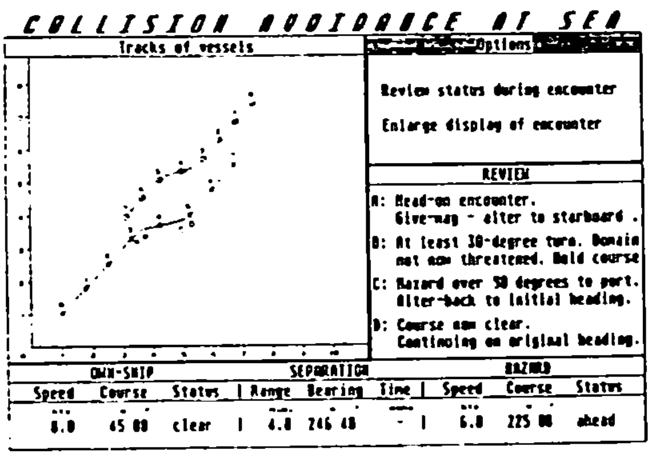

Figure 6

Example of a Completed Head-on Encounter

Figure 7 shows an overtaking manoeuvre when own-ship is to starboard of hazard vessel. Here own ship turns starboard on the same heading, but parallel track of the hazard ship, and then turns to port on original course, passing ahead of hazard vessel.

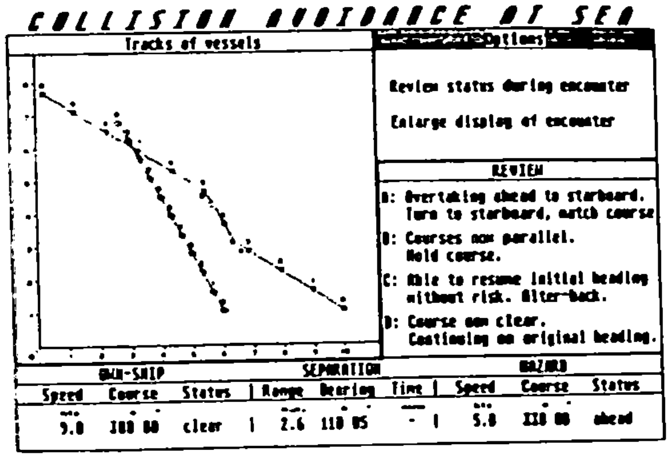

Figure 7

Own Ship being Overtaken by Hazard Vessel

ANALYSIS OF SIMULATION RESULTS

The effectiveness of the present rule structure has been evaluated by a multiple run of 500 simulated encounters. The results of this exercise are shown in the bar chart in Figure 8.

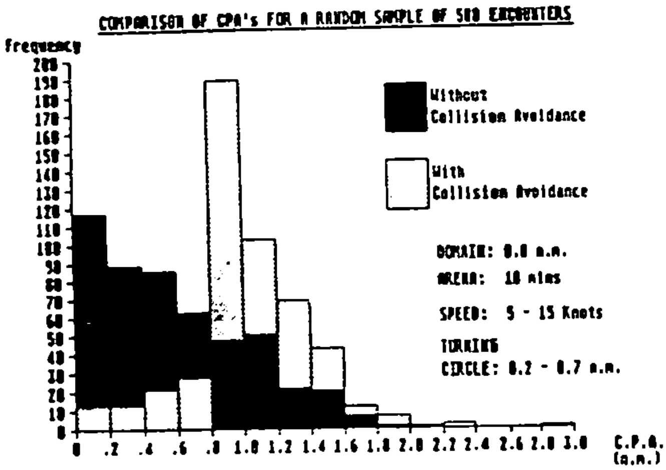

Figure 8

Comparison of CPA's for a Random Sample of 500 Encounters

The results of this multiple simulation illustrate that, in the large majority of cases, potential domain violations were avoided by the expert system invoking appropriate collision avoidance strategies; moreover, the manoeuvres involved did not, in the main, involve excessive course alterations - closest point of approach (CPA) separations for these manoeuvres are clearly bunched just outside the domain boundary. However, a small number of the encounters still show CPA's within the domain boundary - this is clearly unacceptable. The reasons for this are considered to be threefold:

1. The rule base in the initial expert system module was shown to have certain `blind spots`, in which each vessel assessed the other as the give-way vessel, and the required avoidance action was not taken;

2. In using a fairly simple model for generation of simulated vessels, some realism was lost - vessels with unlikely combinations of characteristics were involved in a number of the offending encounters;

3. A small number of these simulations highlighted the need for an increased level of sophistication in the strategy for recognising and handling encounter situations. Specifically, the fixed look-ahead, used to assess the time to initiate avoidance manoeuvres, appeared inadequate in certain cases; an ongoing look-ahead, extrapolating at regular intervals from current status, seemed indicated to ensure action at the optimum time. This is considered to be the most important development to arise from the multiple simulation exercise, and will form a central feature of future developments.

AUTOMATIC GUIDANCE AND COLLISION AVOIDANCE

At present, the collision avoidance system described in this paper acts as an expert advisor. The next stage of development is to link the guidance and collision avoidance systems together to form an expert controller. The advent of such systems should be viewed along-side recent advances in Vessel Traffic Systems (VTS). Discussions in the International Maritime Organisation (IMO) and the European Economic Community (EEC) have contributed to a global dialogue on the future aims and objectives of VTS.

The challenge facing the maritime world over the next decade will be the implementation of VTS in all major ports, using advanced surveillance and communication techniques, together with shipborne automatic guidance and collision avoidance systems. Such systems will increase the efficiency of port operation and at the same time improve safety levels, particularly in poor weather conditions.

NOMENCLATURE

Matrices and Vectors

A	Discrete State Transition Matrix
B	Discrete Control Matrix
C	Discrete Disturbance Matrix
F	Continuous Time System Matrix
G, G_C, G_D	Continuous Time Forcing Matrices
Q	State Error Weighting Matrix
R	Control Weighting Matrix
r	Desired State Vector
u	Control Vector
w	Disturbance Vector
x	Ship Related State Vector

Scalar Symbols

J	Cost Function or Performance Index
k	Integer Counter
t	Continuous Time
t_o, t_1	Initial and Final Times
u_a, u_d	Actual and Desired Forward Velocity
X_o, Y_o	Earth Co-ordinate System
ψ_e	Course Error

REFERENCES

Abkowitz, M. 1964. "Lectures on Ship Hydrodynamics, Steering and Manoeuvrability." Hy-A Report, Hy, 5.?, Denmark.

Blackwell, G.K., R.S. Burns and C.T. Stockel. 1987. "An Expert System Approach to Collision Avoidance." Proceedings of Summer Computer Simulation Conference, Montreal. Society for Computer Simulation, July, 1987, pp. 43-48.

Burns, R.S., M.J. Dove, T.H. Bouncer, C.T. Stockel. 1985. "A Discrete, Time Varying, Non-linear Mathematical Model for the Simulation of Ship Manoeuvres." First International Maritime Simulation Symposium, Munich, June 1985.

Coldwell, T.G. 1981. "Marine Traffic Flow and Casualties on the Humber." Journal of Navigation, Volume 34,000.

Colley, B.A., R.G. Curtis, C.T. Stockel, 1983. "Manoeuvring Times, Domains and Arenas." Journal of Navigation. Volume 36, pp 324-328.

Dove, M.J., R.S. Burns, C.T. Stockel, T.H. Bouncer. 1985. "Simulation of a Digital Filter/Estimator for the Navigation of Large Ships in Confined Waters." Proc. First Intercontinental Symposium on Maritime Simulation, Control Data Corporation. Schliersee, W. Germany, pp 147-160.

International Marine Organisation, 1981. "International Regulations for Preventing Collisions at Sea."

Panel on Human Error in Merchant Marine Safety. 1976. "Human Error in Merchant Marine Safety." National Academy of Science, Washington D.C.

SIMULATION METHODOLOGIES, LANGUAGES AND ARCHITECTURES AND AI AND GRAPHICS FOR SIMULATION

PROCEEDINGS OF THE
1989
EUROPEAN SIMULATION MULTICONFERENCE

JUNE 7 - 9, 1989

AMBASCIATORI PALACE HOTEL
ROME, ITALY

EDITED BY
G. IAZEOLLA
A. LEHMANN
H.J. VAN DEN HERIK

ORGANIZED BY
THE SOCIETY FOR COMPUTER SIMULATION INTERNATIONAL

IN COOPERATION WITH
ASIM, DBSS, ISCS, SIMS, UKSC, JSST AND CSSC

PRINTED IN BELGIUM

Look-Ahead Simulation as a Tool for Real-Time Expert Systems

G.K. Blackwell, C.T. Stockel and B.A. Colley
Plymouth Polytechnic
United Kingdom

ABSTRACT

An expert system (E.S.) should be capable of making rational decisions or inferences, based on available information. One common ploy in E.S. applications is the 'What if..' strategy: the consequences of a particular action are followed through to their logical conclusion, which is then assessed; the original action is then judged on the basis of its final outcome.

In many real-time control systems, it is not possible to predict future events directly; circumstances in a future time-frame may only be evaluated by modelling the intervening states within the E.S., and following through the decision process for each of those states.

The authors are concerned with development of a real-time expert system for marine collision avoidance. A major factor in decision-making by this E.S. has been the projected time to collision (or near-miss) with a potential hazard. In the initial version, a ship would be assigned a set value, according to its performance characteristics; when that critical value was reached, the E.S. would assess the situation and act accordingly.

In the light of results from (simulated) test runs, the system has been enhanced by incorporating a periodic look-ahead simulation, as described above. At each time-step, the linear prediction of time to collision/near-miss is used as a first estimate of time for action; from a predetermined 'safe outer limit' back, multiple simulation runs within the E.S. model possible manoeuvres until a critical time-to-go is reached. Each simulation run consists of a complete sequence of events, from initiating a turn until both vessels are safely clear of each other and back on course. In this way, both the timing and the nature of the manoeuvre are checked in advance to ensure safe clearance without excessive course deviation - assuming no untoward action on the part of the hazard vessel. This Predetermined Safe Manoeuvring Time (PSMT) is used to schedule any action. During the manoeuvre, the E.S. is required to keep a watching brief on the hazard, and adjust its 'game plan' accordingly, if necessary.

The enhanced E.S. is housed in a dedicated microcomputer, with communication links to two other computer systems, simulating actions of own-ship and hazard vessel respectively. A multi-window user interface gives a graphical display of ships' tracks plus information on ships' status, both updated in real time. This E.S. has proved totally reliable in handling every configuration of 2-ship encounter tested to date. Moreover, test runs are carried out at 10 to 20 times real-time speeds, indicating ample scope for expansion to handle more complex scenarios.

Figure 1.

Multiple Look-Ahead Simulation.

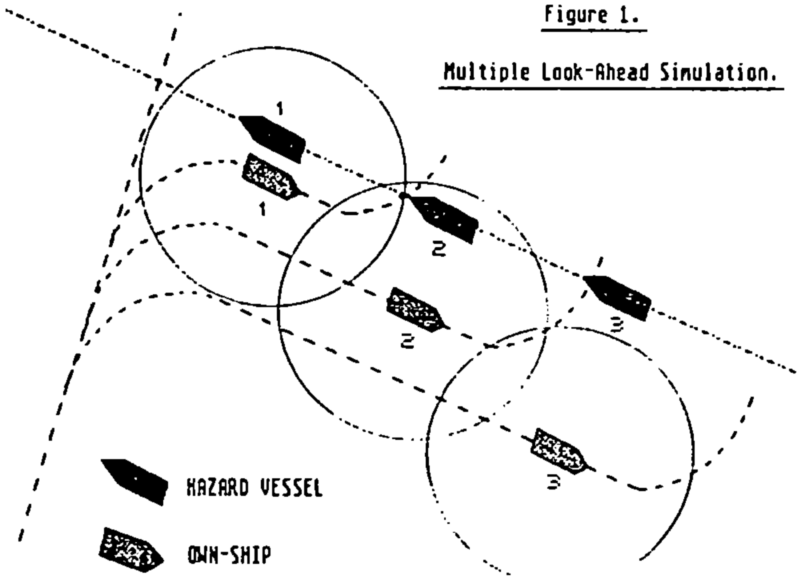

Figure 1 (opposite) shows a sequence of three look-ahead simulations of a crossing encounter, with own-ship taking avoiding action. Starting times for successive simulations work backwards from a predefined minimum.

Simulations 1 and 2 both show own-ship's domain (safety zone) being infringed by the hazard vessel. Simulation 3 identifies a safe time-to-go for the avoidance manoeuvre.

1. INTRODUCTION

The master of a modern vessel, standing on the bridge of his ship, receives a wide range of information from a variety of sources: instruments connected to relevant equipment give him ship's rudder setting and engine rpm; appropriate sensors give information on speed through the water, pitch, roll, and yaw; navigational aids such as Decca, Omega, Loran-C and satellite navigation systems provide accurate data (typically within a few metres) for fixing position and ground speed (Ie. true speed, relative to the sea-bed); detector systems such as radar and sonar warn of obstacles above or below the surface, giving range and direction - an enhancement of radar, the ARPA (Automatic Radar Plotting Aid) system, even gives the (relative) speed and course of any moving hazard in the vicinity, such as another ship. Additionally, of course, he has his own eyes and ears, and those of his crew, to fill in any gaps left by these electronic aids.

The system under development is intended ultimately to be yet another electronic aid on the bridge, one which will receive data from a variety of these sensors (plus chart data from on-board 'electronic charts') and give a reasoned response to prevailing circumstances. Clearly 'eyes and ears' information may be input via a keyboard, when necessary. Rather than proliferating the mass of data to be assimilated by the ship's master, the purpose of this device is to collate relevant data in the field of collision avoidance. Reasoned conclusions and recommendations may then be presented in an easily-digested form, to hopefully reduce the scope for error by omission, misreading, or simply 'information overload'.

The development environment consists of two simulated ships in open-sea conditions. One, termed 'own-ship', is under the guidance of the expert system described here; the other, the 'hazard vessel', is also guided by machine intelligence, with an option to manually override such control. The rule base for the expert system is based on the International Regulations for Preventing Collisions at Sea [1], supplemented by substantial research into good practice, as observed by experienced mariners. Output is presented in the form of a multi-window display, combining graphical, numerical and textual information.

2. PRINCIPLES OF OPERATION

At the present phase of development, the rule base consists of rules for:

(a) Identifying the presence of a potential hazard, and assessing the threat in terms of expected time to infringement of the domain - this time factor is central to the decision on when to take avoiding action, if necessary;

(b) Identifying the type of encounter, including several variants in some cases (the primary types of encounter being head-on, crossing and overtaking), and fixing the status of own-ship and the perceived status of hazard-ship in the encounter (head-on or give-way);

(c) Negotiating the stages of the encounter, with appropriate safety margins.

The rules identifying and handling the three types of encounter are as follows:

(1) Head-on
Two vessels are approaching each other on near- reciprocal bearings (I.e. + or - 6 degrees, in this system). A vessel in such an encounter must consider itself give-way and alter course to starboard (right); it should not begin to come back onto course until the other vessel has passed 90 degrees on the port beam (left)

(2) Crossing
Two vessels are approaching each other at an angle such that each can see one side of the other; more specifically, one vessel is in the sector subtended by the port or starboard light - 112.5 degrees measured from the bow (front) on each side. The ship with the other vessel to starboard should give way, altering course to starboard and going behind. It should not alter back until the stand-on vessel has passed 15 degrees to port of own-ship's original heading;

(3) Overtaking
A ship is deemed to be overtaking if it is approaching from within the stern-light sector, I.e. the rear 135-degree arc. It should give way as follows:
(a) If on near-parallel course (Within 10 degrees, in this system), go behind to starboard; alter back when safe to do so;
(b) If on port quarter of stand-on vessel, action as for (a);
(c) If on starboard quarter, and set to overtake ahead on current course, AND if able to parallel courses without infringing domain, parallel-up; when well past and clear to pull across, return to original heading;
(d) Otherwise, go behind to port, alter-back when safe to do so.

The inference engine for this system rests primarily on previous research in the field of marine collision avoidance. In particular, two well-tried concepts dictate the nature and extent of action taken: the domain, or 'danger zone' around a ship, into which no other vessel or obstacle may be allowed to encroach; and the arena, a larger zone around a vessel, inside which any other vessel or obstacle must be considered a potential hazard.

The concept of the domain was advanced by Goodwin [2] as 'the effective area around a ship which a navigator would like to keep clear with respect to other ships and static objects'. The expert system has as its first objective the protection of the integrity of the domain. The arena concept was introduced by Davis [3] in response to the observed action of mariners manoeuvring to prevent infringement of their domain by ships in their

icinity. This concept was refined by Colley
1] to a time factor, referred to as RDRR:-
ange-to-Domain/Range-Rate - basically, the
ime-to-go (before calculated domain infringe-
ent) at which the mariner will assess the
tatus of his vessel and take appropriate
ction. For example, on the basis of a ship's
andling characteristics, a mariner may choose
o take action 15 minutes before potential
omain infringement - a 15-minute RDRR.

olley's RDRR is used to define minimum action
ime; the look-ahead facility described in
ection 5 then modifies that RDRR value in the
ight of current circumstances - guaranteeing
afe, but not excessive, manoeuvring time.
his Predetermined Safe Manoeuvring Time
PSMT) replaces the previous 'best guess' with
certainty (barring untoward incidents -
ection 5 considers this point).

nother basic principle of collision avoidance
anoeuvres is that not only must they be
arried out, they must also be seen to be
arried out. No other form of communication
etween vessels can be assumed, and so any
voidance action must be clearly identifiable
s such from another vessel. Such manoeuvres,
nder the direction of the expert system, must
e of sufficient magnitude (a 'clear turn')
nd in good time, to indicate to the other
essel that the encounter has been anticipated
nd is being dealt with. Any decision to use
he processing capabilities of the computer to
optimise' course changes purely on the basis
f 'safe clearance' criteria must be moderated
y such considerations.

he current model uses a minimum turn of 30
egrees, a value drawn from previous research
y Colley [5]. The minimum RDRR is defined as
parameter for each vessel - typically 10
inutes, but varied according to the handling
haracteristics of the vessel being simulated.

he rules in the E.S. embody the International
egulations for Preventing Collisions at Sea,
lus clarification of specific practical
etails as noted by Colley and used most
ffectively in an earlier simulation. Using
hese rules, a collision avoidance manoeuvre
s taken through five stages, as follows:

1) 20 minutes before projected domain infringement, look-ahead simulation is used to identify a suitable PSMT, starting from the minimum RDRR.

2) When the PSMT is reached, the give-way vessel initiates a turning action;

3) Once the required turn is effected (minimum 30 degrees), the vessel proceeds on a straight course;

4) When the original heading may be resumed without further risk of domain infringement (and any other conditions are met), alter-back is initiated, necessitating a turn in the opposite direction;

5) When back on the original heading, vessel proceeds on that heading; manoeuvre over.

Figure 2.

Rule Structure.

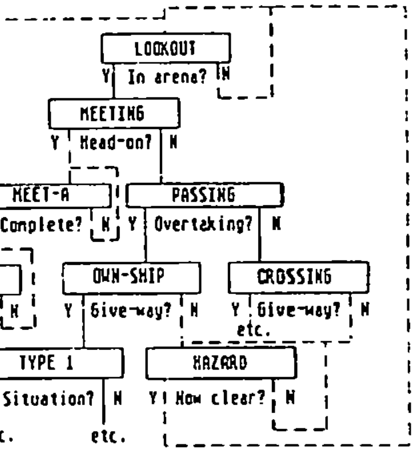

3. INFERENCE ENGINE AND RULE STRUCTURE

The function of the inference engine at each time-step is to:

(a) receive the parameters from the simulator units (speed, heading, position), and from them generate a set of secondary parameters - relative bearing, relative velocity components, time to projected domain infringement, etc;

(b) Apply the appropriate rule to these parameters - depending on status of vessel as perceived at the previous time-step - to ascertain whether the status of the vessel has changed, and if so, what that new status should be; Each rule involves a test on these system values, and may involve evaluation of further mathematical functions of these values - such functions form part of that rule.

(c) If the status of the vessel has changed, flags are set for ship simulator action; also, a change of status may lead to further rules being invoked immediately; the inference engine will continue to process successive rules until a 'defer' flag is reached, indicating that no further action on the rule base is required until the next time-step.

The rule base is a hierarchical structure, based on the structure of a binary tree. However, the left and right links from each node in this 'tree' structure are rather more flexible, in that two or more links may lead to a common node at a lower level, or a link may loop back to a node at a higher level. A schematic of this structure, showing most of the initial rule base, is given in Figure 2.

The rule base is also in effect a state table, each node representing a possible state of own-ship - no encounter currently in progress, second stage of a crossing encounter, etc. Each node contains:

(1) Left and right links to other nodes in the rule structure (or looping back to the same node) - one of these links will be followed at every decision step;

(2) A mathematically-based decision function, used to determine which link is taken at this stage;

(3) A flag for each link, indicating in each case whether the next rule is to be processed immediately, or deferred until the next time-step;

(4) Ship status flags, to pass information to simulator and display.

Enhancement of the rule base involves creation of the appropriate new node(s), and resetting of the relevant links so as to position the node(s) correctly within the rule structure.

Each of the rules involves a yes-no decision function, using inequalities based on current values of displacement and velocity vectors. The majority of processing is thus concerned with solution of triangles, and intersections of straight lines with circles - substantial evaluation of trigonometrical functions. It is anticipated that this aspect of the E.S. will escalate as more complex situations are taken into account; this has been a major factor in determining choice of language and form of rule structure for this time-critical system.

4. EVALUATION OF INITIAL RULE STRUCTURE

A multiple run of 500 simulated encounters was used to evaluate the effectiveness of the rule base used in the first prototypes [6]. For the purposes of this exercise, the RDRR criterion used for avoidance decision/action in each case was a fixed 10-minute interval before anticipated domain infringement.

Each simulation was generated as follows: Two vessels, each with random course, speed and turning circle (See Fig. 3 for bounds on latter two parameters), were placed within a circle of radius 1 nautical mile. Given a fixed domain radius of 0.8 N.M. for each vessel in this exercise, this ensured that each was very likely to consider the other a hazard, according to their terms of reference. The two vessels, one designated own-ship, the other hazard-ship, were then effectively taken back in time ten minutes along their tracks, (or more, if an encounter situation already existed by that time), and the simulation allowed to proceed from that point. The separation at the closest point of approach (CPA) between the two vessels was calculated (i) with no collision avoidance manoeuvre attempted, and (ii) subject to the collision avoidance manoeuvre brought into play by the expert systems module. The results of this exercise are shown in the bar chart in Fig. 3.

These results show that, in the large majority of cases, potential domain violations were avoided by the expert system invoking appropriate collision avoidance strategies; moreover, the manoeuvres involved did not involve excessive course alterations - CPA separations for these manoeuvres are clearly bunched just outside the domain boundary. However, a small number of the encounters still show CPAs within the domain boundary.

The offending encounters were picked out from the data-logging file and analysed. Three main causes were identified for such failure:

(a) The rule structure contained certain 'blind spots', i.e. cases where each ship considered the other the give-way vessel, and so each ploughed on regardless. The fact that this would occur very seldom in near-miss situations, and never in collision situations, does not make this oversight in any way acceptable.

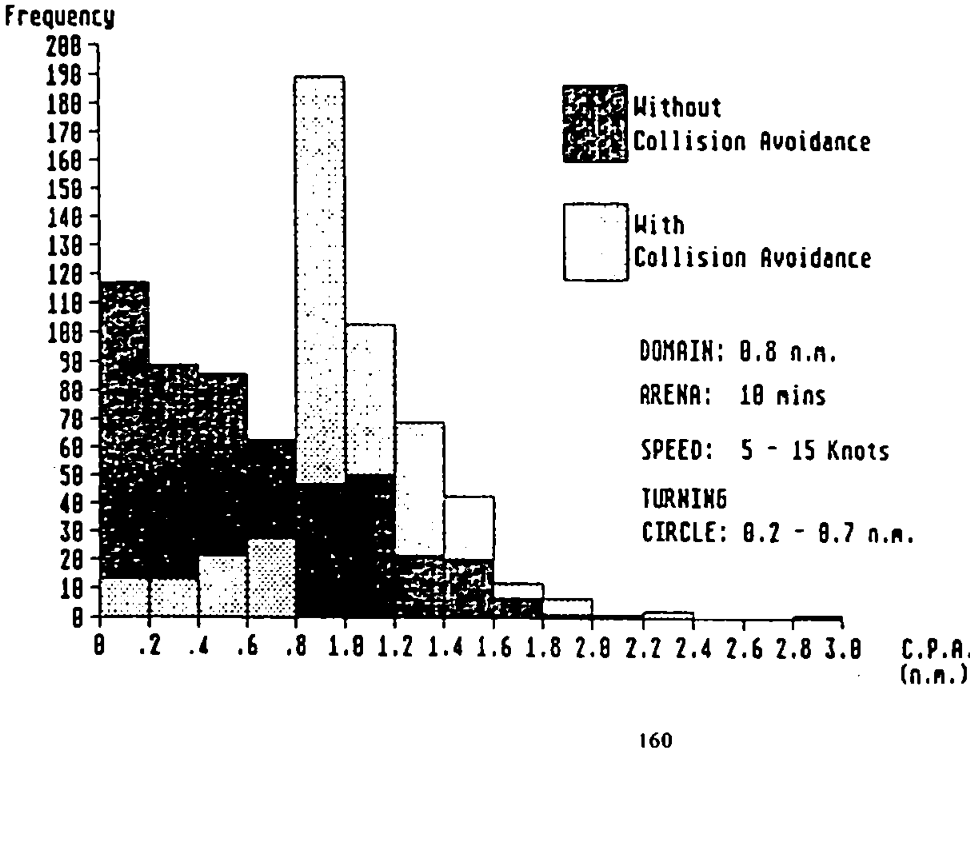

Figure 3.

Comparison of C.P.A.s for a Random Sample of 500 Encounters.

(b) Certain combinations of ship handling characteristics, as thrown up by the 'random encounter generator', gave those vessels inadequate manoeuvrability. A relatively high speed, coupled with a slow rate of turn, prevented a number of vessels from achieving a safe clearance by appropriate manoeuvres within the time allowed for such action (10 minutes).

(c) Certain juxtapositions of vessels, in certain types of encounters, again rendered inadequate the fixed time-to-go criterion of 10 minutes before anticipated domain infringement. For example, ship a, set to cross ahead of ship b from b's port side, may not be quite far enough ahead to clear the domain - i.e. an encounter situation exists. Ship a is required to turn to starboard and pass behind b. This will actually bring a closer to b initially, possibly even exacerbating the problem if the action is not taken in good time (see Figure 1) - one of the simulated encounters in this multiple run showed a CPA of 0.6 N.M. reduced to 0.2 N.M. by collision avoidance manoeuvring!

The first two problems may be resolved by amendment and enhancement of the rule base, and refinement of the ship model selection process for test runs, respectively. The situation described in the third case could be resolved by changing the RDRR criterion to 15 or 20 minutes. However, this simply changes the problem to one of unduly early, and thus excessive, manoeuvring in less critical cases. A more flexible approach was required.

5. LOOK-AHEAD SIMULATION

A characteristic of expert systems is that intuitive reasoning by the human expert must be emulated in a computer system by an explicit logical sequence of cause-and-effect. Such reasoning may be difficult to identify, let alone rationalise and then emulate. A classic example is the decision on when to evaluate an encounter and initiate collision avoidance action. Extensive observation would suggest that a competent mariner has in mind a fairly well-defined time-to-go before domain infringement, at which point in time the nature of the encounter is identified and appropriate avoidance action decided upon; too soon, and the course alteration will be excessive; too late, and domain infringement may not be avoided.

However, results from the multiple encounter simulation above show that such a fixed 'time-to-go' criterion is not adequate for all possible encounter situations. It seems most probable that the expert mariner would readily identify those cases which require earlier-than-usual action; an expert system, of the type under consideration here, should likewise be capable of doing so. This has prompted the inclusion of just such a facility, one which will identify the optimum decision time for any particular scenario. Such a dynamic time-to-go is pinpointed by consideration of the eventual consequences; this means inclusion of a look-ahead facility as an integral part of the decision-making process.

20 minutes before projected domain infringement, the Expert System simulates a complete encounter as it would proceed if action were to be taken at, say, 10 minutes before domain infringement (This minimum figure is variable, being input as an initial parameter). If this time interval proves inadequate, the system will run another simulation, this time for a 12-minute time-to-go. The decision threshold is moved successively back, with a new simulation for each, until a satisfactory time-to-go is identified. The Expert System then allows the encounter to proceed to that PSMT (Predetermined Safe Manoeuvring Time), whereupon the decision on encounter type, and appropriate avoidance action, is taken.

The incorporation into the Expert System of such a facility, requiring possibly several complete encounter simulations to achieve its objective, will clearly put pressure on the time constraints of a real-time system. It is essential that the necessary processing be completed well in advance of the critical phase of the encounter. In the present idealised situation - two ships in open sea - processing time is not a problem. However, it is envisaged that in more complex scenarios (E.g. those involving a number of vessels), such processing may require other decision-making logic to be suspended for several time-steps, whilst it is under way.

This technique has proved successful in handling those encounter situations which were inadequately dealt with by previous versions, using a fixed RDRR criterion. The (apparently) logical extension of this principle, to a JIT (Just-in-time) system, which leaves avoidance action to the last possible moment, has not been pursued. There are two reasons for this:

(i) A prime requirement of this system is that it gives clear indication of intent, in good time, to other ships in the vicinity; the JIT approach would not meet such a requirement, but would instead probably lead to masters of other ships taking emergency action in many cases;

(ii) The JIT approach leaves no leeway for emergency action, should this become necessary due either to untoward action by another vessel or some other unforeseen problem (E.g. rudder failure).

Whilst it is impossible to legislate for every mishap which may occur at sea, a suitable minimum RDRR and a tested PSMT appear to mirror that balance of prudence and efficiency practiced by experienced mariners.

With the introduction of look-ahead simulation into the decision process, processing needs for the E.S. have expanded, highlighting the need to move towards the final objective - a dedicated computer housing the E.S., acting on data input from other sources. The development environment now in use [7] consists of three linked micro-computers (see figure 4): one houses the E.S. for own-ship; a second acts as own-ship simulator, the third provides 'radar' data for a hazard ship, driven by a simulator module and simplified E.S. module - with manual override, to simulate a 'rogue' ship.

6. USER INTERACTION

The complete system comprises three Atari ST computers, all programmed in the 'C' language. Hazard-ship and own-ship-E.S. computers show displays of the encounter as seen from opposing standpoints. The format of the multi-window display is shown with annotations in the figure opposite (Figure 5).

Initial parameters for both vessels are input to the relevant 'simulator' computers. Domain size and RDRR value for own-ship are input to its E.S. computer. The encounter is then initiated, and control inputs (Eg. review requests) made via own-ship's E.S. computer.

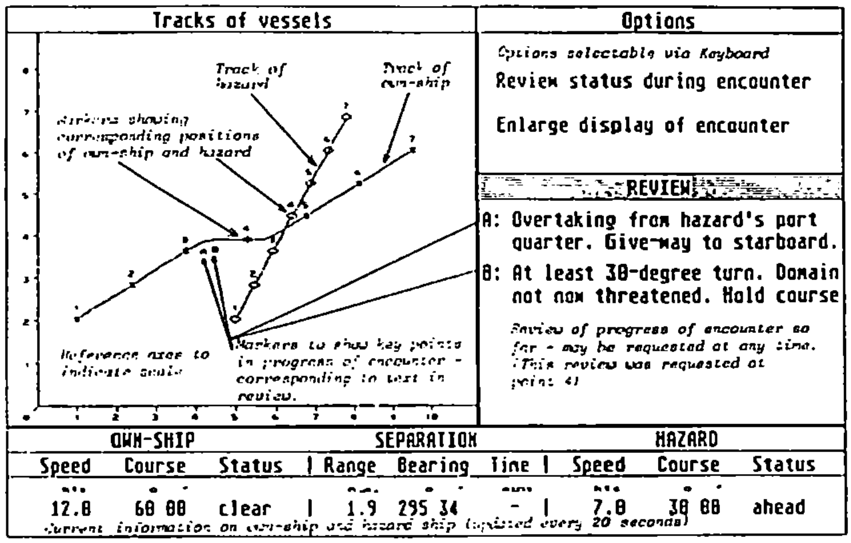

The problem of defining an appropriate time for avoidance action has been a subject of consideration by a succession of researchers in this field. Look-ahead simulation, leading to a proven PSMT, would seem to come as close as possible to guaranteeing domain integrity without undue loss of efficiency. On a more general front, the technique of look-ahead simulation has shown its worth as a decision tool in the real-time expert systems field: forward planning in any time-variant rule-based system must benefit by such an approach.

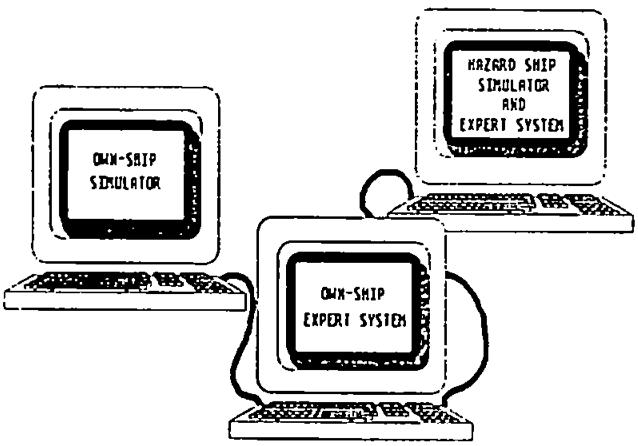

Figure 4.

7. CONCLUSIONS

The inference engine and rule structure developed for this application appear to meet the perceived need. The rule base itself is at present fairly simplistic, dealing only with open-sea encounters between two vessels. However, the system has been designed with a view to very substantial expansion within the current framework; this should accomodate any necessary increase in complexity of the rules.

The real-time nature of the task places severe time constraints on processing; these would appear to be adequately met (present and anticipated future needs) by the design and implementation of the E.S.. The review process is also strongly influenced by the time-based decision process; the optional review provides a brief chronological account of reasons for decisions at stages in the encounter. It is likely that this feature will need revision to handle more complex sequences of rules.

8. REFERENCES

1. International Marine Organisation. 1981. "International Regulations for Preventing Collisions at Sea".

2. Goodwin,E.M.. 1975. A Statistical Study of Ship Domains. Journal of Navigation. Volume 28. pp328-344.

3. Davis,P.V.. Dove,M.J. & Stockel,C.T.. 1980. A Computer Simulation of Marine Traffic Using Domains and Arenas. Journal of Navigation. Volume 33. pp215-222.

4. Colley,B.A.. Curtis,R.G. & Stockel,C.T.. 1983. Manoeuvring Times, Domains & Arenas Journal of Navigation. Vol. 36. pp324-328.

5. Colley,B.A.. 1985. Computer Simulation of Marine Traffic Systems'. PhD Thesis.

6. Blackwell,G.K.. Colley,B.A. & Stockel,C.T.. 1988. An Intelligent Knowledge-Based System for Marine Collision Avoidance. Proceedings of Conference on Maritime Communications and Control. Institute of Marine Engineers. October 1988. pp167-174.

7. Blackwell,G.K. & Stockel,C.T.. 1989. Simulation of Ship Encounters using Multiple Linked Intelligent Systems. To be presented at 1989 Summer Computer Simulation Conference. SCS International. Austin (Texas). July 1989.

PROCEEDINGS OF THE 1989 SUMMER COMPUTER SIMULATION CONFERENCE

JULY 24-27, 1989
THE STOUFFER AUSTIN HOTEL
AUSTIN, TEXAS

TWENTY-FIRST ANNUAL SUMMER
COMPUTER SIMULATION CONFERENCE

Edited by
Joe K. Clema

CONFERENCE SPONSOR

SCS

THE SOCIETY FOR COMPUTER SIMULATION

Printed in the United States of America

Simulation of Ship Encounters Using Multiple Linked Intelligent Systems

G.K.Blackwell and C.T.Stockel
Department of Computing, Plymouth Polytechnic
Drake Circus, Plymouth, Devon, England

ABSTRACT

A computer-based simulation of any process or system requires definition of formulae governing interaction between different entities or sub-systems, to a level of precision considered adequate for the purposes of the exercise. Once such formulae have been derived, either analytically or empirically, the simulation may generally then be programmed as a single composite unit.

However, such an approach presupposes that none of the entities taking part in the simulation possesses any innate intelligence, and that interaction will take place on a purely mechanistic level. If the system being simulated includes an intelligent entity, capable of (and responsible for) more complex decision-making, then other factors must be considered. Notably, a clear distinction must be drawn between the full set of system state variables and those which are available to this intelligent entity for decision-making purposes. Not only must this be so, it must be seen to be so, to maintain the validity of the simulation exercise.

With the advent of the Expert System (E.S.), this becomes more than an academic problem. Particularly in control applications, where an E.S. may be part of a process being simulated, it is essential to draw a clear line between the knowledge base for the E.S. and the parameter set for the simulation. This may be most effectively and demonstrably achieved by housing the E.S. in separate hardware; communication links may then pass appropriate data to the E.S., and control signals back to the main simulation database.

Development of an earlier E.S., for marine collision avoidance, has now reached a stage where it will operate on a stand-alone basis in a dedicated microcomputer system. The hypothetical vessel guided by this E.S. is simulated in a second microcomputer, linked to the E.S. by 2-way communication channels. A third microcomputer simulates the 'hazard' vessel, guided by a simplified E.S. module; links between the knowledge bases of the two vessels carry data which would normally be available via radar and other navigational aids.

This multi-computer simulation appears to be a very effective way of checking the performance of such an E.S. prior to actual field trials. For its part, the E.S. has proved highly satisfactory in handling encounter situations.

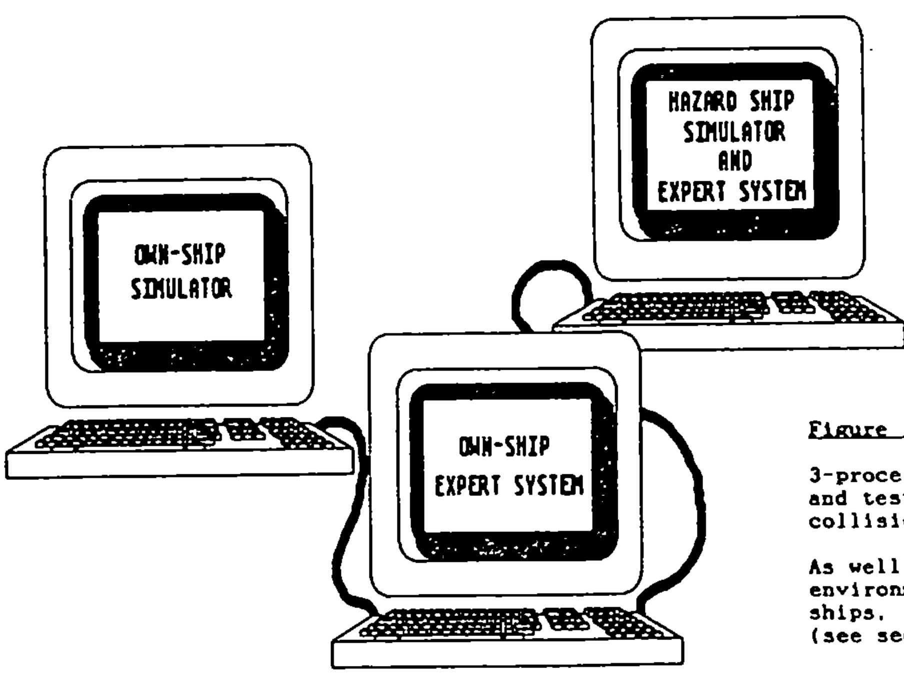

Figure 1.

3-processor network - development and testing environment for marine collision avoidance expert system.

As well as the expert system, the environment inludes two simulated ships, own-ship and hazard vessel (see section 3).

0-911801-58-X/89/0000-0766$02.00©1989SCi

INTRODUCTION

mputer simulation is well established as a sign tool and planning aid in a wide variety fields. It may be used to save time, save ney, reduce risk, or facilitate detailed spection of processes. As a first step, the levant subsystems must be identified. The les governing interaction between subsystems ould then be defined to a level of precision propriate to the purposes of the simulation. ce those rules have been refined into a pure rm - generally as mathematical formulae - e composite process may then be regarded as single unit, and programmed as such. Interlationships between elements of the system pear in the final package as functions of ate vectors, or state transition tables.

is general technique can be applied in some rm to any simulation where the link between use and effect can be viewed as mechanistic. e transition function approach breaks down, owever, when one element of the simulation s an intelligent entity; in such cases, action to a given set of conditions cannot nerally be programmed as a set of athematical functions or decision matrices. he reasoned response of such an entity to revailing circumstances is itself a complex rocess, requiring a totally different pproach.

his is not to say that all simulations which nvolve humans (who are all, to some extent, ntelligent) present such problems: population rowth, actuarial figures, and many other such odels depend on a statistical approach to umanity 'en masse' which fits observed facts; uman operators of ship or flight simulators hould be regarded as actuators of simulation equences, rather than part of the simulation rocess itself. The problem only arises when n intelligent entity plays an integral part n the simulation process.

This has not generally presented a problem in the past, since one object of simulation has been to eliminate the human factor. With the advent of intelligent knowledge-based systems (IKBS), a reasoning element has been brought into machine processes; as real-time control IKBSs are developed, a suitable simulated test environment is required. Many applications of such IKBSs rely on rigorous simulation testing prior to actual field trials. IKBS techniques may also be used to undertake simulation of a kind not previously feasible, where the human decision-making process forms an integral part of a larger picture; here machine intelligence may be used to model response at an individual level, rather than on the macro scale referred to previously.

As with other real-time response systems, the validity of the model or the test environment requires that timing is true-to-life. Timing considerations in most mechanistic systems can be dealt with by reference to a system clock or interrupt timer. Timing of intelligent systems cannot generally be pinned down in this way; in addition, processing demands of all but the simplest rule structure are likely to be substantially greater than those in most mechanistic systems - particularly for control applications making use of mathematical library functions (trig, logs, stats, etc).

Equally important, in considering the validity of a simulation exercise, is the nature and scope of the information being passed to the intelligent system, on which it bases its decisions. It is clearly inappropriate to make available all of the data currently in use in the simulation, if the final control system, or the human counterpart of this IKBS, would only have access to part of that information - i.e. a clear separation between the database for the simulation and the knowledge base for the machine intelligence is an essential feature of such a simulation environment.

Given such considerations, a logical approach is to treat the IKBS element of the model in the same way as the 'human actuator' referred to earlier - physically distanced from the rest of the simulation, with suitable interfaces for transfer of information and control signals. Timing of the IKBS is independent of other processing, and with adequate dedicated CPU power should not be a problem; transfer of information to and from the IKBS is explicit, and therefore open to verification; there is (or should be) no risk of over-simplification of the decision process by 'short-cuts' inadvertently introduced into the simulation logic, as might happen in a single composite package.

This approach has been used, for the reasons given above, in a simulation incorporating an Expert System (E.S.) for marine collision avoidance. The Rule Base for this system is founded on the International Regulations for Preventing Collisions at Sea [1], supplemented by substantial research into good practice, as observed by experienced mariners. The complete simulation environment comprises three microcomputers, as shown in figure 1. The eventual objective of this project is a bridge-based expert adviser for collision avoidance. This system will receive data from various sensory inputs (navigational aids, radar, ship's speed and rudder position, etc), and give reasoned recommendations as to appropriate action.

2. PRINCIPLES OF OPERATION

At the present phase of development, the rule base consists of rules for:

(a) Identifying the presence of a potential hazard, and assessing the threat in terms of expected time to infringement of the ship's domain (danger zone) - this time factor is central to the decision on when to take avoiding action, if necessary;

(b) Identifying the type of encounter, including several variants in some cases (the primary types of encounter being head-on, crossing and overtaking), and fixing the status of own-ship and the perceived status of hazard-ship in the encounter (head-on or give-way);

(c) Negotiating the stages of the encounter, with appropriate safety margins.

The rules identifying and handling the three types of encounter are as follows:

(1) Head-on
Two vessels are approaching each other on near- reciprocal bearings (I.e. + or - 6 degrees, in this system). A vessel in such an encounter must consider itself give-way and alter course to starboard (right); it should not begin to come back onto course until the other vessel has passed 90 degrees on the port beam (left)

(2) Crossing
Two vessels are approaching each other at an angle such that each can see one side of the other; more specifically, one vessel is in the sector subtended by the port or starboard light - 112.5 degrees measured from the bow (front) on each side. The ship with the other vessel to starboard should give way, altering course to starboard and going behind. It should not alter back until the stand-on vessel has passed 15 degrees to port of own-ship's original heading;

(3) Overtaking
A ship is deemed to be overtaking if it is approaching from within the stern-light sector, I.e. the rear 135-degree arc. It should give way as follows:
(a) If on near-parallel course (Within 10 degrees, in this system), go behind to starboard; alter back when safe to do so;
(b) If on port quarter of stand-on vessel, action as for (a);
(c) If on starboard quarter, and set to overtake ahead on current course, AND if able to parallel courses without infringing domain, parallel-up; when well past and clear to pull across, return to original heading;
(d) Otherwise, go behind to port, alter-back when safe to do so.

The inference engine for this system rests primarily on previous research in the field of marine collision avoidance. In particular, two well-tried concepts dictate the nature and extent of action taken: the domain, or 'danger zone' around a ship, into which no other vessel or obstacle may be allowed to encroach; and the arena, a larger zone around a vessel, inside which any other vessel or obstacle must be considered a potential hazard.

The concept of the domain was advanced by Goodwin [2] as 'the effective area around a ship which a navigator would like to keep clear with respect to other ships and static objects'. The expert system has as its first objective the protection of the integrity of the domain. The arena concept was introduced by Davis [3] in response to the observed action of mariners manoeuvring to prevent infringement of their domain by ships in their vicinity. This concept was refined by Colley [4] to a time factor, referred to as RDRR:- Range-to-Domain/Range-Rate - basically, the time-to-go (before calculated domain infringement) at which the mariner will assess the status of his vessel and take appropriate action. For example, on the basis of a ship's handling characteristics, a mariner may choose to take action 15 minutes before potential domain infringement - a 15-minute RDRR.

Colley's RDRR is used to define minimum action time; the look-ahead facility described in section 6 then modifies that RDRR value in the light of current circumstances - guaranteeing safe, but not excessive, manoeuvring time. This Predetermined Safe Manoeuvring Time (PSMT) replaces the previous 'best guess' with a certainty (barring untoward incidents - section 6 considers this point).

Another basic principle of collision avoidance manoeuvres is that not only must they be carried out, they must also be seen to be carried out. No other form of communication between vessels can be assumed, and so any avoidance action must be clearly identifiable as such from another vessel. Such manoeuvres, under the direction of the expert system, must be of sufficient magnitude (a 'clear turn') and in good time, to indicate to the other vessel that the encounter has been anticipated and is being dealt with. Any decision to use the processing capabilities of the computer to 'optimise' course changes purely on the basis of 'safe clearance' criteria must be moderated by such considerations.

The current model uses a minimum turn of 30 degrees, a value drawn from previous research by Colley [5]. The minimum RDRR is defined as a parameter for each vessel - typically 10 minutes, but varied according to the handling characteristics of the vessel being simulated.

The rules in the E.S. embody the International Regulations for Preventing Collisions at Sea, plus clarification of specific practical details as noted by Colley and used most effectively in an earlier simulation. Using these rules, a collision avoidance manoeuvre is taken through five stages, as follows:

(1) 20 minutes before projected domain infringement, look-ahead simulation is used to identify a suitable PSMT, starting from the minimum RDRR.

(2) When the PSMT is reached, the give-way vessel initiates a turning action;

(3) Once the required turn is effected (minimum 30 degrees), the vessel proceeds on a straight course;

(4) When the original heading may be resumed without further risk of domain infringement (and any other conditions are met), alter-back is initiated, necessitating a turn in the opposite direction;

(5) When back on the original heading, vessel proceeds on that heading; manoeuvre over.

3. ADVANCED DEVELOPMENT ENVIRONMENT

As previously stated, the system has been developed within an integrated environment, comprising two simulated ships (each with their own knowledge base) in addition to an expert system for decision-making. In the early stages, a single microcomputer housed the total package, alternately actioning a ship simulator module (to 'drive' the two ships) and an E.S. module (providing collision avoidance decision-making for each vessel in turn, independently) at each time-step.

This setup, though adequate for initial prototyping, was considered to have certain drawbacks. The authenticity of the simulation process was felt to be compromised by the coexistence side-by-side of the knowledge bases for the two ships, and the common use of decision-making and simulation routines (despite the independent processing of the two vessels). More important, a single processor was responsible for both simulation and decision-making for two ships. Whilst this has not proved a problem with simple models, increasing complexity makes progressively greater demands on processing time.

With the introduction of look-ahead simulation into the decision process, processing requirements for the E.S. have expanded somewhat. This has highlighted the need to move towards the final objective, a dedicated computer housing the E.S., acting on data input from other sources. The development environment now in use consists of three linked micro-computers (see figure 1): one houses the E.S. for own-ship; a second acts as own-ship simulator; the third provides 'radar' data for a hazard ship, driven by a simulator module and simplified E.S. module - with optional manual override, to simulate a 'rogue' ship.

4. INFERENCE ENGINE AND RULE BASE

The function of the inference engine at each time-step is to:

(a) receive the parameters from the simulator units (speed, heading, position), and from them generate a set of secondary parameters - relative bearing, relative velocity components, time to projected domain infringement, etc;

(b) Apply the appropriate rule to these parameters - depending on status of vessel as perceived at the previous time-step - to ascertain whether the status of the vessel has changed, and if so, what that new status should be: Each rule involves a test on these system values, and may involve evaluation of further mathematical functions of these values - such functions form part of that rule.

(c) If the status of the vessel has changed, flags are set for ship simulator action; also, a change of status may lead to further rules being invoked immediately; the inference engine will continue to process successive rules until a 'defer' flag is reached, indicating that no further action on the rule base is required until the next time-step.

The rule base is a hierarchical structure, based on the structure of a binary tree. However, the left and right links from each node in this 'tree' structure are rather more flexible, in that two or more links may lead to a common node at a lower level, or a link may loop back to a node at a higher level. A schematic of this structure, showing most of the initial rule base, is given in Figure 2.

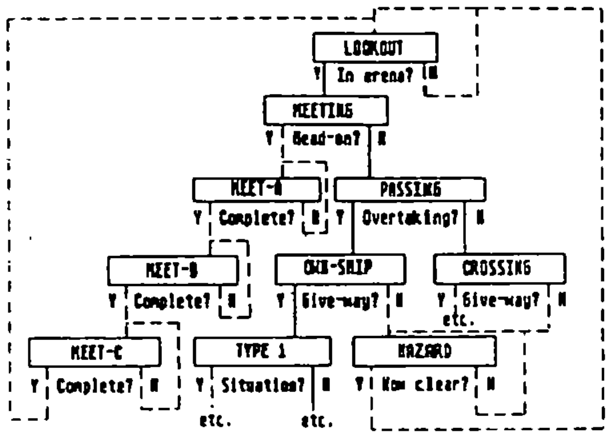

Figure 2.

Decision Network for Rule Base

The rule base is also in effect a state table, each node representing a possible state of own-ship - no encounter currently in progress, second stage of a crossing encounter, etc. Each node contains:

(1) Left and right links to other nodes in the rule structure (or looping back to the same node) - one of these links will be followed at every decision step;

(2) A mathematically-based decision function, used to determine which link is taken at this stage;

(3) A flag for each link, indicating in each case whether the next rule is to be processed immediately, or deferred until the next time-step;

(4) Ship status flags, to pass information to simulator and display.

Enhancement of the rule base involves creation of the appropriate new node(s), and resetting of the relevant links so as to position the node(s) correctly within the rule structure.

Each of the rules involves a yes-no decision function, using inequalities based on current values of displacement and velocity vectors. The majority of processing is thus concerned with solution of triangles, and intersections of straight lines with circles - substantial evaluation of trigonometrical functions. It is anticipated that this aspect of the E.S. will escalate as more complex situations are taken into account; this has been a major factor in determining choice of language and form of rule structure for this time-critical system.

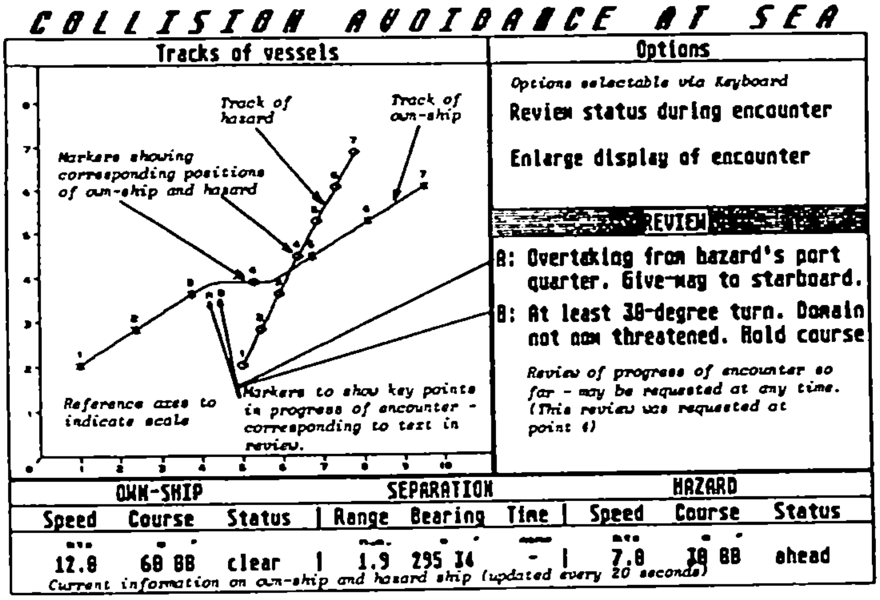

Figure 3.

Annotated example of expert system screen display. An overtaking encounter is shown, with own-ship giving way. A review of the first part of the encounter has been given, on operator's request.

5. DISPLAY AND USER INTERACTION

The complete system comprises three Atari ST computers, all programmed in the 'C' language. 'Hazard vessel' computer and own-ship 'expert system' computer show screen displays of the encounter as seen from opposing standpoints. The format of the multi-window display is shown with annotations in Figure 3.

The graphics window shows tracks of own-ship and hazard; markers at 8-minute (ship time) intervals indicate corresponding points on the two tracks - asterisks for own-ship, diamonds for hazard. The status window at the bottom of the screen gives detail on own-ship's status, and information as far as possible on observed status of hazard vessel. The review window (own-ship only) gives a time-based summary of decisions and actions so far (on request), corresponding to markers on track of own-ship.

Initial parameters for both vessels are input to the relevant 'simulator' computers. Domain size and RDRR value for own-ship are input to own-ship's 'expert system' computer. The encounter is then initiated, and subsequent control information (E.g. Review requests) input, via own-ship's E.S. computer.

The rules for emergency avoidance of 'rogue' hazard vessels (I.e. those which do not obey the collision avoidance regulations) have not yet been added to the rule base. However, the facility for simulating such a hazard has been incorporated into the multi-computer system: The simplified expert system module (used to control the actions of the hazard ship in an encounter situation) is housed in the hazard vessel simulation computer, and linked to the simulator module for that vessel; this E.S. may be overridden, and the actions of the hazard vessel controlled, by an operator at the keyboard of that computer. This facility will be used to test 'emergency action' rules as they are added to the rule structure in the primary expert system.

6. LOOK-AHEAD SIMULATION

Figure 4 (below) shows a sequence of three look-ahead simulations of a crossing encounter, with own-ship taking avoiding action. Starting times for successive simulations work backwards from a predefined minimum RDRR value.

Simulations 1 and 2 both show own-ship's domain being infringed by the hazard vessel. Simulation 3 identifies a safe time-to-go for the avoidance manoeuvre.

Figure 4.

Multiple Look-Ahead Simulation.

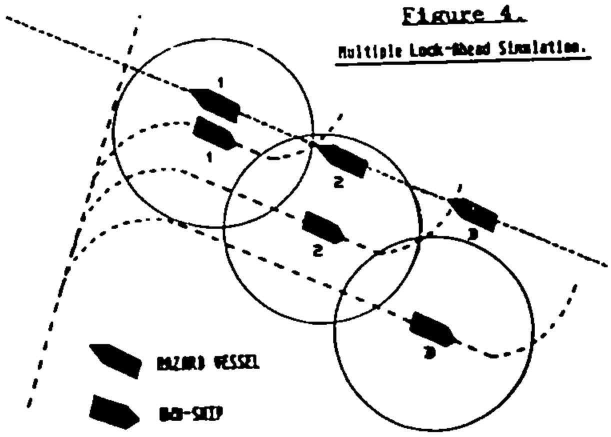

7. EXAMPLE

Figures 5a, 5b, show the screen displays for own-ship and hazard ship respectively after an overtaking encounter. Own-ship has drawn parallel with hazard vessel, overtaken on the starboard (right) side, then altered back onto original course well ahead of hazard. A review of the encounter is given on the screen for own-ship's E.S.

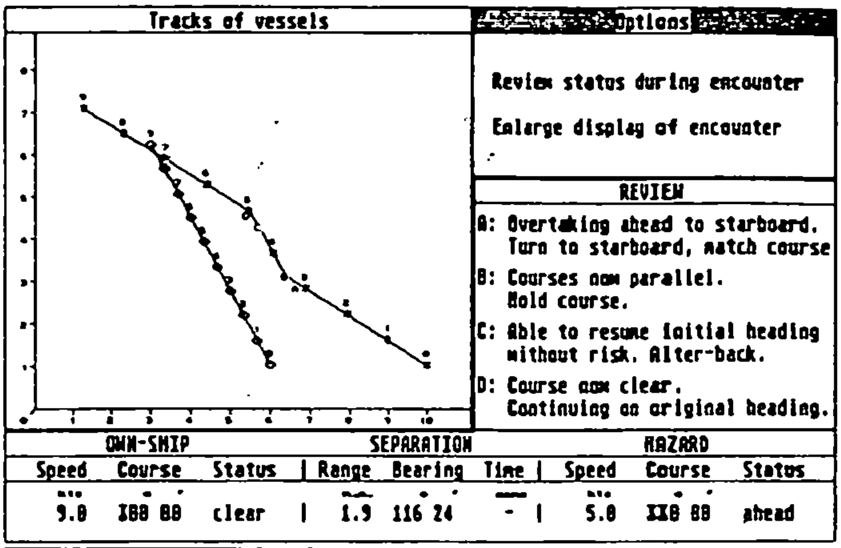

Figure 5a.

Display of overtaking manoeuvre as shown on own-ship's expert system computer. Own-ship is give-way vessel in this encounter.

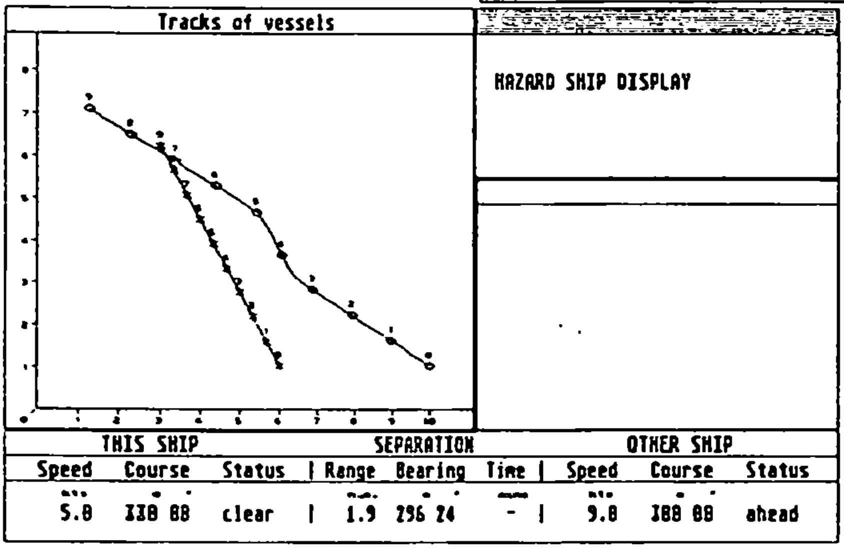

Figure 5b.

Display of manoeuvre shown in figure 5a, as seen on display of hazard ship simulator computer. Hazard ship is stand-on vessel.

8. CONCLUSIONS

Multiple simulation runs have shown the expert system effective in dealing with simple 2-ship encounters; the structure of the system allows expansion of the rule base to handle complex encounter situations. Look-ahead simulation, as a decision tool within the expert system, ensures that action is taken at an appropriate time to guarantee a successful manoeuvre without risk of domain infringement - this topic is considered more fully elsewhere [6].

The combination of expert systems technology with established simulation techniques has led to a need for a new-style approach. The use of a number of computers, with interchange of information at an appropriate level, would appear to offer interesting possibilities in the field of intelligent simulation processes. Specifically in respect of a test vehicle for an intelligent collision avoidance system, this approach has provided an effective model of real-world circumstances, giving a sound basis for continued system development.

9. REFERENCES

1. International Marine Organisation, 1981. "International Regulations for Preventing Collisions at Sea".

2. Goodwin, E.M., 1975. A Statistical Study of Ship Domains. Journal of Navigation. Volume 28, pp328-344.

3. Davis, P.V., Dove, M.J. & Stockel, C.T., 1980. A Computer Simulation of Marine Traffic Using Domains and Arenas. Journal of Navigation. Volume 33, pp215-222.

4. Colley, B.A., Curtis, R.G. & Stockel, C.T., 1983. Manoeuvring Times, Domains & Arenas. Journal of Navigation. Vol. 36, pp324-328.

5. Colley, B.A., 1985. "Computer Simulation of Marine Traffic Systems". PhD Thesis.

6. Blackwell, G.K., Colley, B.A. & Stockel, C.T., 1989. Look-Ahead Simulation as a Tool for Real-Time Expert Systems. Simulation Multiconference, Society for Computer Simulation, Rome (Italy), June 1989.

MODELLING AND CONTROL OF MARINE CRAFT

Edited by

M. M. A. POURZANJANI

School of Engineering, University of Exeter,
Exeter, UK

and

G. N. ROBERTS

Royal Naval Engineering College,
Manadon, Plymouth, UK

ELSEVIER APPLIED SCIENCE
LONDON and NEW YORK

Onboard Systems for Automatic Ship Guidance and Intelligent Collision Avoidance

G K Blackwell and R S Burns
Polytechnic South West, Plymouth, U.K.

Abstract

A potential encounter situation between two or more vessels requires an intelligent, reasoned response from the master of each of those vessels. A computer-based expert system for marine collision avoidance should also be capable of emulating all of these functions. Notably, the decision-taking function requires considerably more than a simple mechanistic application of the anti-collision regulations.

This paper describes the development of such a system, supported by a simulation environment incorporating two ancillary computers to simulate own-ship and hazard vessel. To date, a very simplistic ship model has been used, comprising simple constant velocity motion plus arcs of circles at constant speed for turning curves.

Multivariable system theory is employed to describe the dynamic characteristics of surface vessels to provide a more elegant, and accurate ship model. Optimal control theory is employed to provide a completely automatic guidance system.

The possibility of creating an intelligent guidance system is considered, whose normal task is to lock a vessel onto a predefined track at a given speed and heading. If an encounter situation with another vessel arises, the collision avoidance module will flag the control system as to the necessary course of action to take so that the international regulations are not breached.

1. INTRODUCTION

The master of a modern vessel, standing on the bridge of his ship, receives a wide range of information from a variety of sources; instruments connected to relevant equipment give him ship's rudder setting and engine rpm; appropriate sensors give information on speed through the water, pitch, roll and yaw; navigational aids such as Decca, Omega, Loran-C and satellite navigation systems provide accurate data (typically within a few metres) for fixing position and ground speed (i.e. true speed, relative to the sea-bed); detector systems such as radar and sonar warn of obstacles above or below the surface, giving range and direction - an enhancement of radar, the ARPA (Automatic Radar Plotting Aid) system, even gives the (relative) speed and course of any moving hazard in the vicinity, such as another ship. Additionally, of course, he has his own eyes and ears, and those of his crew, to fill any gaps left by those electronic aids.

The system under development is intended ultimately to be yet another electronic aid on the bridge, one which will receive data from a variety of these sensors (plus chart data from on-board 'electronic charts') and give a reasoned response to prevailing circumstances. Clearly 'eyes and ears' information may be input via a keyboard, when necessary. Rather than proliferating the mass of data to be assimilated by the ship's master, the purpose of this device is to collate relevant data in the field of collision avoidance. Reasoned conclusions and recommendations may then be presented in an easily-digested form, to hopefully reduce the scope for error by omission, misreading, or simply 'information overload'.

The development environment consists of two simulated ships in open-sea conditions. One, termed 'own-ship', is under the guidance of the expert system described here; the other, the 'hazard vessel', is also guided by machine intelligence, with an option to manually override such control. The rule base for the expert system is based on the International Regulations for Preventing Collisions at Sea (1), supplemented by substantial research into good practice, as observed by experienced mariners. Output is presented in the form of a multi-window display, combining graphical, numerical and textual information - see annotated example in Figure 1.

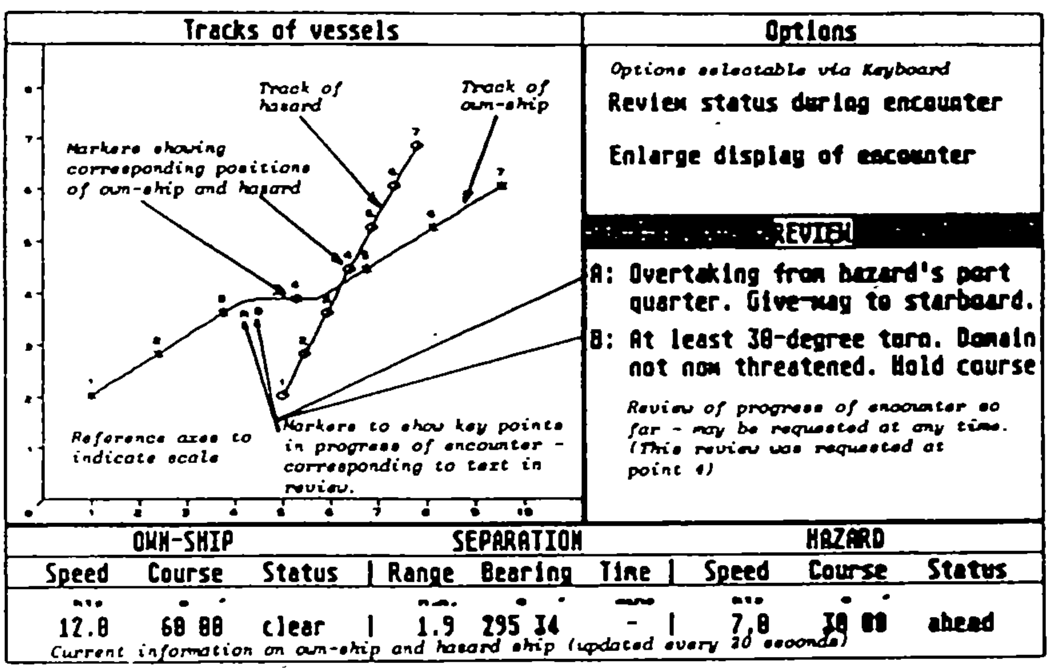

Figure 1. Annotated Example of Screen Display

The expert adviser is currently being transferred to an Acorn Archimedes RISC (Reduced Instruction Set Computer) System, a multitasking WIMPS (Windows-

Icons-Mouse-Printer System) environment, with capability for substantial sensory inputs and control outputs. The combination of processing power, communication interfacing capability, presentation style and user interaction well suit this system for such a task.

2. PRINCIPLES OF OPERATION

At the present phase of development, the rule base consists of rules for:

a) Identifying the presence of a potential hazard, and assessing the threat in terms of expected time to infringement of the domain - this time factor is central to the decision on when to take avoiding action, if necessary;

b) Identifying the type of encounter, including several variants in some cases (the primary types of encounter being head-on, crossing and overtaking), and fixing the status of own-ship and the perceived status of hazard-ship in the encounter (head-on or give-way);

c) Negotiating the stages of the encounter, with appropriate safety margins.

The rules identifying and handling the three types of encounter are as follows:

1) Head-on - Two vessels are approaching each other on near-reciprocal bearings (i.e. $\pm 6^o$, in this system). A vessel in such an encounter must consider itself give-way and alter course to starboard (right); it should not begin to come back onto course until the other vessel has passed 90 degrees on the port beam (left);

2) Crossing - Two vessels are approaching each other at an angle such that each can see one side of the other; more specifically, one vessel is in the sector subtended by the port or starboard light -112.5^o measured from the bow (front) on each side. The ship with the other vessel to starboard should give way, altering course to starboard and going behind. It should not alter back until the stand-on vessel has passed 15^o to port of the own-ship's original heading;

3) Overtaking - A ship is deemed to be overtaking if it is approaching from within the stern-light sector, i.E. The rear 135^o arc. It should give way as follows:

   a) If on near-parallel course (within 10^o, in this system), go behind to starboard; alter back when safe to do so;

   b) If on port quarter of stand-on vessel, action as for (a);

   c) If on starboard quarter, and set to overtake ahead on current course, AND if able to parallel courses without infringing domain, parallel-up; when well past and clear to pull across, return to original heading;

d) Otherwise, go behind to port, alter-back when safe to do so.

The inference engine for this system rests primarily on previous research in the field of marine collision avoidance. In particular, two well-tried concepts dictate the nature and extent of action taken: the domain, or 'danger zone' around a ship, into which no other vessel or obstacle may be allowed to encroach; and the arena, a large zone around a vessel, inside which any other vessel or obstacle must be considered a potential hazard.

The concept of the domain was advanced by Goodwin (2) as 'the effective area around a ship which a navigator would like to keep clear with respect to other ships and static objects'. The expert system has as its first objective the protection of the integrity of the domain. The arena concept was introduced by Davis (3) in response to the observed action of mariners manoeuvring to prevent infringement of their domain by ships in their vicinity. This concept was refined by Colley (4) to a time factor, referred to as RDRR:- Range to Domain/Range-Rate - basically, the time-to-go (before calculated domain infringement) at which the mariner will assess the staus of his vessel and take appropriate action. For example, on the basis of a ship's handling characteristics, a mariner may choose to take action 15 minutes before potential domain infringement - a 15 minute RDRR.

Colley's RDRR is used to define minimum action time; the look-ahead facility described in section 4 then modifies that RDRR value in the light of current circumstances - guaranteeing safe, but not excessive, manoeuvring time. This Predetermined Safe Manoeuvring Time (PSMT) replaces the previous 'best guess' with a certainty (barring untoward incidents - section 4 considers this point).

Another basic principle of collision avoidance manoeuvres is that not only must they be carried out, they must also be seen to be carried out. No other form of communication between vessels can be assumed, and so any avoidance action must be clearly identifiable as such from another vessel. Such manoeuvres, under the direction of the expert system, must be of sufficient magnitude (a 'clear turn') and in good time, to indicate to the other vessel that the encounter has been anticipated and is being dealt with. Any decision to use the processing capabilities of the computer to 'optimise' course changes purely on the basis of 'safe clearance' criteria must be moderated by such considerations.

The current model uses a minimum turn of 30°, a value drawn from previous research by Colley (5). The minimum RDRR is defined as a parameter for each vessel - typically 10 minutes, but varied according to the handling characteristics of the vessel being simulated.

3. INFERENCE ENGINE AND RULE STRUCTURE

The function of the inference engine at each time-step is to:

a) Receive the parameters from the simulator units (speed, heading, position) and from them generate a set of secondary parameters - relative bearing, relative velocity components, time to projected domain infringement, etc.;

b) Apply the appropriate rule to these parameters - depending on status of vessel as perceived at the previous time-step - to ascertain whether the status of the vessel has changed, and if so, what the new status should be. Each rule involves a test on these system values, and may involve evaluation of further mathematical functions of these values - such functions form part of that rule.

c) If the status of the vessel has changed, flags are set for ship simulator action; also, a change of status may lead to further rules being invoked immediately; the inference engine will continue to process successive rules until a 'defer' flag is reached, indicating that no further action on the rule base is required until the next time-step.

The rule base is a hierarchical structure, based on the structure of a binary tree. However, the left and right links from each node in this 'tree' structure are rather more flexible, in that two or more links may lead to a common node at a lower level, or a link may loop back to a node at a higher level. A schematic of this structure, showing most of the initial rule base, is given in Figure 2.

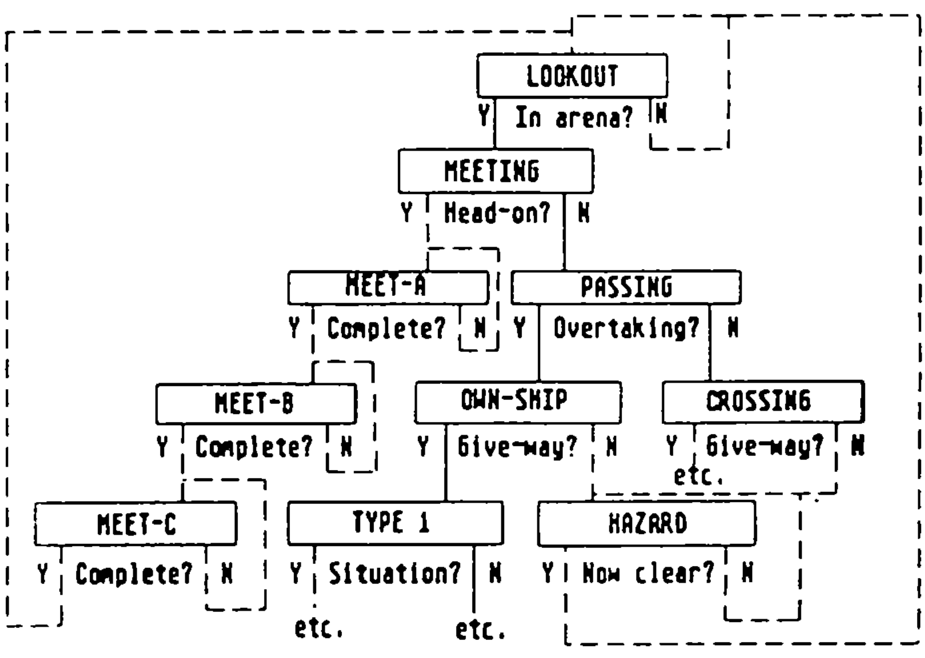

Figure 2. Rule Structure

The rule base is also in effect a state table, each node representing a possible state of own-ship - no encounter currently in progress, second stage of a crossing encounter, etc. Each node contains:

1) Left and right links to other nodes in the rule structure (or looping back to the

same node) - one of these links will be followed at every decision step;

2) A mathematically-based decision function, used to determine which link is taken at this stage;

3) A flag for each link, indicating in each case whether the next rule is to be processed immediately, or deferred until the next time-step;

4) Ship status flags, to pass information to simulator and display;

Enhancement of the rule base involves creation of the appropriate new node(s), and resetting of the relevant links so as to position the node(s) correctly within the rule structure.

The rules embody the International Regulations for Preventing Collisions at Sea, plus clarification of specific practical details as noted by Colley and used most effectively in earlier simulation. Using these rules, a collision avoidance manoeuvre is taken through five stages, as follows:

1) 20 minutes before projected domain infringement, look-ahead simulation is used to identify a suitable PSMT, starting from the minimum RDRR.

2) When the PSMT is reached, the give-way vessel initiates a turning action;

3) Once the required turn is effected (minimum 30°), the vessel proceeds on a straight course;

4) When the original heading may be resumed without further risk of domain infringement (and any other conditions are met), alter-back is initiated necessitating a turn in the opposite direction;

5) When back on the original heading, vessel proceeds on that heading; manoeuvre over.

Each of the rules involves a yes-no decision function, using inequalities based on current values of displacement and velocity vectors. The majority of processing is thus concerned with solution of triangles, and intersections of straight lines with circles - substantial evaluation of trigonometrical functions. It is anticipated that this aspect of the E.S. Will escalate as more complex situations are taken into account; this has been a major factor in determining choice of language and form of rule structure for this time-critical system.

4. LOOK-AHEAD SIMULATION

A characteristic of expert systems is that intuitive reasoning by the human expert must be emulated in a computer system by an explicit logical sequence of cause-and-effect. Such reasoning may be difficult to identify, let alone rationalise and then emulate. A classic example is the decision on when to evaluate an encounter and initiate collision avoidance action. Extensive observation would suggest that a competent mariner has in mind a fairly well-defined time-to-go before domain infringement, at which point in time the nature of the encounter is identified and appropriate avoidance action decided upon; too soon, and the course alteration will be excessive; too late, and the domain infringement may not be avoided.

However, results from the multiple encounter simulation above show that such a fixed 'time-to-go' criterion is not adequate for all possible encounter situations. It seems most probably that the expert mariner would readily identify those cases which require earlier-than-usual action; an expert system, of the type under consideration here, should likewise be capable of doing so. This has prompted the inclusion of just such a facility, one which will identify the optimum decision time for any particular scenario. Such a dynamic time-to-go is pinpointed by consideration of the eventual consequences; this means inclusion of a look-ahead facility as an integral part of the decision-making process.

20 minutes before projected domain infringement, the Expert System simulates a complete encounter as it would proceed if action were to be taken at, say, 10 minutes before domain infringement. (this minimum figure is variable, being input as an initial parameter). If this time interval proves inadequate, the system will run another simulation, this time for a 12-minute time-to-go. The decision threshold is moved successively back, with a new simulation for each, until a satisfactory time-to-go is identified. The Expert System then allows the encounter to proceed to that PSMT (Predetermined Safe Manoeuvring Time), whereupon the decision on encounter type, and appropriate avoidance action, is taken. See Figure 3.

Figure 3 shows a sequence of three look-ahead simulations of a crossing encounter, with own-ship taking avoidance action. Starting times for successive simulations work backwards from a predefined minimum.

Simulations 1 and 2 both show own-ship's domain (safety zone) being infringed by the hazard vessel. Simulation 3 identifies a safe time-to-go for the avoidance manoeuvre.

The incorporation into the Expert System of such a facility, requiring possibly several complete encounter simulations to achieve its objective, will clearly put pressure on the time constraints of a real-time system. It is essential that the necessary processing be completed well in advance of the critical phase of the encounter. In the present idealised situation - two ships in an open sea - processing time is not a problem. However, it is envisaged that in more complex scenarios (e.g. those involving a number of vessels), such processing may require other decision-making logic to be suspended for several time-steps whilst it is under way.

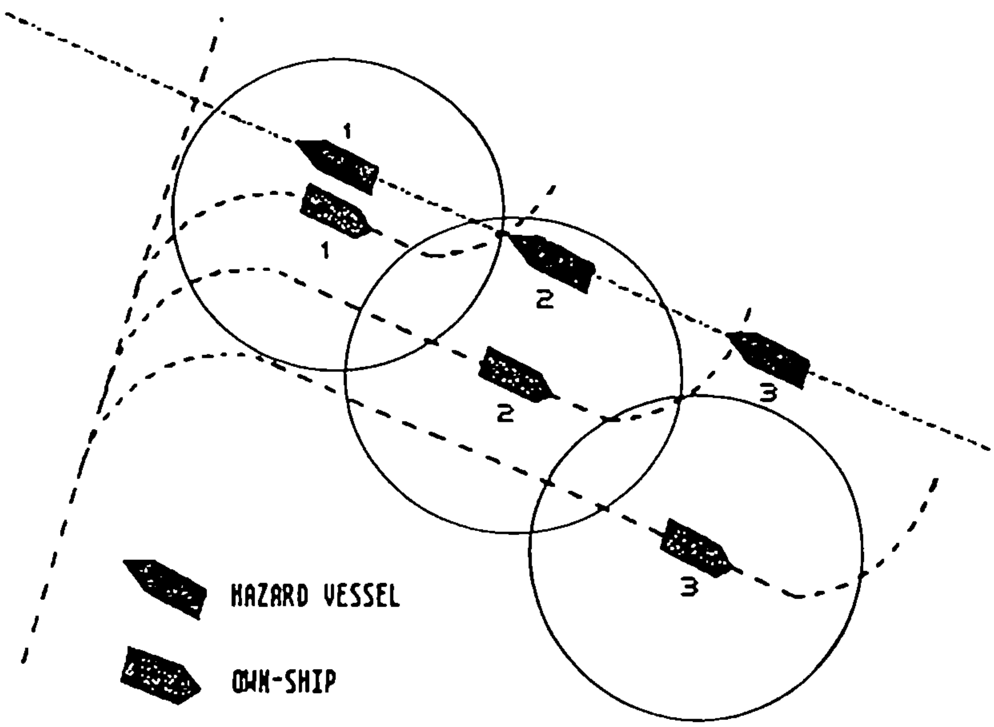

Figure 3. Multiple Look-Ahead System

5. THE DEVELOPMENT ENVIRONMENT

With the introduction of look-ahead simulation into the decision process, processing needs for the E.S. Have expanded, highlighting the need to move towards the final objective - a dedicated computer housing the E.S., acting on data input from other sources. The development environment now in use (6) consists of three linked micro-computers (see Figure 4): one houses the E.S. For own-ship; a second acts as own-ship simulator; the third provides 'radar' data for a hazard ship, driven by a simulator module and simplified E.S. module - with manual override, to simulate a 'rogue' ship.

The new implementation, on the Acorn RISC machine, will allow three modes of operation:

1) Programs simulating own-ship and hazard vessel, running as separate tasks on the same computer, with software communication links representing sensor/control data and radar input (both ways) respectively;

2) Simulator programs on separate processors, as at present, with parallel and serial hardware links to own-ship and hazard, respectively.

3) Real data, input via parallel and serial links from ship's sensors and radar, on board a suitably-fitted research vessel.

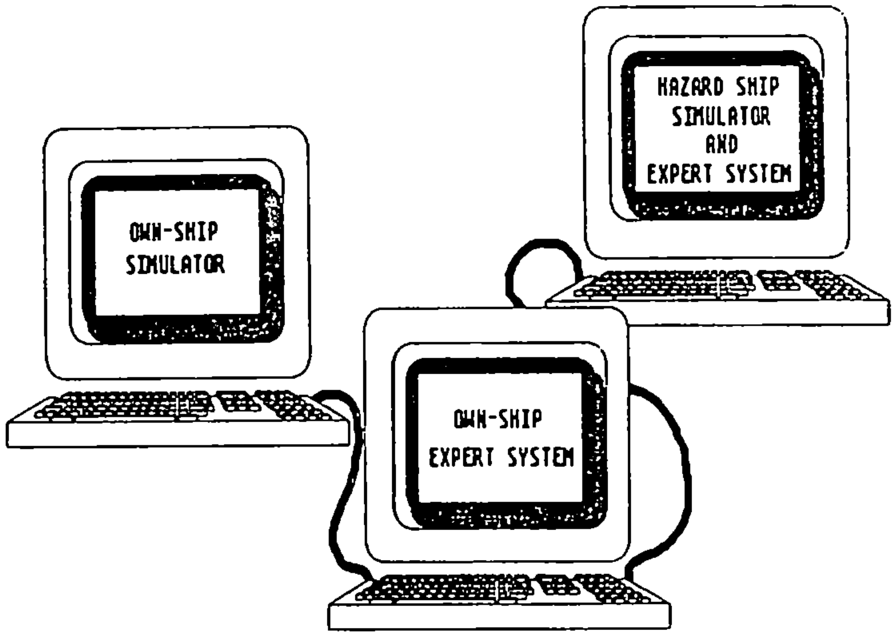

Figure 4.

6. THE SHIP GUIDANCE PROBLEM

Most existing commercial ship autopilots are single input, single output control systems, the output variable usually being heading. Such a system is unsuitable for automatic collision avoidance, where it is essential to know at any instant the position of the vessel, together with its forward velocity and heading.

Multivariable system theory may be used to model such a situation and optimal control theory employed to formulate a control policy.

System Model - It has been shown (7) that the equations of motion for a ship in a manoeuvring situation may be expressed in surge, sway and yaw by the state matrix vector

$$\dot{X}(t) = F(t)X(t) + G_C(t)U(t) + G_D(t)w(t) \tag{1}$$

which has a state vector:

$$X^T = (\delta_A n_A x u y v \psi r) \tag{2}$$

a control vector:

$$U^T = (\delta_D n_D) \tag{3}$$

and a disturbance vector:

$$w^T = (u_c v_c u_a v_a) \tag{4}$$

Guidance System - The essential features of a ship automatic guidance system are shown in Figure 5.

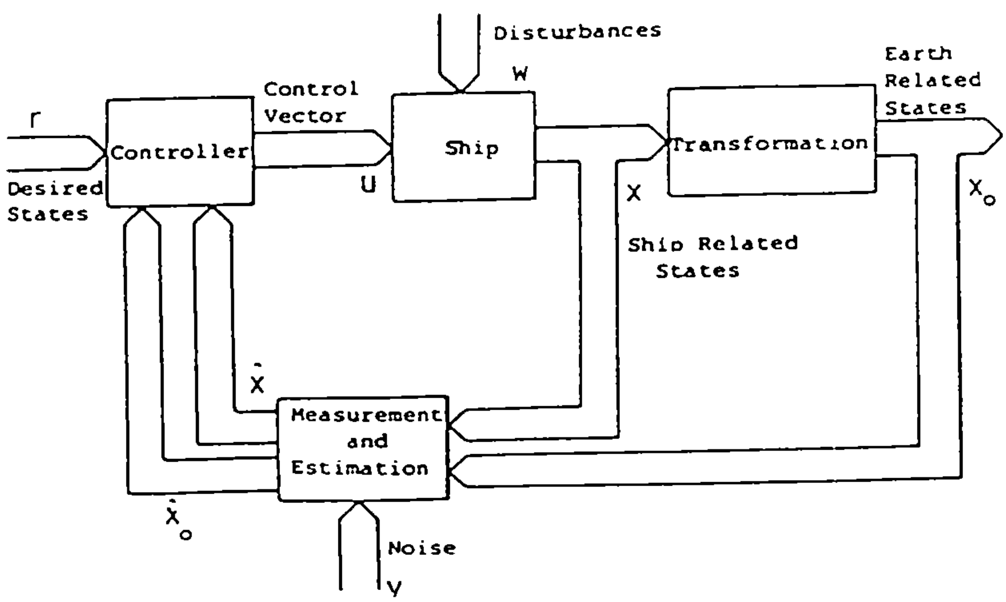

Figure 5. Ship Automatic Guidance System

Deterministic Optimal Control - The quadratic criterion to be minimised is:

$$J = \int_{t_0}^{t_1} \left[(X - r)^T Q (X - r) + U^T R U\right] dt \tag{5}$$

where r is the desired value of the state vector, supplied by the collision avoidance module. It can be shown that constrained functional minimalisation yields the simplified matrix Riccati equations:

$$WF + F^T W + Q - WGR^{-1}G^T W = 0 \tag{6}$$

This provides the optimal control law:

$$U_{opt} = r - R^{-1}G^TWX \tag{7}$$

or,

$$U_{opt} = r - SX \tag{8}$$

7. LINEAR OPTIMAL, CLOSED-LOOP POLE ASSIGNMENT

Consider the time-invarient form of the state equations (1) together with the optimal control law (equation 7). Upon substitution of (7) into (1) for no disturbances:

$$\dot{X} = (F - GR^{-1}G^TW)X + Gr \tag{9}$$

or

$$\dot{X} = (F - GS)X + Gr \tag{10}$$

were the term $(F - GS)$ may be identified as the closed-loop state matrix of the optimal system. The optimal closed-loop eigenvalues (poles) are then given by:

$$|sI - (F - GS)| = 0 \tag{11}$$

It is apparent that when F and G are time-invariant the location of the optimal closed-loop poles depend upon the value of the feedback matrix S, which in turn is dependent upon weighting matrices Q and R.

Using equation (11) Kouvaritakis (8) has demonstrated that the optimal eigenvalues may be placed as shown in Figure 6. The effect of changes in forward velocity are considered elsewhere.(9)

8. COMPUTER SIMULATION OF GUIDANCE SYSTEM

Having established an optimal control law, it is necessary to assess its effectiveness. An evaluation was undertaken by conducting a computer simulation study of the automatic guidance system implemented on the trail ship in the approaches to Plymouth, U.K., using a pre-determined track.

In Figure 7, the vessel is travelling at 7.717 m/s (15 knots) and it can be seen that there is an overshoot at each way-point. This may be corrected for by (a) reverse-time integration or (b) dual-mode operation.

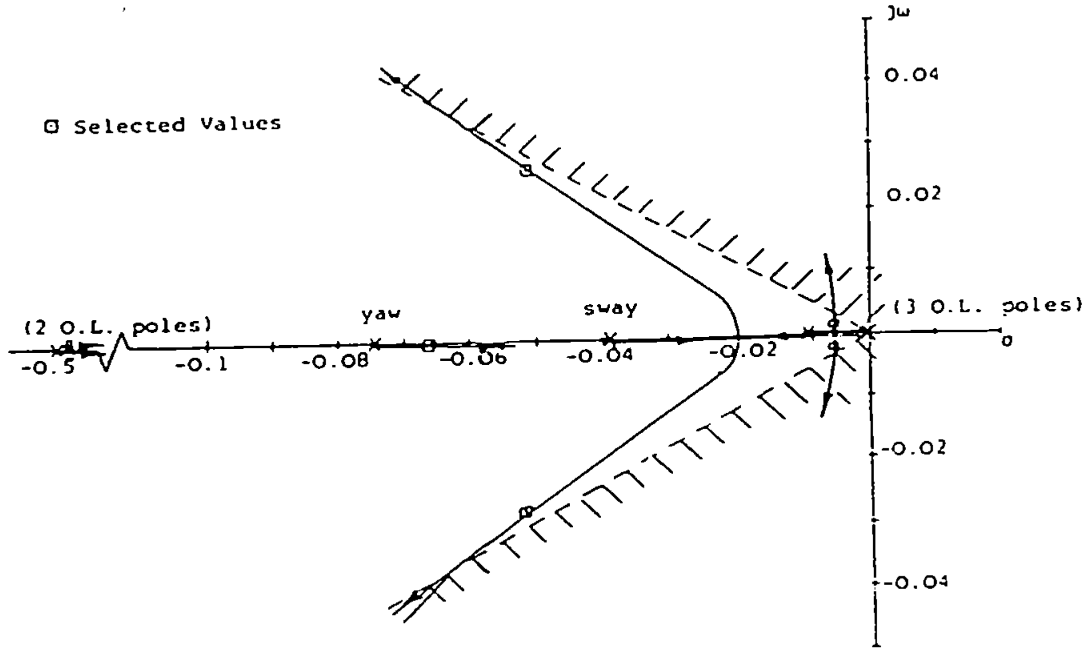

Figure 6. Optimal Eigenvalues

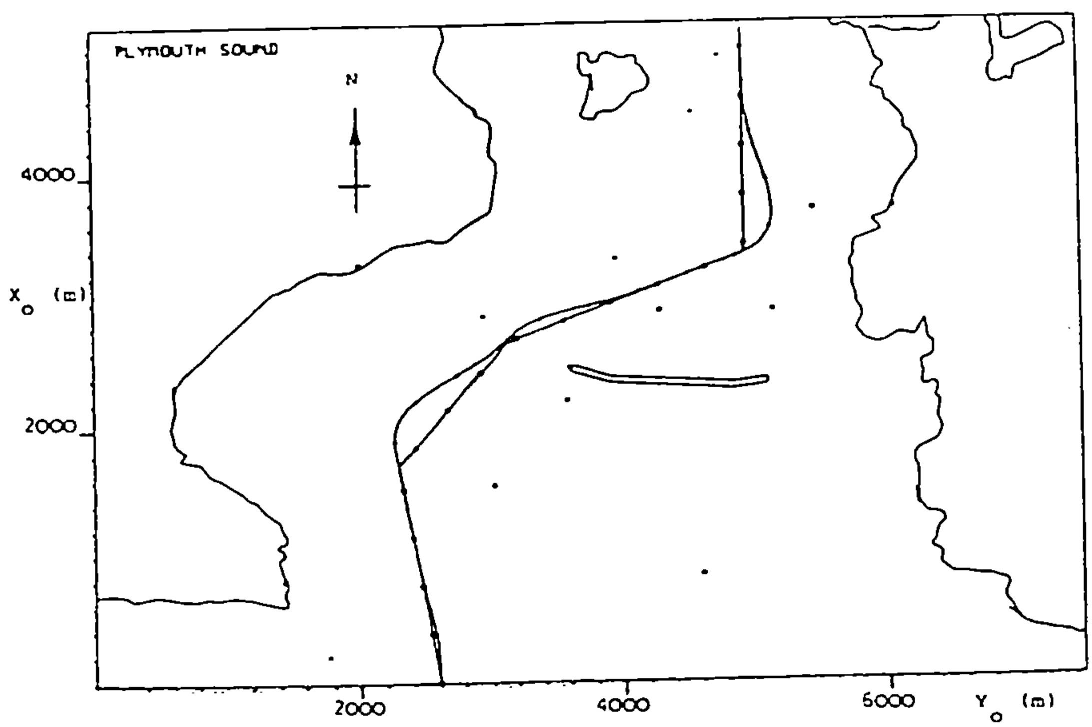

Figure 7. Effect of Transient Overshoot

With the former, a reverse-time command vector is generated from knowledge of the ship's closed-loop dynamics together with the desired state vector. The technique was found to possess good transient anticipation but suffered from a lack of steady-state accuracy, this being sensitive to controller setting.

Dual-mode control, on the other hand, is a simple concept that allows the command vector to be replaced by a method of advancing way-points so that track-changing occurs without overshoot. At the advanced way-point, a switch is made from track to course-keeping. This initiates a turning manoeuvre and when the course error is less than a prescribed amount, track-keeping mode is re-established. Dual-mode control was found to possess excellent track-keeping capability as can be seen in Figure 8.

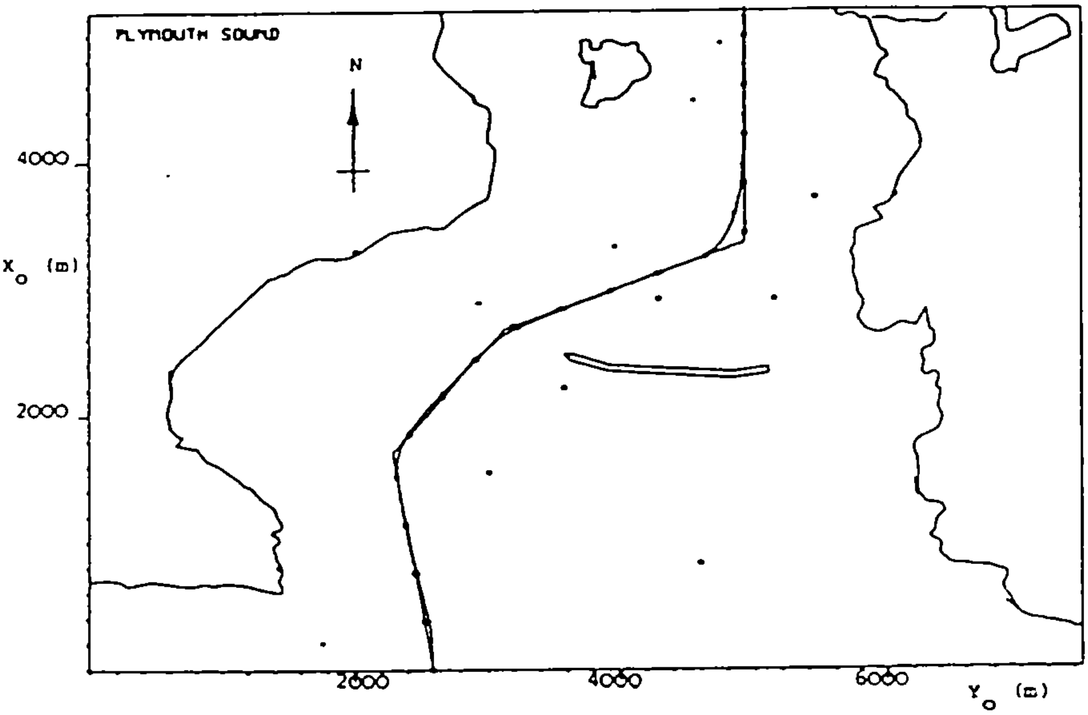

Figure 8. Dual-Mode Control with Way-Point Anticipation

9. CONCLUSIONS

The concept of a completely automatic system that:

a) makes decisions on whether an avoidance manoeuvre is necessary;

b) decides on what the correct avoidance measure is;

c) Implements the manoeuvre;

would be viewed with suspicion by most ship's masters. A more acceptable solution is one whereby the sensory system is used in an advisory role, with the automatic system as a backup. This is the direction taken by the aerospace industries, where piloted landings are the norm but automatic landings are employed under bad weather conditions. It is still an undisputed fact however, that most collisions at sea are still the result of humna error, and any system, whether advisory or automatic that can reduce the probability of such events occurring is worthy of investigation.

NOMENCLATURE

Matrices and Vectors

F	Continuous Time System Matrix
G, G_C, G_D	Continuous Time Forcing Matrices
I	Identity Matrix
Q	State Error Weighting Matrix
R	Control Weighting Matrix
r	Desired State Matrix
S	Feedback Gain Matrix
U	System Input Vector
W	Riccati Matrix
w	Disturbance Vector
X, X_0	Ship and Earth Related State Vectors

Scalar Symbols

K^z	Integer Counter
n_A, n_D,	Actual and Demanded Engine Speeds
r	Yaw Rate
s	Laplace Operator
t	Continous Time
T	Sampling Time Interval
u, u_c, u_a	Forward Velocity of Ship, Current and Air
v, v_c, v_a	Lateral Velocity of Ship, Current and Air
x, y, z	Ship Related Cartesian C0-ordinate System
δ_A, δ_D	Actual and demanded Rudder Angles
ψ	Heading of Ship

References

1. International Marine Organisation, 1981. International Regulations for Preventing Collisions at Sea.

2. Goodwin, E. M.,"A Statistical Study of Ship Domains", Journal of Navigation. Vol. 28, pp 328-344.

3. Davis, P. V., Dove, M. J. And Stockel, C. T.,"A Computer Simulation of Marine Traffic Using Domains and Arenas", Journal of Navigation. Vol. 33, pp 215-222.

4. Colley, B. A., Curtis, R. G. And Stockel, C. T.,"Manoeuvring Times, Domains and Arenas", Journal of Navigation. Vol. 36, pp 324-328.

5. Colley, B. A.,"Computer Simulation of Marine Traffic Systems", Ph.D Thesis.

6. Blackwell, G. K. And Stockel, C. T. "Simulation of Ship Encounters using Multiple

Linked Intelligent Systems", Proc. Of Summer Computer Simulation Conference, SCS International, Austin (Texas), July, 1989.

7. Burns, R. S., Dove, M. J., Bouncer, T. H. And Stockel, C.T.,"A Discrete Time Varying, Non-Linear Mathematical Model for the Simulation of Ship Manoeuvres", First Intercontinental Maritime Simulation Symposium and Mathematical Modelling Workshop, Munich, w. Germany, June 1985.

8. Kouvaritakis, B.," The Optimal Root Loci of Linear Multivariable Systems", Int. J. Control, 1978. Vol. 28, pp 33-62.

9. Burns, R. S. and Dove, M. J. "An Automatic Guidance System for Large Ships", International Confernece on Computer Aided Design, Manufactur and Operation in the Maritime and Offshore Industries. Computational Mechanics Institute, Washington D.C., U.S.A., Sept. 1986.

Modelling and Simulation

Proceedings of the
1990 European Simulation Multiconference

June 10-13, 1990
Atrium Hotel
Nuremberg
Federal Republic of Germany

Edited by
Bernd Schmidt

Organized by
The Society for Computer Simulation International

In Cooperation with
EUROSIM, JSST and CSSC

ISBN 0-911801-73-1

Printed in the United States of America

OPTIMAL STRATEGIES FOR A MARINE REAL-TIME EXPERT SYSTEM

G.K. Blackwell and C.T. Stockel
Department of Computing, Polytechnic South West
Plymouth, Devon, England

1. FITTING-OUT

The quality and applicability of advice given by an intelligent knowledge-based system (IKBS) is bounded by a number of constraints: first, any limitations in the scope and flexibility of the IKBS to non-standard situations; second, certain relevant information may be unavailable to the knowledge base; third, variations in the quality (accuracy) of information passed to the system may affect the reliability of the response given by such an expert advisor.

This last point is particularly pertinent in the case of a real-time IKBS with sensory input; any sensor will be subject to a degree of error (e.g. temperature readings to plus or minus 2 per cent), and an IKBS should be 'aware' of such limitations in order to include appropriate safety margins in its reasoning and its conclusions. A sensor-based system may also be prone to certain 'blind spots' - any such likelihood of partial information should also feature in the rule base.

A limitation which is unique to real-time IKBS is that of timeliness: the value of a decision or recommendation is dependent on its being available in time for it to affect the outcome of the process being monitored. Depending on the time-frame available for evaluation of the relevant facts, and the level of complexity of such evaluation, this constraint may apply pressure on the expert system in two ways: the design and functioning of the inference engine (i.e. its handling of the rule base) must be so streamlined as to fire the appropriate sequence of rules as time-efficiently as possible; and the rule structure itself must provide for optimal analysis of a situation - 'optimal' in the sense of speed, concomitant with full evaluation of all relevant information.

By the nature of such a system, all analysis and planning is forward-looking in time; any proposed strategy must take into account the possible future consequences of such action. Relating this to games theory, each possible 'move' must be considered in the light of the likely future 'game state', or even alternative possible game states. The scope for such alternatives expands rapidly as extrapolation into future time-frames is extended; this look-ahead process is one of the major factors which must be optimised, without loss of significant definition. A previous paper presented by the authors [1] dealt with the subject of look-ahead simulation for real-time expert systems.

Continuing the analogy: the game also becomes more complex, and the number of options grows rapidly, as the number of pieces on the board increases. This expansion is certainly not linear; in general, it follows a factorial, or even exponential, growth pattern. A system which is fairly simple to analyse when limited to two entities may require substantially more processing when further entities are taken into consideration. The rule base must be structured to accomodate such extensions as efficiently as possible, without loss of significant detail.

An optimal strategy will in general be derived by assessing the alternatives (i.e. simulating the effects of each successively along the time-line), and selecting the 'best' - on the basis of a number of criteria, appropriately weighted. Some strategies may disqualify themselves by leading to an unacceptable outcome. A very high degree of parallelism is evident in this task, indicating that substantial benefits may be gained by use of parallel processing techniques.

The authors are concerned with development of on-board IKBS for marine collision avoidance, linked to radar and similar sensory input systems, as well as providing keyboard/VDU communication for the user. Previous work was based on two-ship encounters in the open sea; current developments centre on extension of the rule base to handle multi-ship encounters. From the outset, the above considerations have been taken into account in both the structure of the inference engine and the underlying principles of the rule base.

The first, fairly simplistic, approach is to identify the hazard vessel posing the greatest immediate threat, deal with that as an isolated unit, then re-evaluate the new situation in a similar manner. Some multi-ship encounters could yield to this strategy; however, there is the possibility that avoidance of Hazard 1 could result in an increase in risk from Hazard 2, necessitating even greater or more urgent manoeuvring [2]. There is even the possibility of 'oscillation', with each of the two hazards alternately becoming the more threatening. Full resolution of the multi-ship encounter problem must take account of the total situation - as a competent master mariner would. One strategy, for example, observed as common practice by masters of cross-channel ferries, is an early course change to avoid a group of vessels presenting a potential hazard well ahead - see figure 1. In general, a simple manoeuvre at an early stage is much preferred to a complex 'slalom' course at the latest possible moment. Such considerations must naturally be balanced against deflection from planned course, with consequent loss of time and use of extra fuel.

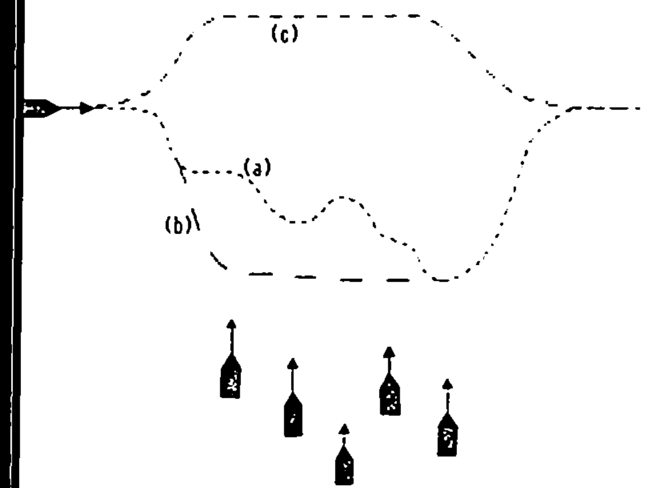

gure 1

lti-ship encounter.

ree possible courses for ownship to avoid
approaching group of hazard vessels:
shows avoidance of each ship in turn;
delays alter-back so as to avoid the group;
is an early course change to pass ahead, minimising track divergence.

PRELIMINARY MANOEUVRES

e rapid advance of microcomputer hardware chnology in recent years has been paralleled major developments in the field of software ols and environments. Less conspicuous, but ually significant, has been the progress on a de range of electronic sensors and actuators; th in terms of versatility and accuracy, and se of hardware and software linking to a nventional microcomputer system, such control vices now form an integral part of the total mputer applications toolkit. The intelligent, er-friendly, real-time advisory/control stem is an idea whose time has come.

o established criteria for conventional nagement information systems are:

) the system should be capable of meeting users on their own ground; such users are unlikely to be computer specialists, and should not be expected to learn complex new computer-oriented skills simply in order to make use of this facility;

) an information system worthy of the name should inform - succinctly, upon request (but with capability for drawing the user's notice to items in urgent need of attention), highlighting pertinent data rather than overwhelming the user with an avalanche of facts which may or may not be relevant.

ch criteria must surely demand far greater phasis in a real-time management advisory stem: time-critical decisions should not be ndered by unfamiliar procedures or irrelevant detail. Of the numerous disasters attributed to human error, two major contributory factors have frequently been (a) lack of familiarity with operating practice, and (b) inability to extract and collate the relevant material from the mass of data available, then formulate and implement an appropriate strategy, in the limited time available. An expert system should emulate its human counterpart, whose expertise lies in an ability to pinpoint the significant, discard the irrelevant, and present well-reasoned conclusions in a form, and on a time-scale, which enables prompt implementation of those conclusions by one relying on that expertise.

Nowhere is this more so than in the maritime setting. The master of a modern ship has access to a wide variety of information: via bridge-mounted electronic instruments (some of them repeaters from other parts of the ship); via his own senses; and relayed from other sources, near and far. Not all of this information is of major significance: he may be in radio contact with his shore base, hundreds of miles away, but not with the master of a vessel a mile away which could constitute a very immediate threat of possible collision. Even visual information, relied upon so heavily at sea, as elsewhere, may give a false impression: misinterpretation of observations of navigational lights has been cited at many a Board of Trade inquiry as the root cause of misjudged actions. Conditions of poor visibility obviously compound such errors, and the rare bonus of radio contact between vessels can play one false; the reassuring voice assumed to emanate from a ship half a mile through the fog has on more than one tragic occasion turned out (after the event) to represent another vessel several miles distant.

In brief, there is ample scope for 'information overload', where less relevant (or even misleading) data distracts attention from vital signs warning of impending danger, and crucial to avoidance of that danger. Small wonder that shipping tragedies (generally tagged 'human error') continue to regularly claim human lives and wreak ecological mayhem - quite apart from the valuable craft and cargo tonnage lost on an almost weekly basis. Perhaps it is worth adding that some of these incidents have occurred in circumstances where pressures were anything but extreme, where disaster seemed so unthinkable that it wasn't even considered - until far too late; here, too, a dispassionate electronic sentinel could redress the balance, pre-empting regretful hindsight with preventive foresight.

An integrated bridge-based vessel management system would be responsible for receiving all electronically-sensed data, such as inputs from navigational and radar systems, and advising on various ship management functions on the basis of that data. Essential data not available via on-board sensors could be input through the keyboard or other suitable manual input device. Output may be displayed on a single screen, using a multiple windowing environment, or a network of terminals, some possibly in other parts of the ship (eg. the engine room); specialist display units may also be required.

It is anticipated that a WIMPS (Windows-Icons-Mouse-Pointer System) 'desktop' environment would provide a suitable user interface. Pop-up or drop-down menus would provide easy access to system facilities as required, and adjustable text and graphics windows could selectively display information relevant to the tasks in hand; a multi-tasking operating system would oversee concurrent processing of such tasks. Ultimately, a distributed WIMPS environment, such as X-Window ((c)M.I.T.), might give remote access to some of these management functions, in addition to other facilities.

3. UNDER WAY

Intelligent collision avoidance decision-making, whether by electronic or human agency, presupposes comprehensive knowledge of certain fixed, or relatively invariant, characteristics of the vessel being directed ('own-ship'): engine performance, rudder response, overall length, beam and draught, stopping distance at various speeds, etc - such data may be termed static information. In addition, the current state of the ship's affective environment must be known - own speed, heading, rudder setting and associated status information, plus data on proximity, velocity and nature (such as can be ascertained) of any potential hazards in the vicinity; this continually variable, or dynamic information, may be collected via sensors, radar and related electronic instrumentation.

The decision process involves submitting this compound of static and dynamic information for evaluation by a system of rules, which should identify and resolve any potential conflict situations. This process must be completed within a very limited time-frame, since:

(a) a new scenario will present itself for evaluation almost immediately (in practice re-evaluation two or three times per minute is adequate for such a purpose);
(b) to be effective, a decision must be based on current data, as otherwise it works from a situation which no longer exists;
(c) a solution to an impending disaster, unless offered almost immediately, is no solution at all, merely a post-mortem!

An initial framework for the rule base already exists, in the International Regulations for Preventing Collisions at Sea [3]. By themselves, however, the regulations are inadequate for such a task, since they leave the onus for certain major decisions with the mariner. Such phrases as 'in good time', 'a clear turn', and 'a safe distance' are open to a wide variety of individual interpretations which may themselves vary with different contexts; sea state, nature of vessel and cargo, possible restrictions (such as coastal features or shipping lanes), traffic density- all affect decisions on collision avoidance manoeuvres. As with any competent human agent, therefore, an electronic adviser would have to be conversant with that broader spectrum of 'rules' (i.e. situational responses) which derive from applied common-sense and years of experience; such experience, gleaned from human experts, must be formalised into an inference structure which gives substance to the bare bones of the official regulations. To be fully effective, the system should be unobtrusive, providing information only as requested, in a meaningful way - but with the capability for drawing attention to matters in need of immediate action.

Clearly, an electronic perception of the environment will differ in certain respects from that arrived at through human senses, and electronic appraisal of a given scenario may necessarily take a different path from a human line of thought. For example, a human observer may spot almost immediately that another ship is turning, by its change in visual aspect; a computer would need several successive ARPA readings before reaching the same conclusion. Conversely, the human lookout would need to observe another vessel for some time before deducing that it was on a constant bearing, and thus on a collision course - a fact which would be immediately apparent to the computer, giving even time and position of expected collision by extrapolation from velocity and displacement vectors. It follows that the requirement is for a system which, in using its processing capabilities to the optimum, parallels human reasoning, without necessarily duplicating it in every detail. The final outcome should nevertheless correspond to the human expert's decision process, combining observed good practice with self-evident commonsense.

4. INTO ACTION

A customised expert system shell has been developed initially on an Atari microcomputer, and recently transferred to an Acorn Archimedes RISC (Reduced Instruction Set Computer) system. After extensive consideration of suitable programming languages for such a task, the 'C' language was chosen for a number of reasons:

(1) the constraints of real-time processing dictate the need for optimum efficiency, in terms of speed, compatible with the complexities of program structure required for such a rule-based system;

(2) many of the rules are expressed in terms of trigonometrical considerations, giving a substantial bias towards mathematical calculation, and thus favouring a language which deals efficiently with such needs;

(3) C is the native language of most WIMPS environments, giving access to facilities at all levels;

(4) entities of various types - rules, ships, encounters (not to mention windows, icons etc. within the WIMP system) - may all be handled very effectively as 'structures' in C; this versatile C data type meets all perceived needs (particularly under the newly-defined ANSI standard) without the apparent overheads of object-oriented or declarative languages.

Both computers were chosen on the basis of processing power at relatively low cost, good WIMPS and graphics facilities, and capability for a variety of input/output options, both digital and analog. A decision to transfer to the Acorn RISC machine was based on the need for compatibility with other ship management functions currently under development by the Ship Control Group at Plymouth, as well as its superiority on a number of fronts. The RISC OS multitasking environment, coupled to versatile I/O handling, are well suited to the planned integrated system.

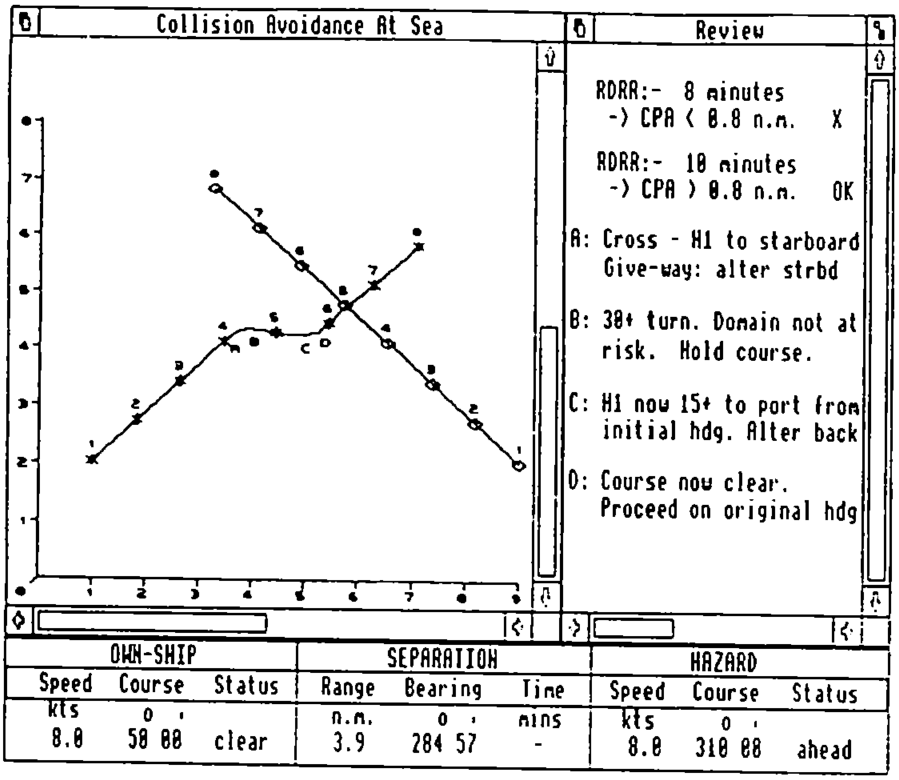

Figure 2

Multiple-window screen display showing scroll bars and window control icons. Bottom 'situation' window may optionally be hidden; 'review' window may be extended or scrolled to give more detailed information; 'tracks' window may give (scrollable) large scale view, or complete small scale view of current situation.

The screen display is in the form of three windows, as shown in figure 2. A square graphics window, filling approximately half of the screen area, displays the tracks and positions of one's own vessel (own-ship) and any other vessels in the vicinity, with markers to show positions at corresponding times; own-ship is shown as a green track, others in various colours - at the current stage of development, only one other vessel (the most threatening) will be 'processed' by the rule structure as a potential hazard; the graphics window display may be enlarged by reference to a pop-up menu (invoked and actioned by mouse buttons), and any part of the scene studied in detail by scrolling the window horizontally or vertically. A second window optionally displays a review of the progress of the situation to date, including the decisions made by the expert system at various stages; this window may also be scrolled, or expanded over the graphics window temporarily, to show further detail of the stages of an encounter. A third window may be selected and de-selected as required, again by a mouse-operated pop-up menu; this window, appearing below the other two, displays data on the course, speed and current manoeuvreing (if any) of own-ship, plus such data as is available (via radar, etc.) on any other vessel in the vicinity considered to be a hazard - future developments will allow selection from the graphics window of which potential hazard is to be 'shadowed' by this window at any time.

Early development comprised the expert system logic, plus ship simulator modules for own-ship and a hazard vessel, all housed in one computer; a later version made use of three microcomputers, one each for expert system, own-ship simulator, and hazard ship simulator.

Having thus verified the integrity of the expert system, the current implementation is designed to operate in any of three modes, two being simulator trials, the third real data:

(1) a semi-intelligent hazard vessel, using a simplified expert system for control, and a dumb own-ship controlled by the current system, are both simulated as additional tasks under the multitasking RISC_OS, on the same computer as the expert system; software communication channels between the tasks simulate radar data to and from hazard vessel, and sensor/control signals between own-ship and expert system;

(2) hazard and own-ship are simulated as above, but in separate computers from the expert system; communications are via serial and parallel links respectively;

(3) inputs via serial and parallel ports, as in (2), are from genuine radar and sensors aboard own-ship, i.e. a research vessel with the computer system on board; such input data has as yet been used only to test the efficacy of this form of input, and observe the response of the system; however, it is envisaged that the control outputs from the system will in due course be connected to the automatic controls already fitted to the vessel (previously used with marked success in computerised ship guidance [4]).

Simulation of own-ship and hazard for (1) and (2) are based on the linear ship model [5]. A more accurate model of own-ship is required in the decision logic of the expert system itself for reference in the look-ahead module outlined in section 5. The non-linear modular model derived by Tapp [6] is used for this purpose.

5. RULES OF ENGAGEMENT

Whilst the prime objective of any collision avoidance protocol is to avoid collisions, this must be done in such a way as to satisfy an important secondary objective: to inform other mariners, through positive action, that the possibility of such an incident has been recognised and is being dealt with. Both of these aims are served by manoeuvres which seek to maintain a zone of clear water around one's vessel, generally referred to as the domain. This observed practice of experienced mariners [7] has been taken as the prime objective of the expert system; a circle, of radius suited to the vessel and circumstances, centred on ownship, has proved an adequate representation of the domain to date.

Such considerations give rise to a broader area of concern, encompassing those vessels or other hazards which are to be assessed as potentially threatening, leading to appropriate action as necessary. This 'area', being ruled by closing speed of ownship relative to any given hazard, is somewhat elastic in its physical dimensions; the crucial consideration is expected time to domain infringement, or RDRR (Range-to-Domain/Range-Rate). A ship's master applying a 12-minute RDRR criterion, for example, (based on handling characteristics for his vessel) would evaluate a potential encounter 12 minutes prior to domain infringement, and initiate avoidance action at that time.

A fixed RDRR for every encounter has been found to be inappropriate as a general strategy; some manoeuvres may be left too late to clear a hazard, others may make an excessive detour through altering course unduly early. Figure 3 illustrates the point: in both cases, the relative velocity of ownship is at 90 degrees to the track of the hazard vessel; in both cases the required avoidance manoeuvre is identical; however, in case (b), this manoeuvre must commence approximately twice as long before anticipated domain infringement as in case (a) - i.e. the ideal RDRR for case (b) is double that for case (a). Paradoxically, the passing distance without any course change would be considerably smaller (0.2 n.m.) for (a) than for (b) (0.6 n.m.); anticipated closest point of approach (CPA) is clearly not a simple guide to required manoeuvring time.

A flexible time-constraint, matching decision time to interval needed for a safe manoeuvre, ensures optimum manoeuvring time: adequate, but not excessive. Predetermined Safe Manoeuvring Time (PSMT) is found by simulating a projected encounter well in advance, for increasing RDRR values (starting from some preset minimum, say 10 minutes), until safe clearance is achieved. The simulation exercise is an integral element of the decision logic, and corresponds to a human appraisal of whether or not a situation requires particularly early remedial action. A fixed minimum RDRR is used in preference to a free-floating JIT (Just In Time) strategy, in recognition of the need to take action in good time, so as not to panic masters of other ships into emergency action.

The rule structure is designed to:

(a) note the presence of a potential hazard, assess the threat in terms of expected time to domain violation (if applicable) and derive the PSMT for avoidance action;

(b) At PSMT: identify the type of encounter, fix the status of ownship and perceived status of hazard ship (give-way or stand-on) at this time in the encounter; once decided, status is maintained throughout the encounter, unless circumstances change - the regulations rule out changes in status due solely to changes in relative positions through avoidance manoeuvres;

(c) Negotiate the stages of the encounter, with appropriate safety margins.

It is anticipated that (c) will be extended to incorporate the possible need for a change in strategy in response to untoward action by other vessels, including emergency manoeuvres; such additional rules may be edited into the existing rule base, as described below

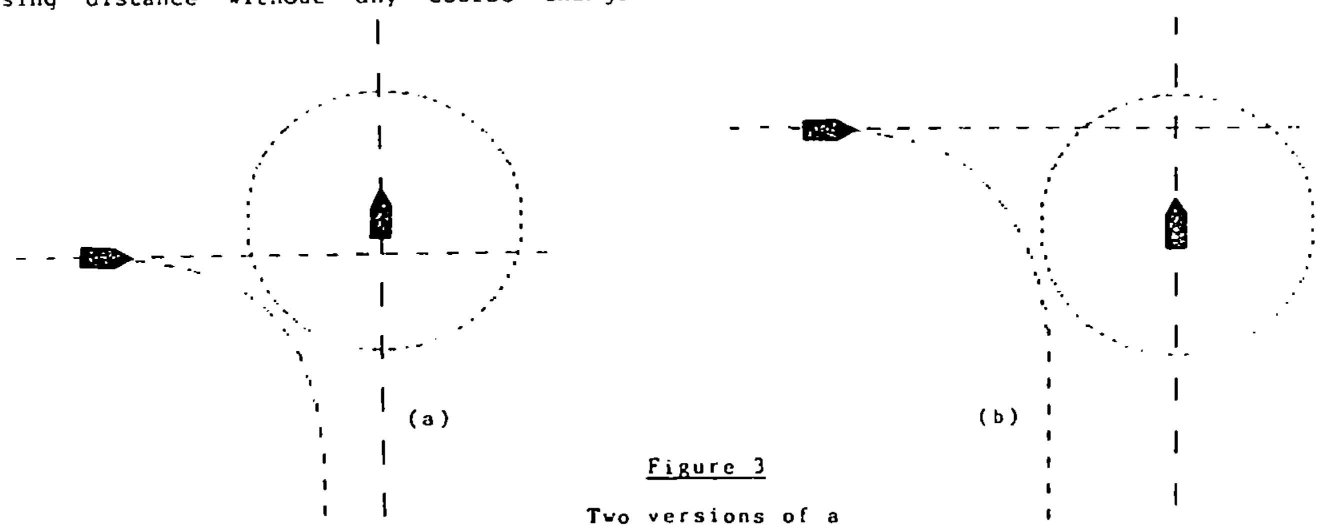

Figure 3

Two versions of a crossing encounter.

A domain of 0.8 n.m. is shown surrounding the hazard vessel in each case.

The three primary types of encounter are:

(1) Head-on
Two vessels are approaching each other on near-reciprocal bearings (within +/- 6 degrees, in this system);

(2) Crossing
Two vessels are approaching each other at such an angle that each can see one side of the other; more specifically, each vessel is in the sector subtended by the port or starboard light, 112.5 degrees measured from the bow (front) on each side, of the other;

(3) Overtaking
A ship is deemed to be overtaking if it is approaching another from within the stern-light sector (the rear 135 degree arc).

The rules as implemented embody the anti-collision regulations, plus clarification of specific practical detail as noted by Colley [op cit] and used to good effect in an earlier simulation. The inference engine operates on a discrete time-interval of 20 seconds. At each step, the inference engine will:

(a) take dynamic information via communication channels from radar and own-ship sensors (or simulators); from these parameters (speed, course, position) it will generate secondary data: relative bearing, relative velocity components, etc. - for use in rule evaluation;

(b) apply the appropriate rule to these data, to ascertain the new situation; each rule involves a test on these system variables, and may entail evaluation of further functions of these variables - such functions form part of that rule;

(c) trigger display and control/simulator outputs in response to any change in status; invoke any new rules indicated by such a change, until a 'defer' flag is reached, inhibiting any further action on the rule base until the next time-step.

The rule base is an hierarchical structure, based on a binary tree, but with considerably more flexibility in its links to left and right 'sub-trees': in this rule structure, links may lead to any other node, up or down; two or more links may lead to a common node; a link may loop back to a node at a higher level. The rule base is also in effect a state table, each node representing a possible state of ownship: no encounter currently in progress, second stage of a 'parallel-up' overtaking encounter, etc. Extension of the rule base is achieved by creating the relevant new nodes and resetting the necessary links to insert them at the appropriate points in the structure. It is envisaged that this task will ultimately be simplified by a suitable software utility.

Each rule incorporates a Boolean function, using inequality tests on combinations of displacement and velocity vectors. The majority of processing is thus centred on evaluation of trigonometrical functions. This factor becomes increasingly dominant as situations increase in complexity - a major consideration in choice of language and form of rule structure for a time-critical system. Floating-point hardware can yield significant benefit in such a task.

6. FULL CIRCLE

The authors have highlighted the need for rationalisation of information reaching the bridge, particularly as it relates to the task of collision avoidance. They have demonstrated the ability of current technology to elicit much of this information directly, and to present it in a manageable form. Further, they have demonstrated a process whereby prompt and efficient analysis of that information may provide timely advice to the mariner on matters vital to ship safety. En route it is hoped that non-seafarers will have gained some insight into the uncertainties which can precipitate maritime disasters, and that seafarers will perceive potential benefits in the supportive role to be played by such technology.

Maritime disasters down the ages have provided much food for thought on misunderstandings and misinterpretations of data. Various ingenious, and usually highly impractical, ideas have been put forward after most major incidents, aimed at reducing either the risk or the consequences of such events. Perhaps the last word should go to one C.E. Kelway, on the sinking of the 'Empress of Ireland' in 1914, with the greatest loss of passenger lives in peacetime of any liner [8]. His proposed system involved the use of stopwatches, a mechanical computer on the bridge for solving problems in trigonometry, and the use of 'Hertzian waves' (radio waves); three-quarters of a century does not seem an ungenerous gestation period for such a life-saving concept.

7. REFERENCES

1. Blackwell,G.K., Colley,B.A., & Stockel,C.T.
Look-Ahead Simulation as a Tool for Real-Time Expert Systems.
SCSI European Multiconference, Rome, Italy, June 1989. pp 157-162

2. Colley,B.A., 1985
Computer Simulation of Marine Traffic Systems.
Ph.D. Thesis, Plymouth Polytechnic (CNAA)

3. International Maritime Organisation, 1981
International Regulations for Preventing Collisions at Sea.

4. Miller,K.M., 1990
A Navigation and Automatic Collision Avoidance System for Marine Vehicles.
Ph.D. Thesis, Polytechnic South West (CNAA)

5. Dove,M.J., 1984
A Digital Filter/Estimator for Control of Large Ships in Confined Waters.
Ph.D. Thesis, Plymouth Polytechnic (CNAA)

6. Tapp,N.J., 1989,
A Non-Dimension Mathematical Model for use in Marine Simulators.
M.Phil. Thesis, Plymouth Polytechnic (CNAA)

7. Goodwin,E.M., 1975
A Statistical Study of Ship Domains.
Journal of Navigation, vol. 28, pp328-344

8. Croall,J., 1981
Disaster at Sea.
Stein & Day, New York, USA.

Proceedings of the 1990 Summer Computer Simulation Conference

July 16-18, 1990
The Skyline Hotel
Calgary, Alberta, Canada

Twenty-Second Annual Summer
Computer Simulation Conference

Edited by
Bill Svrcek
John McRae

CONFERENCE SPONSOR

SCS

THE SOCIETY FOR COMPUTER SIMULATION

PRINTED IN THE U.S.A.

Strategic Planning for Intelligent Real-Time Management

G.K. Blackwell and C.T. Stockel

Department of Computing, Polytechnic Southwest
Drake Circus, Plymouth, Devon, England

ABSTRACT

An expert advisory system is naturally limited in the quality of advice which it may offer, by both the quality and completeness of the information provided for its consideration. As with its human counterpart, such a system can not function fully effectively with partial or inaccurate data. Likewise, an over-simplistic rule base could lead to erroneous conclusions in complex situations. It is essential that an intelligent advisory system should clearly indicate its own limitations, and moderate its recommendations by appropriate safety margins.

Such caveats apply particularly to real-time sensor-based systems; most sensors are subject to a certain level of tolerance (error), and many systems with sensory input have 'blind spots'. Moreover, the decision function of any real-time expert system is subject to a unique constraint, that of timeliness: the advice or decision of such a system must be delivered within a limited time-frame (often a matter of seconds) if it is to affect the outcome of the process being monitored. Both the structure of the rule base, and the action of the inference engine in firing those rules, must be tailored to meet such deadlines.

The forward-planning nature of such systems dictates that proposed strategies be evaluated in terms of future consequences. The use of simulation as a predictor mechanism within an expert system is described in a previous paper by the authors - which also covers an example of an intelligent simulation environment. In a simple situation, with a small number of component elements, the look-ahead requirement may not pose serious problems; however, as with a board game, the number of alternative strategies increases factorially in relation to the number of 'pieces'. Both the rule base, and the inference engine which actions it, must be so structured as to manage situations to the required level of complexity within the time available - without loss of definition. The optimal strategy will be the one that leads to the 'best' outcome, measured against a number of criteria, suitably weighted.

The authors are concerned with the development of an on-board intelligent knowledge-based system for marine collision avoidance, linked to radar and other sensory input systems. Work has previously centred on two-ship encounters in the open sea; current developments concern extension of the rule base to handle multi-ship encounters. In the first instance, such scenarios may be resolved to a sequence of two-ship encounters, identifying which hazard ship poses the greatest threat at each stage. However, action to avoid Hazard 1, taken in isolation, may lead to an increased risk from Hazard 2, necessitating even greater or more urgent manoeuvring. There is even possibility of 'oscillation', with each of the two hazards alternately becoming more threatening. A full solution of the multi-ship problem must take account of the total situation, as a competent master mariner would. A simple manoeuvre at an early stage is much preferred to a complex 'slalom' course at the latest possible moment. Such considerations must naturally be balanced against excessive course deviations.

Specific developments now in hand include:
extension of the rule base to incorporate branch-points in the decision sequence;
application of Games Theory techniques to optimise process of evaluation of alternative strategies;
inclusion in rule base of early course change strategy, and dynamic change of strategy with circumstances.
The system uses an Acorn 'Archimedes' RISC machine, with simulated sensory inputs. It has also been tested on board a research vessel, with inputs from radar and associated marine electronic equipment.

1. INTRODUCTION

Recent advances in microcomputer technology offer exciting possibilities in the synthesis of computer simulation of complex mechanical systems with expert controller systems based on proven artificial intelligence concepts. A desktop computer offering processing power hitherto associated with mainframes, providing a true multi-tasking environment, front-ended by a sophisticated WIMPS (Windows-Icons-Mouse-Pointer System) user interface and a multi-colour graphics display, forms a readily-accessible environment for this type of application. A variety of software tools are available for use on such a system, combining with a range of analog and digital, serial and parallel, communication ports to cover a wide array of tasks.

In the maritime environment, this suggests the provision of a real-time computer advisory system for the master of an ocean-going vessel. Such tasks as the monitoring of engine and steering gear, track-keeping (including complex manoeuvres in difficult conditions) and on-board weather routeing are all well-suited for advisory support, or even possibly automatic control, by such an expert system.

One specific area in which automated expert support could yield significant benefit - as evidenced by numerous recent shipping disasters - is the recognition and avoidance of potential collisions or groundings. Given relevant background information, and suitable input channels for continuously updating dynamically-changing factors in the vessel's environment, an Intelligent Knowledge-Based System (IKBS) would maintain a continuous 'watching brief' on ships and other potential hazards in the vicinity. On detecting the need for alterations of course and/or speed to eliminate the danger of hazardous potential collision situations, a suitably-tailored inference process would make recommendations of appropriate actions to ensure safety. This implies the development of an accurate model for the behaviour of the vessel undertaking such manoeuvres, and for the perception of a hazard at a sufficiently future time in order for the manoeuvre to be carried out.

In order to provide meaningful assistance in the form of an electronic adviser, it is necessary first to have some understanding of the situation at the control station (the bridge) of a modern vessel; a competent advisory system will be expected to handle that same situation as expertly, and within the same time constraints, as the ship's master. He is constantly in receipt of a wide range of information from a variety of sources: instruments relaying rudder setting and engine rpm, as well as speed through the water, roll, pitch and yaw; positional information and speed-made-good (i.e. relative to the seabed) are provided with the aid of such navigational systems as Decca, Omega, Loran-C and Sat-Nav; any obstacles above or below the surface are brought to notice by radar and sonar detection systems; the relative speed and course of other vessels in the vicinity is given by ARPA, the Automatic Radar Plotting Aid. On top of this continuous stream of electronic data, visual and auditory information, including that passed on by other members of the crew, must be processed in real time, and appropriate responses made. The consequences of inadequate processing, generally referred to as 'human error', have been all too evident of late.

Certain prerequisites for the proposed expert system are immediately evident. First, it must be capable of receiving and processing all of the electronically-sensed data which is normally displayed on the bridge instruments, plus any additional information which may be input by the user. Second, it must be cognisant, as is the ship's master, of performance characteristics of the vessel, plus any other relevant parameters - length, draught, etc. Third, it must be conversant with the International Regulations for Preventing Collisions at Sea [2], and the ramifications of those regulations in a multiplicity of situations. Fourth, it must be capable of arriving at a reasoned response to current circumstances, and presenting that response clearly, together with a succinct and unconfusing explanation of its reasoning if required. Fifth, it must be able to re-evaluate changing circumstances, and give its conclusions within a matter of seconds.

As indicated, the advisory task under consideration here is that of collision avoidance. Development is taking place within the framework of a total voyage management system; it is intended that the collision avoidance task will ultimately be integrated with other elements of the system under the control of a WIMPS 'desktop' multi-tasking environment. For the present, the expert advisory system for collision avoidance is being developed as one application under such an environment. The requirements for such a system, as outlined above, are explored in detail below.

2. THE EXPERT SYSTEM AS A TOOL

Society today is suffering from 'information overload'. Every aspect of life, domestic or business, is beset by a continuous stream of data from a range of sources. The significance of this data to the task in hand varies from vital through irrelevant to positively obstructive. One important role of any top executive's personal assistant is to filter incoming data, ensuring that management's attention can be given wholly to the relevant facts (hopefully in an easily-digestible form), and not swamped or blocked by extraneous 'noise'.

An expert system for voyage management should be seen in such a role. One clear reason for human error in the maritime setting is an inability to extract and collate the relevant material from the mass of data available, then formulate and implement an apposite strategy, all within a strictly limited time-frame. A rule-based system could rapidly identify pertinent facts, discard irrelevancies, and construct a plan of action (or possibly a set of alternatives) for consideration by the manager, i.e. the ship's master.

In the specific instance of marine collision avoidance, an initial framework for the rule base already exists, in the aforementioned anti-collision regulations. By themselves, however, these regulations are inadequate for such a task, since they leave the onus for certain major decisions with the mariner. Such phrases as 'in good time', 'a clear turn', and 'a safe distance' are open to a wide variety of individual interpretations, which may themselves vary with different contexts; sea state, nature of vessel and cargo, possible manoeuvring restrictions (such as coastal features or shipping lanes), traffic density - all affect decisions on collision avoidance manoeuvres. As with any competent human P.A., therefore, the electronic adviser would have to be conversant with that broader spectrum of 'rules' (i.e. situational responses) which derive from applied commonsense and years of experience; such experience, gleaned from human experts, must be formalised into an inference structure which gives substance to the bare bones of the official regulations. Continuing the analogy, the system should be unobtrusive, providing information only as requested, in a meaningful way - but with the capability for drawing attention to matters in need of immediate action.

Clearly, an electronic perception of the environment will differ in certain respects from that arrived at through human senses, and electronic appraisal of a given scenario may necessarily take a different path from a human line of thought. For example, a human observer may spot almost immediately that another ship is turning, by its change in visual aspect; a computer would need several successive ARPA readings before reaching the same conclusion. Conversely, the human lookout would need to observe another vessel for some time before deducing that it was on a constant bearing, and thus on a collision course - a fact which would be immediately apparent to the computer, even down to time and position of impending collision, by extrapolation from velocity and displacement vectors. It follows that the requirement is for a system which, in using its processing capabilities to the optimum, parallels human reasoning, without necessarily duplicating it in every detail. The final outcome should nevertheless correspond to the human expert's decision process, combining observed good practice with self-evident commonsense.

A graphical representation of the current scenario is clearly a prerequisite; a well-planned display, invoking good use of colour and scale, would allow the immediate situation to be taken in by a cursory glance, without excessive detail. Such detail could be offered as 'optional extras', for example:

(a) a status report, giving current speed and bearing, details of any current manoeuvre, plus any pertinent data on nearby ships or other hazards (notably, projected time to collision or near-miss, if applicable);

(b) an appraisal of the current situation, indicating advised course of action, with supporting rationale for such advice available on demand; in a real-time expert system, that rationale would necessarily include reference to prior, and likely future, events - a dimension absent from most IKBS.

Such options could be provided at the press of a button in a menu-driven WIMPS environment, with no requirement for keyboard dexterity or other new skills. Information would be to hand exactly as and when needed, without confusion or complication, particularly bearing in mind that this facility is likely to be most needed at times of greatest stress .

3. THE PLYMOUTH DEVELOPMENT SYSTEM

A customised expert system shell has been developed initially on an Atari ST micro-computer, and recently transferred to an Acorn Archimedes RISC (Reduced Instruction Set Computer) system. After much consideration of suitable programming languages for such a task, the 'C' language was chosen for a number of reasons:

(1) the constraints of real-time processing dictate the need for optimum efficiency, in terms of speed, compatible with the complexities of program structure required for such a rule-based system;

(2) many of the rules are expressed in terms of trigonometrical considerations, giving a substantial bias towards mathematical calculation, and thus favouring a language which deals efficiently with such needs;

(3) C is the native language of most WIMPS environments, giving access to facilities at all levels;

(4) entities of various types - rules, ships, encounters (not to mention windows, icons etc. within the WIMP system) - may all be handled very effectively as 'structures' in C; earlier intentions of using an object-oriented front-end to this system have been set aside, since this versatile C data type meets all perceived needs (particularly under the new ANSI standard) without the apparent overheads of other, object-oriented or declarative languages.

Both computers were chosen on the basis of processing power at relatively low cost, good WIMPS and graphics facilities, and capability for a variety of input/output options, both digital and analog. A decision to transfer to the Acorn RISC machine was based on the need for compatability with other ship management functions currently under development by the Ship Control Group at Plymouth, as well as its superiority on a number of fronts. The RISC_OS multitasking environment, coupled to versatile I/O handling, are well suited to the planned integrated system.

The screen display is in the form of three windows, as shown in figure 1. A square graphics window, filling approximately half of the screen area, displays the tracks and positions of one's own vessel (own-ship) and any other vessels in the vicinity, with markers to show positions at corresponding times; own-ship is shown as a green track, others in various colours - at the current stage of development, only one other vessel (the most threatening) will be 'processed' by the rule structure as a potential hazard; the display in the graphics window may be enlarged by reference to a pop-up menu (invoked and actioned by mouse buttons), and any part of the scene studied in detail by scrolling the window horizontally or vertically. A second window optionally displays a review of the progress of the situation to date, including the decisions made by the expert system at various stages; this window may also be scrolled, or expanded over the graphics window temporarily, to display further detail of the stages of an encounter. A third window may be

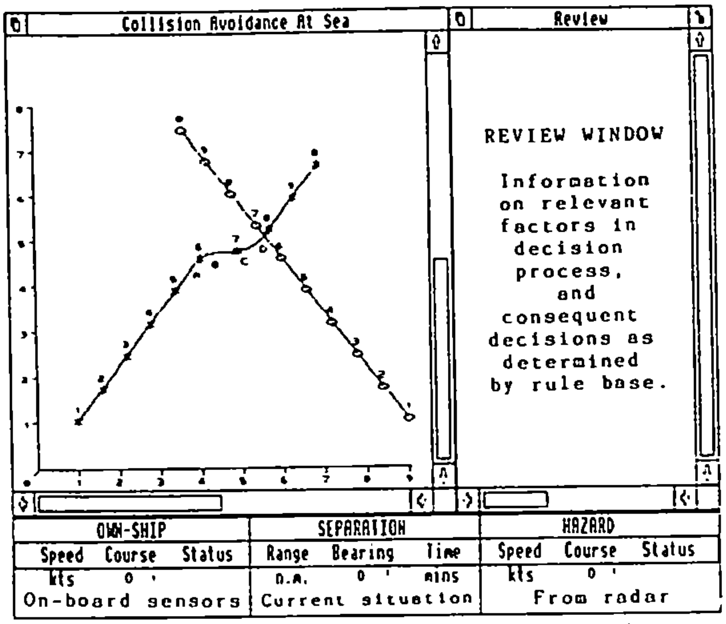

Figure 1

Multiple-window screen display, showing scroll bars and window control icons. Bottom 'situation' window may optionally be hidden; 'review' window may be extended or scrolled to give more detailed information; 'tracks' window may give (scrollable) large scale view, or complete small scale view of current situation.

selected and de-selected as required, again by a mouse-operated pop-up menu; this window, appearing below the other two, displays data on the course, speed and current manoeuvring (if any) of own-ship, plus such data as is available (via radar, etc.) on any other vessel in the vicinity considered to be a hazard - future developments will allow selection from the graphics window of which potential hazard is to be 'shadowed' by this window at any time.

Early development comprised the expert system logic, plus ship simulator modules for own-ship and a hazard vessel, all housed in one computer; a later version made use of three microcomputers, one each for expert system, own-ship simulator, and hazard ship simulator. Having thus verified the integrity of the expert system, the current implementation is designed to operate in any of three modes, two being simulator trials, the third real data:

(1) A semi-intelligent hazard vessel, using a simplified expert system for control, and a dumb own-ship controlled by the current system, are both simulated as additional tasks under the multitasking RISC_OS, on the same computer as the expert system; software communication channels between the tasks simulate radar data to and from hazard vessel, and sensor /control information between own-ship and expert system;

(2) Hazard and own-ship are simulated as above, but in separate computers from the expert system; communications are via serial and parallel links respectively;

(3) Inputs via serial and parallel ports, as in (2), are from genuine radar and sensors aboard own-ship, i.e. a research vessel with the computer system on board; such input data has as yet been used only to test the efficacy of this form of input, and observe the response of the system; however, it is envisaged that the control outputs from the system will in due course be connected to the automatic controls already fitted to the vessel (previously used with marked success in computerised ship guidance [3]).

Simulation of own-ship and hazard for (1) and (2) are based on the linear ship model [4]. A more accurate model of own-ship is required in the decision logic of the expert system itself for reference in the look-ahead module described in section 4. The non-linear modular model derived by Tapp [5] is used for this purpose.

4. CRITERIA FOR DECISION-MAKING

The ultimate objective of any collision avoidance procedure is to avoid collisions; a secondary objective is to inform other mariners, through positive action, that the possibility of such an incident has been recognised and is being dealt with. Both of these aims are served by manoeuvres which seek to maintain a zone of clear water around one's vessel, generally referred to as the domain. This observed practice of experienced mariners [6] has been taken as the prime objective of the expert system; a circle, of radius suited to the vessel and circumstances, centred on ownship, has proved an adequate representation of the domain to date. From this stems a broader area of concern, identified by Davis [op. cit.] as the arena, namely the zone of proximity within which another vessel or object should be considered as a potential hazard, evaluated as such, and appropriate avoidance action initiated if necessary. Such proximity considerations were refined by Colley [7] to a time-based criterion, the RDRR (Range-to-Domain/Range-Rate); this represents the expected time to domain infringement, calculated from current velocities. A ship's master applying a 12-minute RDRR, for example (based on handling characteristics for his vessel) would evaluate a potential encounter 12 minutes before domain infringement, and implement appropriate action at that time.

Results from an early version of the system, using the RDRR concept, indicate that such a fixed 'decision scheduler' is not appropriate to every encounter situation. A flexible time-constraint, matching decision time to interval needed for a safe manoeuvre, ensures optimum manoeuvring time: adequate, but not excessive. This Predetermined Safe Manoeuvring Time (PSMT) is found by simulating any projected encounter well in advance, for increasing RDRR values (starting from some preset minimum, say 10 minutes), until safe clearance is achieved. The simulation exercise is an integral element of the decision logic, and corresponds to a human appraisal of whether or not a situation requires particularly early remedial action. Figure 2 shows the principle of the look-ahead simulation used within the expert system to select an appropriate RDRR, i.e. the PSMT, for any encounter. A fixed minimum RDRR is used in preference to a free-floating JIT (Just In Time) strategy, in recognition of the need to take action in good time, so as not to panic masters of other ships into emergency action.

The rule structure is designed to:

(a) Note the presence of a potential hazard, assess the threat in terms of expected time to domain violation (if applicable) and derive the PSMT for avoidance action;

(b) At PSMT: identify the type of encounter; fix the status of own-ship and perceived status of hazard ship (give-way or stand-on) at this time in the encounter; once decided, status is maintained throughout the encounter, unless circumstances change - the regulations rule out changes in status due solely to changes in relative positions through avoidance manoeuvres;

(c) Negotiate the stages of the encounter, with appropriate safety margins.

It is anticipated that (c) will be extended to incorporate the possible need for a change in strategy in response to untoward action by other vessels, including emergency manoeuvres; such additional rules may be edited into the existing rule base, as described in section 5.

The rules identifying and handling the three primary types of encounter are as follows:

(1) Head-on
Two vessels are approaching each other on near-reciprocal bearings (within +/- 6 degrees, in this system). A vessel in such an encounter must consider itself give-way and alter course to starboard (right); it should not begin to come back onto course until the other vessel has passed 90 degrees on the port beam (left side);

Figure 2 Multiple Look-Ahead Simulation

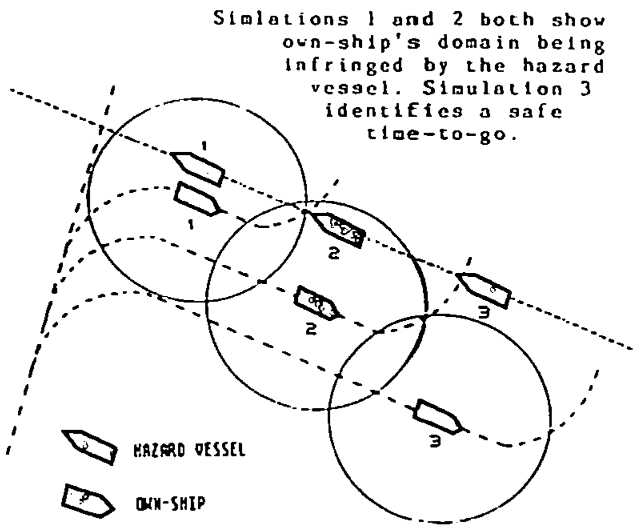

(2) Crossing
Two vessels are approaching each other at an angle such that each can see one side of the other; more specifically, one vessel is in the sector subtended by the port or starboard light, 112.5 degrees measured from the bow (front) on each side. The ship approaching the port side of the other vessel should give way, altering course to starboard and going behind. It should not alter back until the other vessel has passed 15 degrees to port of own-ship's original heading;

(3) Overtaking
A ship is deemed to be overtaking if it is approaching from within the sternlight sector (the rear 135 degree arc). It should give way as follows:
(a) If on near-parallel course (within 10 degrees, in this system), go behind; alter back when safe to do so;
(b) If on port quarter of stand-on vessel, action as for (a);
(c) If on starboard quarter, and set to overtake ahead on current course, AND if able to parallel courses without infringeing domain, parallel-up; when well past and clear to pull across, return to original heading;
(d) Otherwise, go behind to port, alter-back when safe to do so.

The rules embody the International Regulations for Preventing Collisions at Sea, plus clarification of specific practical detail as noted by Colley and used most effectively in an earlier simulation. Using these rules, a collision avoidance manoeuvre is taken through five stages, as follows:

(1) 20 minutes in advance of projected domain infringement, look-ahead simulation is used to identify a suitable PSMT;

(2) When the PSMT is reached, the give-way vessel initiates a turning action;
(3) Once the required turn is effected (minimum 30 degrees in this system, to conform with 'clear turn' requirement, but more if necessary to clear the domain), the vessel proceeds on a straight course;
(4) When the original heading may be resumed without further risk of domain infringement (and any other conditions are met), alter-back is initiated, necessitating a turn in the other direction;
(5) When back on the original heading, vessel proceeds on that heading; manoeuvre over.

5. INFERENCE ENGINE AND RULE STRUCTURE

The inference engine operates on a discrete time-interval of 20 seconds. At each step the inference engine will:

(a) take dynamic information via communication channels from radar and own-ship sensors (or simulators); from these parameters (speed, course, position) it will generate secondary data: relative bearing, relative velocity components, etc. - for use in rule evaluation;

(b) apply the appropriate rule to these data, to ascertain the new situation; each rule involves a test on these system variables, and may entail evaluation of further functions of these variables - such functions form part of that rule.

(c) trigger display and control/simulator outputs in response to any change in status; invoke any new rules indicated by such a change, until a 'defer' flag is reached, inhibiting any further action on the rule base until the next time-step.

The rule base is an hierarchical structure, based on a binary tree, but with considerably more flexibility in links to left and right 'sub-trees': in this rule structure, links may lead to any other node, up or down; two or more links may lead to a common node; a link may loop back to a node at a higher level. Figure 3 shows a schematic of this structure, illustrating part of the initial rule base.

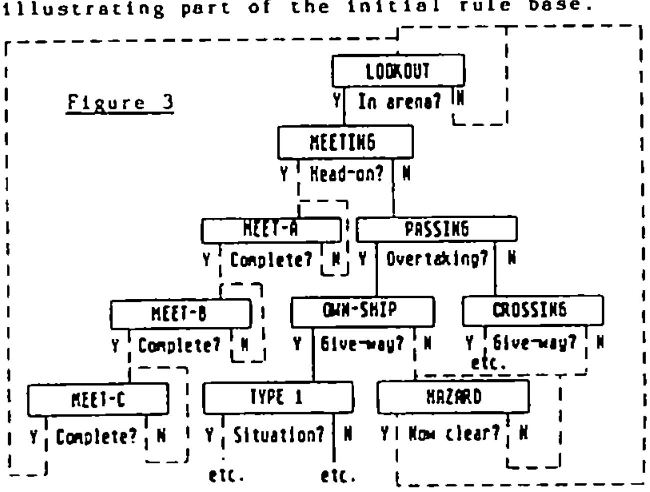

Figure 3

The rule base is also in effect a state table, each node representing a possible state of own-ship: no encounter currently in progress, second stage of a 'parallel-up' overtaking encounter, etc. Each node contains:
(1) left and right links to other nodes (or looping back to the same node) in the rule structure; one of these links will be followed at every decision step (i.e. every time-step);
(2) a mathematically-based decision function, used to determine which link is followed at each stage;
(3) an 'immediate/defer' flag for each link, indicating action time for the next rule;
(4) ship status flags, to pass display and control information.

Extension of the rule base is achieved by creating the relevant new nodes and resetting the necessary links to insert them at the appropriate points in the structure. It is envisaged that a software utility will ultimately simplify this task for the non-computer-specialist.

Each rule incorporates a Boolean function, using inequality tests on combinations of displacement and velocity vectors. The bulk of processing is thus centred on evaluation of trigonometrical functions. This factor becomes increasingly dominant as situations increase in complexity - a major consideration in choice of language and form of rule structure for a time-critical system. Not surprisingly, floating-point hardware can yield significant benefit in such a task!

6. REFERENCES

1. Blackwell, G.K. and Stockel, C.T., 1988
Simulation of Ship Encounters Using Multiple Linked Intelligent Systems.
Proceedings SCSC89, Austin, Texas pp766-771

2. International Marine Organisation, 1981.
International Regulations for Preventing Collisions at Sea.

3. Miller, K.M., 1990
A Navigation and Automatic Collision Avoidance System for Marine Vehicles.
Ph.D. Thesis, Polytechnic Southwest (CNAA)

4. Dove, M.J., 1984
A Digital Filter/Estimator for Control of Large Ships in Confined Waters.
Ph.D. Thesis, Plymouth Polytechnic (CNAA)

5. Tapp, N.J., 1989
A Non-Dimensional Mathematical Model for use in Marine Simulators.
M.Phil. Thesis, Plymouth Polytechnic (CNAA)

6. Goodwin, E.M., 1975
A Statistical Study of Ship Domains.
Journal of Navigation, Vol. 28, pp328-344

7. Colley, B.A., 1985
Computer Simulation of Marine Traffic Systems.
Ph.D. Thesis, Plymouth Polytechnic (CNAA)

Maritime Communications and Control

London
21 – 23 November 1990

Maritime Communications and Control
21 – 23 November 1990

Paper 7

The use of a mathematical model in a collision avoidance system

G K Blackwell, BTech, PGCE, J Chudley, BSc and M J Dove, MSc, PhD, CEng, MRINA, FRIN
Ship Control Group, Polytechnic South West

SYNOPSIS

The work of the Ship Control Group at Polytechnic South West encompasses a number of research areas. The overall objective is an integrated navigation system, which will automatically steer a vessel along a predetermined track, dealing with hazards competently and in accordance with the regulations.[1] This has, to date, included substantial research on both ship models and automatic control and guidance systems.

The authors are concerned with the development of an on board intelligent knowledge-based system (IKBS) for marine collision avoidance, as part of an overall voyage management system. A purpose-built expert system shell and rule structure has been implemented in the context of a windows-icons-mouse-pointer system (WIMPS) environment on a multitasking microcomputer suitable for installation on board ship.

Sensory inputs from radar and associated instrumentation may be supplemented by data via the keyboard. Output is in the form of advisory messages and alarms where appropriate. An auto-control option facilitates computer controlled manoeuvring through encounter situations. Simulation runs in conjunction with a computerised track-keeping system, already in operation aboard the polytechnic research vessel, confirming the ability of the combined system to divert around potential hazards and then return the vessel onto track. Field trials are anticipated early in the new year, following the installation of the necessary instrumentation.

INTRODUCTION

Recent accidents at sea, together with a series of crises that have increased the price of fuel oil, have made shipowners and operators more safety and economy conscious: this in turn has made the requirements on ship control more demanding, particularly in confined waters, where extensive manoeuvring is needed. The ability to predict the path of the ship with a high degree of precision is clearly of major importance.

It has been suggested that 85% of all marine collisions and groundings are due to human error, and of these 90% occur in coastal waters.[2] On this evidence alone there is a case for research into, and development of, automated control and guidance systems. Mariners on a sea passage are likely to experience periods of relatively uneventful sailing, interspersed with periods requiring careful attention and substantial decision-making, such as traversing a busy seaway or entering port. The potential hazards of such a regime are twofold: on the quiet stretches a false sense of security can lead to impending danger being overlooked until it is too late; conversely, information input at busy times may overload the decision-making process (ie the officer of the watch), leading to ill-judged actions or dangerous delays in manoeuvring. Both of these problems could be obviated by an electronic monitoring system, which would analyse sensory inputs such as radar and navigational information, giving reasoned and pertinent advice to the mariner on the bridge. If such a system were available to advise the officer of the watch, perhaps disasters such as those involving the *Exxon Valdez* and the *Marchioness* could be avoided.

Grahame Blackwell gained a BTech degree in maths, statistics and computing, and a PGCE from Brunel University. He has since been engaged in a variety of computer applications, notably in the scientific and real-time fields. He later moved into education, and played an active part in the development of educational computing, and is now a senior lecturer in computing at Polytechnic South West. He is currently registered for a part-time PhD, the subject of which is 'an expert systems approach to collision avoidance'.

Whilst serving an apprenticeship with a marine engineering company, John Chudley gained an HNC in mechanical and production engineering. This was followed by a first class honours degree in nautical studies from Plymouth Polytechnic (now Polytechnic South West). He has since undertaken one year of full time research into mathematical modelling of ships and is now a lecturer at the Polytechnic.

Before becoming a lecturer, Dr M J Dove served 20 years as a navigator in the Merchant and Royal Navy. During this period he obtained an MSc in control technology. He recently gained a PhD for work in 'automatic control guidance'. This work led to the development of the Ship Control Group at Polytechnic South West. This group specialises in R&D in marine navigation and automatic guidance. Dr Dove is a member of RINA and was elected to a Fellowship of the Institute of Navigation for services to education development in navigation. He currently manages 400 undergraduates taking degrees in marine studies at the Polytechnic.

Much has been written about the fully automated unmanned ship, which may only feasibly be put into operation in the distant future. However, the Ship Control Group at Polytechnic South West have developed, and are improving, a system to assist the mariner. Dove,[3] investigated the use of Kalman filters for improvements to position fixing in the approaches to a port. In joint research with Dove, Burns,[4] has studied the guidance problem; the results from these twin projects form an optimal filter together with an optimal controller, thus breaking the optimal guidance problem down into two distinct phases. The system is now installed and operational aboard the Polytechnic survey vessel.[5] An expert system for collision avoidance, currently under development will be interfaced to the system at a later date.[6]

DEFINITION AND GENERAL FORM OF THE MATHEMATICAL MODEL

Mathematical models of ship dynamics may be broadly split into three categories, covering applications in:

1. ship manoeuvrability analysis:
   a. ship design;
   b. waterway improvement and port facilities;
   c. safety regulations and casualty studies.
2. training and research simulators;
3. shipboard manoeuvring predictors.

The research group in Plymouth is developing a shipboard manoeuvring predictor to provide assistance to the navigator in track control and path-keeping, and otherwise planning the trajectory of the vessel so as to minimise the likelihood of mishap.

Development of the model starts with a set of generalised equations to express the dynamics of a rigid body in a fluid medium, derived from Newton's second law of motion. These equations are then extended to model the complex hydrodynamic forces and moments experienced by a hull manoeuvring in response to the control inputs of rudder and propeller. By integrating through small time steps, the motions of the vessel can be solved. Further forces and moments are then introduced in response to the disturbance inputs of wind and tide.

A ship at sea has six degrees of freedom of motion: translation along three orthogonal axes and rotation about each of those axes. Three of these – heave, roll and pitch – may be ignored, since they make no contribution to motion in the horizontal plane (the area of concern for all practical purposes in ship manoeuvring predictors). By taking the ship's centre of mass as the origin for the co-ordinate system, the remaining three forms of motion may be represented by:

Surge: $m(\dot{u} - vr) = X$	(forces generating surge)	
Sway: $m(\dot{v} - ur) = Y$	(forces generating sway)	
Yaw: $I_z\dot{r} = N$	(moments generating yaw)	(1)

The forces and moments on the left side of these equations represent the hydrodynamic and aerodynamic reactions on the hull of the ship in response to the applied control forces, in addition to any other disturbance inputs. The mathematical model for ship manoeuvring should cater for the following characteristics of ship dynamics:

1. realistic turning for all rudder angles including helm delay and loss of speed in the turn; the response to rudder action should be asymmetric for a single screw ship;
2. realistic acceleration and deceleration including inertial effects and engine delays;
3. ahead and astern motion;
4. reduction in ahead motion, effective helm and squat effect in shallow water;
5. drift caused by a variable tidal stream;
6. drift and yaw caused by a wind, variable in both magnitude and direction, acting on the hull and superstructure of the ship;
7. single or twin screw operation, including independent control of each screw in both directions, and turning rate;
8. variable pitch operation of the screw;
9. ship motion due to waves.

The ship model should also be amenable to alteration in order to simulate a wide range of hull forms and vessel sizes, ranging from a fishing boat to a supertanker.

Current research on ship manoeuvring modelling tends to favour the modular model, in which the individual elements such as the hull, propeller, rudder, engines and external influences on a manoeuvring ship, are each represented as separate interactive modules. Each self-contained module (relating to hydrodynamic or control forces, or to external effects) is constructed by reference to a detailed physical analysis of the process being modelled. The system as a whole is then modelled by combining the individual elements.

The equations of motion for a modular manoeuvring model are generally expressed in the form:

$$m\dot{u} - mrv = X_H + X_P + X_R + X_E$$
$$m\dot{v} + mru = Y_H + Y_P + Y_R + Y_E$$
$$I_z\dot{r} = N_H + N_P + N_R + N_E \qquad (2)$$

where suffixes H, P, R and E denote components of hull, propeller, rudder and external forces respectively.

One advantage of this approach is its potential for expansion. Extra modules may be added to those above to simulate bow and stern thrusters, for example, represented in the equations by additional terms of the form $X_B, X_S, Y_B, Y_S, N_B, N_S$. Hence effects such as those referred to above may be investigated by 'customising' the model.

MATHEMATICAL MODELLING AT POLYTECHNIC SOUTH WEST

The hydrodynamic forces of a rigid body, with forward speed in free surface waves, present a difficult problem; accuracy of prediction of the motion depending significantly on the ability to determine these forces. The hydrodynamic coefficients required in the equations of motion of a body moving through a fluid, are usually classified into three general categories:

1. *Static*. Due to the components of linear velocities of the body relative to the fluid.
2. *Rotary*. Due to components of angular velocity;
3. *Acceleration*. Due to either linear or angular acceleration components (also termed 'added mass').

The number and type of hydrodynamic coefficients required will vary according to the complexity of the problem being investigated; the type of mathematical model; and the extent to which various effects are included in the representation. Research over a number of years at Polytechnic South West has culminated in a modular manoeuvring model developed by Tapp for use in a marine simulator;[7] the text includes an overview of the various methods for deriving the required coefficients for such a model, taking the form of equation (2).

SHIP MODELLING IN COLLISION AVOIDANCE

The nature and complexity of ship models incorporated in collision avoidance systems depends primarily on whether one is concerned with shore-based Vessel Traffic Control (VTC) or bridge-based advisory systems. Research by the authors is advancing on both fronts: the prototype ship-based IKBS described two years ago, has progressed substantially in various respects and is the subject of most of the latter part of this paper[6]; and collaborative work with the University of Rome II is yielding interesting developments in the field of parallel processing and shore-based marine traffic control.[9]

A shore-based advisory/control system requires a working knowledge of the performance characteristics of all vessels under surveillance at any time, since the ability to give helpful directives presupposes a capability to accurately predict the likely outcome of manoeuvres by one or more vessels. In general, those vessels will each fall into one of three categories:

1. known vessel, for which hydrodynamic coefficients are documented in detail, and for which an accurate model may be defined;
2. known or unknown vessel, for which hydrodynamic coefficients are not available, but which may be fitted into one of a number of 'classes' of ship, and thus matched approximately to a less tightly defined model;
3. unknown vessel for which little or no performance data is available (apart from immediate observation), and which must therefore be judged on a very broad basis, with a wide margin for error allowed around any assumptions which have to be made.

Such considerations lead naturally to the need for:

1. a file of 'modelling' data for all vessels for which such data is available, and which frequent the waters under surveillance;
2. a library of 'standard' vessel types, to which the majority of ships may be fitted with reasonable accuracy;
3. a rule base which, in so far as is possible, avoids the need for suppositions about any ship not susceptible to handling by (1) or (2) – this may well mean preferential treatment for such vessels, on the basis of 'when in doubt, keep clear'.

Shore-based VTC, with its requirement for detailed modelling of a number of ships simultaneously, demands substantial computing power to operate effectively in real-time; hence the reference to parallel processing. In most other respects the problems of shore-based collision avoidance, and the function of ship models therein, mirror those of bridge-based systems as considered further in the paper. Further references to this topic are to be found in Colajanni,[9] Degré,[10] and Bootsma and Poldermann.[11]

The requirement for ship models in the shipboard Collision Avoidance System (CAS) is twofold: firstly, the expert system logic must be able to model ship behaviour in order to predict likely outcomes of manoeuvres: secondly, the development environment for such a system for much of the time will consist of a simulated environment in which own-ship and the various hazard vessels are represented by appropriate computer models.

In the latter case, all of the vessels may be represented as accurately as current modelling techniques permit. There is no problem of 'unknown identity', since vessels for the simulation exercise may be chosen from those for which full sets of hydrodynamic coefficients are available. The only hindrance is the collection of a suitably representative cross-section of sample ship data. The gathering of hydrodynamic coefficients for a single ship is not a trivial task,[12] and few such sets are as yet readily available.[13] However, in the absence of specific data, parameters for a 'typical' ship of a specific type are adequate for such a task – as proposed above for 'standard vessel types'. The question of adequate computing power for such a task need not present a problem, since separate vessels may if necessary be simulated on separate processors – such a technique has been shown to have various benefits at certain stages of development.[14] Alternatively, a multitasking environment with substantial processing power (using one or more processors) yields comparable benefits; such a setup is described later in this paper. A less satisfactory solution is to run the simulation slower than real-time to circumvent processor limitations; such a technique could aid in early development, but may mask operational limitations of the system if relied upon too heavily.

Modelling of vessel behaviour, as an integral element of the expert system decision logic, falls into two very distinct categories: modelling of own-ship's response to applied controls and prevailing environment; and prediction of expected behaviour of a hazard vessel. There is no problem with the former, since installation of shipboard CAS presupposes prior evaluation of the hydrodynamic coefficients for the vessel in question. The latter, however, poses a very real pragmatic question: given that one cannot read (still less guide) the thoughts of the master of another vessel, what benefit may be derived from assessing the consequences if a specific control were applied at a specific time on that vessel? The situation is very different from shore-based VTC, where advice (or directives) may be communicated to any vessel, in the manner of air traffic control. At best, one may take into account the perceived manoeuvrability of the other vessel, and thus its ability to resolve a tricky situation; at worst, one has to admit the possibility of an adverse manoeuvre on his part. Until proven otherwise (by unreasonable procrastination or other 'rogue' action), it may be reasonably assumed that the other vessel will conform with the 'rules of the road'; but any attempt to model the other vessel's likely avoidance manoeuvres would be purely speculative, and should not form the basis for action on the part of own-ship.

Given the above caveats, the modelling of a hazard vessel has little part to play in the expert system logic, other than consideration of continuation on current course and speed. Turning action may be observed, and intent to avoid possibly inferred, but the extent of the manoeuvre may not be prejudged in detail. Manoeuvrability of the other vessel may be a factor for consideration in the decision logic for emergency situations – this has yet to be considered. Otherwise, any attempt to model the behaviour of another vessel, beyond constant velocity projection, is based on untested assumptions and therefore of dubious value. This is, of course, a separate issue from that of vessel manoeuvring restrictions and priorities as laid down in the anti-collision regulations, which should form an integral part of the decision structure.

THE COLLISION AVOIDANCE SYSTEM

An expert system comprises four major elements:

1. *The user interface.* Through which the user requests or provides information, and the expert system communicates its findings;
2. *The knowledge base.* Comprising all information to be considered in making decisions and formulating strategies;

5. *The rule structure*. That set of logical connectives by which the facts are assessed, and those decisions and strategies inferred;
4. *The inference engine*. The underlying abstract principles by which the rules are applied to the knowledge base, irrespective of the specific application (in this case, marine collision avoidance).

The functional demands of a marine CAS differ from those f more conventional expert systems in all of these respects. onsiderations of the operating environment; the sensory na- ure of much of the input data; real-time processing of rules omparing vector relationships; a need to rationalise current ecisions in terms of earlier events and anticipated future onsequences, all call for specialist techniques not feasible in ff-the-shelf packages. Basic requirements in these four areas, nd continuing development in the Plymouth system, are out- ned below.

'he user interface

In computing circles much is made of the 'man-machine nterface', and nowhere is it more important than in such a ituation, where clear and rapid communication of meaningful nformation may avert potential disaster. Simple selection of vailable options should be matched by non-(computer-) tech- ical presentation, without any need for the user to learn new kills. Input of relevant data may be facilitated by menu selec- ion and 'dialogue boxes'.

A graphical representation of the current scenario is clearly prerequisite; a well-planned display, invoking good use of olour and scale, allows the immediate situation to be taken in y a cursory glance, without excessive detail. Additional data onsidered relevant for optional selection comes in the form f:

1. a status report; giving present speed and bearing, details of any current manoeuvre, plus any pertinent data on nearby ships or other hazards (notably, projected time to collision or near-miss, if applicable);
2. an appraisal of the current situation, indicating advised course of action, with supporting rationale for such advice available on demand; in a real-time expert system, that rationale necessarily includes reference to prior and likely future events – a dimension absent from most IKBS.

These options are provided at the press of a button in a menu- lrive WIMPS environment, with no requirement for keyboard lexterity or other new skills. Information is to hand exactly as ind when needed, without confusion or complication – an mportant consideration for a facility likely to be most needed it times of greatest stress.

The screen display is in the form of three windows, as shown n Fig 1. The contents of these windows are in a broad sense as reviously defined, and as described more fully in the 1988 aper.[5] Whilst there has been no reason to change the nature of he information displayed, substantial advances have been nade in presentation and user control of that display: colour is ised to good effect, notably to identify the tracks of different hips; scrolling and sizing of windows gives greater control over the information displayed; mouse-activated menus and dialogue boxes give instant access to a variety of features – an example of such a 'pop-up' menu is shown.

Figure 1 also shows a 'hazard control' window superimposed upon the expert system display. The development environment includes simulated hazard vessel(s) with optional manual control.

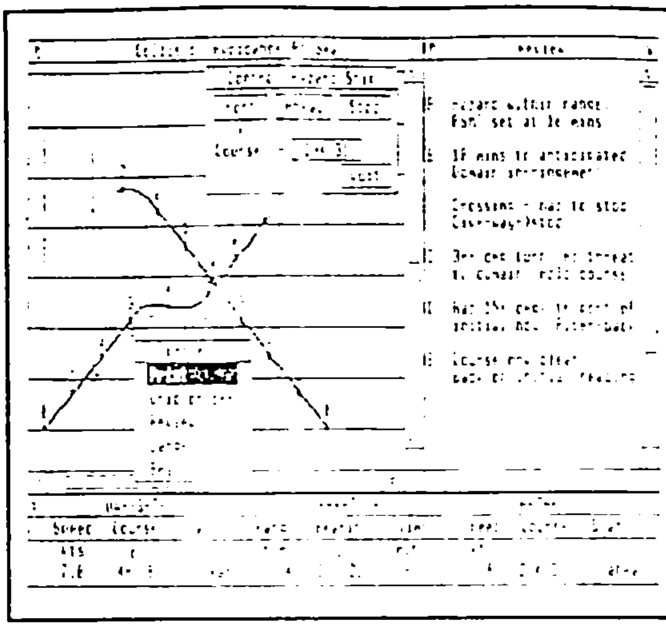

Fig 1: Example screen display, showing text and graphics windows, a pop-up option menu, and a dialogue box used to control a simulated hazard vessel

The knowledge base

Information relevant to the system comprises a combination of: static information, which relates to the vessel and any other relevant factors invariant over a voyage, and dynamic information which changes with time, such as heading, position, speed and data on potential hazards.

The former may be fixed parameters for the vessel (eg beam), or data for a particular voyage (eg gross tonnage); it could include relevant electronic chart data. The latter must be sampled at regular intervals, by sensors attached to engines, rudder etc, and interfaces to instrumentation such as Automatic Radar Plotting Aid (ARPA) and Decca. Secondary data on vector relationships must be calculated for reference by the rule base.

The expert system has been designed in a modular form to facilitate input from a variety of sources, as follows:

1. pop-up dialogue boxes are provided for keyboard input,
2. data from instrumentation is input via a parallel interface to a dedicated communications processor, linked to all sensor/control functions. Decca and ARPA are also accessed via this unit.
3. communication modules enable input from
   a. other processors handling associated tasks, such as automatic track-keeping;
   b. other tasks running under the multitasking operating system on the same processor (notably for simulated test environment, described later in this paper);
   c. associated tasks running in multitasking mode on a shared-memory multiprocessor unit (a planned objective); data acquisition from electronic charts could also be handled by this technique, or by the front-end processor referred to in (b).

Parameters for the ship model constitute static data, whilst behaviour of the model at any time is determined with reference to current and predicted values of dynamic data. In particular circumstances such as shallow-water manoeuvring, alternative ship parameters will apply; a model which adapts dynamically to such conditions is a future aim.

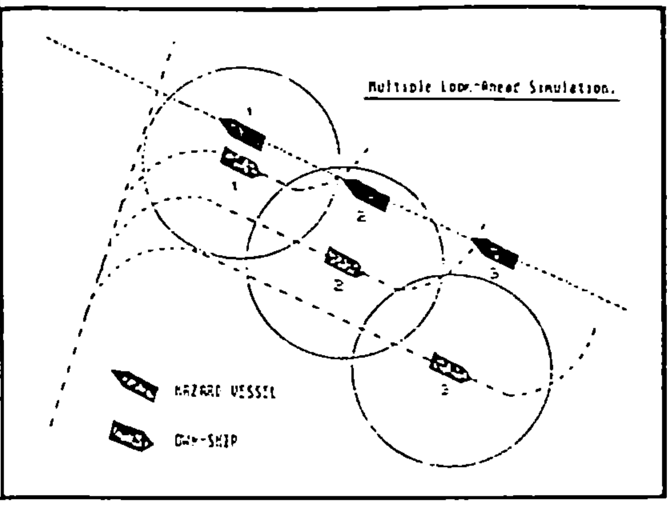

Fig 2: Illustration of multiple look-ahead simulation, showing consequences of starting manoeuvre 10, 12 and 14 min before anticipated domain infringement: In cases 1 and 2 (10 and 12 min respectively), the hazard vessel will still infringe the domain; in case 3, the hazard vessel will safely clear the domain – a PSMT of 14 min

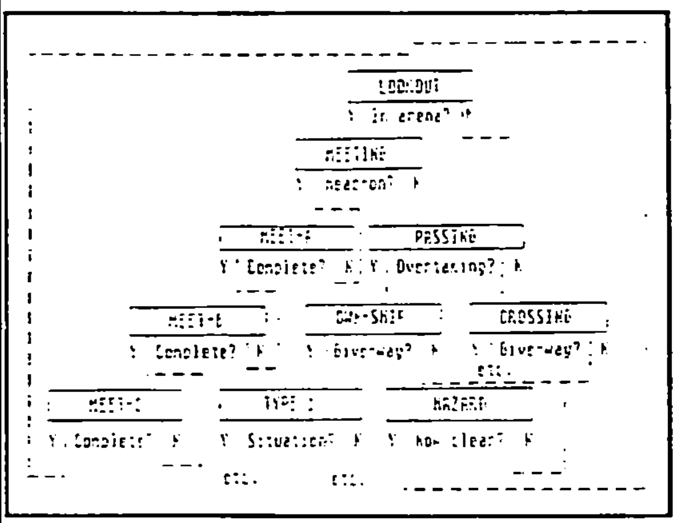

Fig 3: Schematic of rule structure, showing immediate (solid) and deferred (dotted) links between rules

The knowledge base might also include information on other ship types and behaviours, not so much for ship modelling as for recognition purposes in implementing the regulations, eg fishing vessels, vessels constrained by draught, etc. This is likely to need supplementing by manual input, as fully automatic vessel recognition is unlikely to be feasible in the near future. Manual identification, by selection from an option menu, is envisaged initially for the present system once the rule base is extended to include such 'special cases'.

The rule structure

The ultimate objective of any collision avoidance procedure is to avoid collisions; a secondary objective is to inform other mariners, through positive action, that the possibility of such an incident has been recognised and is being dealt with. Both these aims are served by manoeuvres which seek to maintain a zone of clear water (generally termed the domain) around one's vessel. This observed practice of experienced mariners,[15] has been taken as the prime objective of the expert system. A broader area of concern, the arena, corresponds to the zone of proximity within which another vessel or object should be evaluated as a potential hazard, and appropriate avoidance action initiated if necessary.[16] Colley,[17] refined such considerations to a time-based criterion, the Range-to-Domain/Range-Rate (RDRR): this represents the expected time to domain infringement, calculated from current velocities. A ship's master applying a 12 min RDRR, for example (based on his experience of his vessel) would evaluate a potential encounter 12 min before domain infringement, and implement appropriate action at that time.

Results from an early version of the system, using the RDRR concept, indicate that such a fixed 'decision scheduler' is not appropriate to every encounter situation. A flexible time factor, matching decision time to interval needed for a safe manoeuvre, ensures adequate, but not excessive, manoeuvring time. This Predetermined Safe Manoeuvring Time (PSMT) is found by simulating any projected encounter well in advance for increasing RDRR values (starting from some fixed minimum, eg 10 min), until safe clearance is achieved. The simulation exercise is an integral element of the decision logic, corresponding to a human appraisal of whether or not a situation needs particularly early remedial action. Figure 2 shows the principle of the look-ahead simulation used within the expert system to select an appropriate RDRR, ie the PSMT, for any encounter. A fixed minimum RDRR is used rather than a free-floating Just In Time (JIT) strategy, since action must be taken in good time so as not to panic masters of other ships into emergency action. This look-ahead technique is currently being extended to evaluate various options in terms of an overall cost function of speed loss, time delay, fuel usage – subject to an overriding safety constraint.

The rule structure is designed to:

1. Note the presence of a potential hazard, assess the threat in terms of expected time to domain violation (if applicable) and derive the PSMT for avoidance action;
2. At PSMT: identify the type of encounter; fix the status of own-ship and perceived status of hazard ship (give-way or stand-on) at this time in the encounter; once decided, status is maintained throughout the encounter, unless circumstances change – the regulations rule out changes in status due solely to changes in relative positions through avoidance manoeuvres;
3. Negotiate the stages of the encounter, with appropriate safety margins; reactions to adverse action by hazard ship, including emergency manoeuvring, are currently being added to the rule base – hence the 'hazard control' test facility shown in Fig 1.

The structure is based on a binary tree, but with considerably more flexibility in nodal links to left and right sub-trees. Extension of the rule base is achieved by creating the relevant new nodes and resetting the necessary links to insert them at appropriate points in the structure. It is envisaged that a utility program will ultimately be provided to simplify this task.

Figure 3 shows a schematic of this structure, illustrating part of the initial rule base; dotted lines indicate deferred steps.

The inference engine

The inference engine operates on a discrete time-interval of 20s. At each step the inference engine will:

1. take dynamic information via communication channels from radar and own-ship sensors (or simulators). From

these parameters (speed, course, position) it will generate secondary data: relative bearing, relative velocity components, etc, for use in rule evaluation;
2. apply the appropriate rule to these data, to ascertain the new situation. Each rule involves a test on these system variables, and may entail evaluation of further functions of these variables; such functions form part of that rule.
3. trigger display and control/simulator outputs in response to any change in status; invoke any new rules indicated by such a change, until a 'defer' flag is reached, inhibiting any further action on the rule base until the next time-step.

A customised expert system shell, developed initially on an Atari ST microcomputer, has recently been transferred to an Acorn Archimedes Reduced Instruction Set Computer (RISC) system. Software development has been completely in the language C, for a number of reasons:

1. constraints of real-time processing dictate the need for optimum efficiency, in terms of speed, compatible with the complexities of program structure required for such a rule-based system;
2. many of the rules are expressed in terms of trigonometrical considerations, giving a substantial bias towards mathematical calculation, thus favouring a language which deals efficiently with such needs;
3. C is the native language of most WIMPS environments, providing access to facilities at all levels;
4. entities of various types: rules, ships, encounters (not to mention windows, icons etc within the WIMP system), may all be handled very effectively as 'structures' in C. Earlier intentions of using an object-oriented front-end to this system have been set aside, since this versatile C data type meets all perceived needs (particularly under the new ANSI standard) without the overheads of other, object-oriented or declarative languages.

Both computers were chosen on the basis of processing power at relatively low cost, good WIMPS and graphics facilities, and capability for a variety of input/output options, both digital and analog. A decision to transfer to the Acorn RISC machine was based on the need for compatibility with other ship management functions currently under development by the Ship Control Group at Plymouth, as well as its superiority on a number of fronts. The RISC_OS multitasking environment, coupled to versatile Input/Output handling, is well suited to a planned integrated system.

DEVELOPMENT ENVIRONMENT

The system is designed to operate in any of three modes, two being simulator trials, the third real data:

1. a semi-intelligent hazard vessel (guided by a simplified expert system, with manual option) and a dumb own-ship controlled by the current system, are both simulated as additional tasks under the multitasking RISC_OS, on the same computer as the expert system. Software communication channels between the tasks simulate radar data to and from hazard vessel, and sensor/control information between own-ship and expert system;
2. hazard and own-ship are simulated as above, but in separate computers from the expert system. Communications are via serial and parallel links respectively;
3. inputs via serial and parallel ports, as in (2), are from genuine ARPA and sensors aboard own-ship, ie a research vessel with the computer system on board. Control outputs are available for use with the automatic guidance system already in operation.

CONCLUSIONS

Thorough investigation and development of ship models has been incorporated into current development of a collision avoidance system. The ergonomics of the system have also benefited from application of the latest Man-Machine Interface (MMI) techniques. The introduction into the decision logic of a look-ahead simulation module has resolved the question of a safe manoeuvring time, and paved the way for identification of optimal manoeuvres. Not least, the evolution of a totally modular structure provides for simple modification or extension of any aspect of the system: ship model(s); rules; communications and control; and user options.

REFERENCES

1. 'International regulations for preventing collisions at sea', International Maritime Organization (1981).
2. 'Human error in merchant marine safety', Panel on Human Error in Merchant Marine Safety, National Academy of Science, Washington DC (1976).
3. M J Dove, 'A digital filter/estimator for the control of large ships in confined waters', PhD Thesis, Plymouth Polytechnic (1984).
4. R S Burns, 'The automatic control of large ships in confined waters', PhD Thesis, Plymouth Polytechnic (1984).
5. R S Burns, M J Dove and K M Miller, 'A control and guidance system for ships in port approaches', International Conference on Maritime Communications and Control, IMarE, London (1988).
6. G K Blackwell and C T Stockel, 'Optimal strategies for a marine real-time expert system', Proceedings of 1990 European Simulation Multi conference, Society for Computer Simulation, Rome (June 1990).
7. N J Tapp, 'A non-dimensional mathematical model for use in marine simulators', MPhil Thesis, Plymouth Polytechnic (1989).
8. G K Blackwell, B A Colley and R S Stockel, 'An intelligent knowledge-based system for marine collision avoidance', International Conference on Maritime Communications and Control, IMarE, London (1988).
9. M Colajanni, 'Parallel and distributed simulation in vessel maritime traffic surveillance', Proceedings of 1990 Summer Computer Simulation Conference, Society for Computer Simulation, Calgary, Alberta, Canada (July 1990).
10. T Degré, 'Assessment of the effects of vessel traffic services on marine traffic safety by simulation techniques', Proceedings of 1987 Summer Computer Simulation Conference, Society for Computer Simulation, Montreal, Quebec, Canada (July 1987).
11. R Bootsma and K Polderman, 'ATS and VTS – some observations towards a synthesis', The Journal of Navigation, 40, 1 (January 1987).
12. M A Abkowitz, 'Measurement of hydrodynamic characteristics from ship manoeuvring trials by system identification', Transactions SNAME, Vol 88 (1980).
13. C L Crane jnr, 'Manoeuvring trials of a 278 000 dwt tanker in shallow and deep water', Transactions SNAME, Vol 87 (1979).
14. G K Blackwell and C T Stockel, 'Simulation of ship encounters using multiple linked intelligent systems', Proceedings of 1989 Summer Computer Simulation Conference, Society for Com-

MODELLING AND SIMULATION 1991

Proceedings of the 1991 European Simulation Multiconference

June 17-19, 1991
The Panum Institute • Copenhagen, Denmark

Edited by
Erik Mosekilde

SCS

ESM '91 • 1991 EUROPEAN SIMULATION MULTICONFERENCE

Organized by The Society for Computer Simulation (SCS)
in Cooperation with
EUROSIM, SIMS, JSST, and CASS

SHIPS THAT PASS IN THE NIGHT
Intelligent Management of a 'Fuzzy' Event Stream

Blackwell G.K., Rangachari J., Stockel C.T.
Department of Computing
Ship Control Group
Polytechnic South West
Plymouth, Devon, ENGLAND

Introduction

the main, any real-world situation may be sidered, for the purposes of simulation, as her a series of discrete events or a contous process. A minority of undertakings may viewed as a sequence of discrete subcesses, each taking either of the above ms - in extremis, this description may be lied recursively.

task of marine collision avoidance cannot fitted neatly into any such outline. At one el, a maritime voyage may be viewed as a uence of encounters with other individual sels; at another, as a progression of oeuvres to avoid potential collisions or r-misses. Often there is a direct one-one respondence between these two types of nts; sometimes there is not, as one oeuvre may take account of two or more ential hazards. More significantly, this tinction is not clear-cut, as one captain's ltiple encounter' may be another's 'series encounters', with avoidance strategies ected accordingly. Likewise a carefullynned avoidance manoeuvre may be modified at ate stage by one master to take account of new arrival on the scene; another may plete according to plan before commencing ropriate action to deal with the new ard. Individual outlook, coupled with cumstances, together define the boundaries ween successive events and the level of plexity acceptable in a single event.

authors are concerned with development of on-board computer-based expert system (ES) marine collision avoidance. Previous ers [1,2] have considered discrete two-ship ounters: types of manoeuvres, safe manoeuvg times, simulation environments. Recent k has focused on the ongoing nature of the k, with regard both to test environment and imal decision-making (the nature of the play format has also been given serious sideration - see section 3). Under the new ime, the rule base will generate an optimal ategy to encompass all hazard ships in the inity, separating or grouping them as ropriate for manoeuvring purposes. Hazards enter or leave the operational area at egular intervals, any new arrival resulting a review of planned action. Comparison of ernative strategies is based on a multiiate cost function, taking account of such tors as fuel usage, speed loss, time delay, ject to an overriding safety constraint; ghtings of cost factors may be adjusted ording to priorities for a particular age.

The development system accomodates a variable number of hazards. An additional hazard may be activated at any time, preferably outside the radar range of own-ship so as to become effective on entering that range. A hazard is de-activated automatically if it is moving away from own-ship, outside the radar range.

The system is based on the overall philosophy that the most immediate threat is planned for in the most detail; less urgent hazards are incorporated into the strategy planning at a lower level of definition, plans 'firming up' as each in turn comes to the fore. This would seem to mirror the human situation, where a ship's master will deal with the most pressing threat whilst taking into account any future needs to negotiate other impending hazards.

2. The Collision Avoidance Task

The task of marine collision avoidance is, in the first instance, encompassed by the International Regulations for Avoiding Collisions at Sea [3]. These regulations cite specific rules to cover specific types of encounter situations at sea. Such situations divide broadly into three categories:

(1) head-on;
(2) crossing;
(3) overtaking.

The nature of the encounter is identified by the position of one vessel with respect to another, vis-a-vis the navigation lights of each vessel. These navigation lights cover three separate sectors around a vessel, as shown in figure 1. Figure 2 illustrates a variety of encounter situations, as defined by the lights.

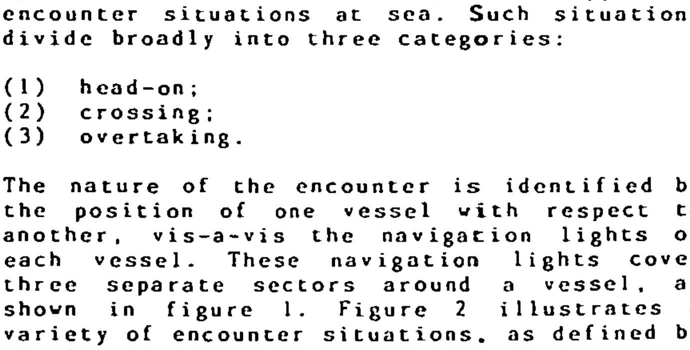

Figure 1. Sectors subtended by a ship's navigation lights.

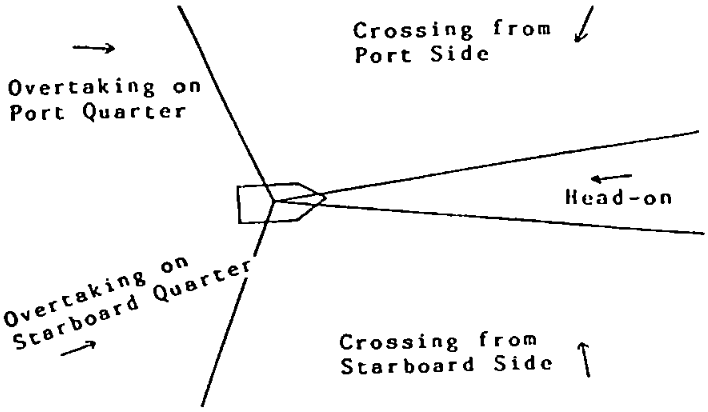

Figure 2. Types of ship encounters

The regulations are adhered to by every responsible mariner, and have the force of law in most countries. However, they are not totally definitive in terms of decisions and actions, and in many cases must be interpreted by mariners in the light of experience and discretion. For example, the master of a vessel on a collision/near-miss course is expected to take action 'in good time' to avert disaster; this may be interpreted as anything from 10 to 30 minutes before closest point of approach (CPA). Likewise, a vessel manoeuvring to avoid another should indicate its intention by making a 'clear turn' (10-30 degrees), and maintaining a 'safe distance' from the other vessel (how long is a piece of string??).

Such flexibility (or vagueness) leads to differing responses by different masters, or even by the same master under different circumstances: a 'safe distance' when carrying a cargo of tractor parts may not be quite so safe with a hold full of volatile chemicals; the luxury of a 20 degree 'clear turn' in the open sea may not be so practical (or necessary to make the point) in a crowded seaway such as the Dover Strait. Clearly decisions on such factors as time to take avoiding action depend crucially on the performance characteristics of the vessel concerned: a cross-channel ferry has a substantially faster response, and rather smaller turning circle than (for example) a Very Large Crude Carrier (VLCC).

These differences in interpretation may prove significant in particular situations, since it is possible for the nature of an encounter to change with time, giving latitude for masters of the vessels in contention to interpret the same situation in two different ways. For example, figure 3 illustrates a situation which, evaluated 20 minutes before CPA would be seen as an overtaking encounter; 10 minutes before CPA it would be recognised as a crossing encounter. According to the regulations, responsibility for avoidance would rest with either one vessel or the other, depending on the interpretation applied. The regulations also direct that, once evaluated, the nature of a situation should be regarded as invariant for manoeuvring purposes (even if a changed scenario evolves, as in fig. 3); hence two masters could each deem the other responsible for avoiding, leading to the necessity for emergency manoeuvring at a late stage.

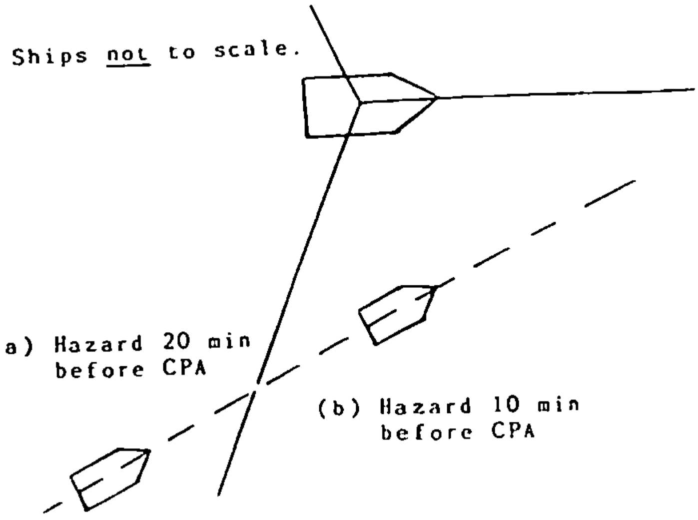

Figure 3. Encounter situation changing from hazard overtaking (hazard give-way) to crossing (own-ship give-way).

The regulations also allow for an early course change, not subject to fixed protocols, if a decision to manoeuvre is made well before any anticipated incident. This 'well before' can likewise be very loosely interpreted, leading to non-standard manoeuvres fairly late on in an encounter. Figure 4 illustrates the classic 'radar-assisted collision', in which master B has chosen a 'course change' to port (being somewhat to port of the other vessel's track) as opposed to the regulation 'avoidance manoeuvre' to starboard. The combination is a recipe for disaster, illustrating the need for clear adherence to well-defined rules. It has even been observed that some mariners will feel free to make an unrestricted 'course change' no matter how late on, if the vessels are not actually on a direct collision course - thus in some cases converting a near-miss into a disaster.

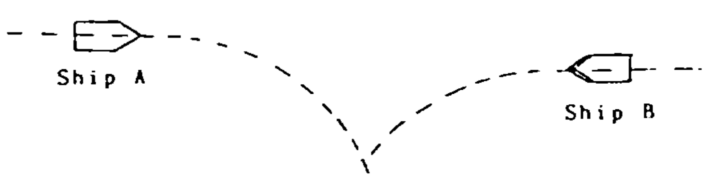

Figure 4. Ship A: regulation turn to starboard
Ship B: 'course change' to port
Head-on near-miss becomes collision.

Previous research [4] has identified the major variants of the three standard encounter situations, and the strategies employed by mariners in those situations. These strategies have been formalised into a set of 'rules' which form the basis of an Intelligent Knowledge-Based System (IKBS) for marine collision avoidance [5]. Criteria relating to time for action, safe clearances and adequate turns are system parameters provided by the user, according to taste and circumstance. The knowledge base for the system comprises a combination of:

(1) 'static information', such as that above plus length, draught and performance characteristics of the vessel;

(2) 'dynamic information' such as current speed and course of own-ship plus radar data on any potential hazards in the vicinity; anticipated future developments include input of satellite position-fixing data and information on fixed navigational hazards from electronic chart digital information systems (ECDIS).

3. The Expert System

The system is built on the anti-collision regulations, filled out by observed good practice among mariners. Successful functioning of the system depends on two fundamental concepts relating to zones around a ship: the domain and the arena.

The domain is that area around own-ship which the master would wish to keep free of any potential hazards. Early research in this field identified the domain as disparate circle sectors [6] or alternatively a circle with own-ship off-centre [7], reflecting:
(a) the need for more sea room to starboard (the normal manoeuvring side); and
(b) comparatively less risk from hazards to the rear (closing speeds being generally less from behind).

For the purposes of this system, a circle centred on own-ship provides an adequate model for the domain; the domain radius, selected by the operator, is of course the representation of the 'clear distance' required by the regulations to be kept from other vessels.

The arena is that boundary around own-ship within which any hazard should be evaluated as a potential threat, and avoiding action initiated if appropriate [7]. This appraisal of a 'threat boundary' has been found to be over-simplistic, as it takes no account of different closing rates by various hazards. A better discriminator is a time-based boundary, triggered by time to domain infringement [4]. This RDRR (Range-to-Domain/Range-Rate) acts as follows: if a 10-minute RDRR is active (for example), a ship with a closing speed of 12 knots would come within this 10-minute threshold at a distance of 2 nautical miles from own-ship's domain boundary; a vessel closing at 6 knots would cross the RDRR threshold only 1 nautical mile from own-ship's domain.

The RDRR criterion has itself been found to have limitations, since every encounter situation is unique in terms of juxtapositions and relative velocities of vessels involved. Section 4 deals with the provision of a look-ahead facility which optimises selection of the appropriate RDRR for a particular encounter.

The expert system (ES) consists of an inference engine, or logical framework, applied to a rule base consisting of directives on preferred/alternative actions at various stages of different types of encounter. The rules are structured into a highly flexible state table, identifying different sequences of manoeuvring stages appropriate to different encounter types; more complex manoeuvres, including various emergency options, may be added to the rules by simply extending the 'binary tree' data structure which holds them. The rules function by reference to the knowledge base referred to previously. Each rule incorporates a decision function, usually involving extensive trigonometrical calculations based on velocity and displacement vectors, plus possibly intersections of turning arcs and domains. The ES operates on a 20-second time interval, evaluating the current scenario at each step and advising on preferred action.

The user interface is based on a Windows-Icons-Mouse-Pointer System (WIMPS) environment, the output display comprising three windows:

(1) a graphics window showing tracks of own-ship and any relevant hazards in the vicinity, in different colours; sequentially-numbered markers are used on each track to show corresponding positions in time;

(2) a text window giving (on request) a summary of circumstances and recommendations to date, cross-referenced by sequential lettering to points on own-ship's track in the graphics window;

(3) an optional 'status' window, displaying current information on speed and course of own-ship, perceived speed and course of most threatening hazard vessel, separation and time to domain infringement, and present encounter state ('clear', 'ready' or stage of current manoeuvre).

Window (1) incorporates a 'pop-up' menu which permits: selection of large/small scale; 1-mile grid on/off (superimposed on display); provision/update of review information in window (2). The contents of window (2) are initially limited to brief statements of fact and advice; this window may be expanded to full screen width (covering window (1)) to allow fuller detail to be given on each stage.

Much consideration has been given to the best format for the graphics window, in terms of the position and motion of own-ship within this window: should the display be Head-up or North-up; should own-ship be static at the centre and relative motion of all other objects shown? Clearly, Head-up display implies a rotational transformation of all other graphical data as own-ship executes a turn. The question of disorientation requires careful consideration, balancing objective/ subjective viewpoints against (for example) distortion of true ship tracks by a display of relative motion. Possibly the most significant factor is the ongoing nature of the task; simple continuous plotting of ship tracks in absolute motion will quickly lead to own-ship being 'lost' off the edge of the display.

The solution proposed, and in the process of implementation, is the plotting of absolute ship tracks (showing true motion) using a variant of the 'rolling road' principle: a boundary (not displayed) is fixed 25% of the way in from each edge of the window (see figure 5); as own-ship reaches this boundary, the whole display is relocated back to a position such that own-ship is at the centre of the display. In this way, at least the 50% of the display around own-ship (horizontally and vertically) at any time is maintained intact, giving continuity and minimising disorientation. It is accepted that periodic 'jumps' in the display may require a degree of 'acclimatisation', but this is considered of minor significance by comparison with either of the two alternatives: relative motion gives a distorted view of manoeuvres, leading to a totally false appraisal of any situation; full window wraparound (off the side - back on the other side, off the top - back on the bottom, and vice versa) could be extremely untidy without being very helpful.

The review text marks a departure from conventional expert systems technology, in that the observations and advice given are time-constrained. The logical analysis of 'this is so and that is so, therefore the following is true' provides an inadequate model for such time-based decision processes. The corresponding logic pattern is of the form 'this was so three minutes ago, so-and-so has occurred since, indicating that the following will be true in twelve minutes - unless such-and-such happens, in which case the consequences will be ...' etc. In short, the added dimension of time broadens the whole decision process, making any clear resumé of that process a very demanding task. This is particularly true in the light of the requirement to present pertinent information in a succinct, easily-digested form. Hence the format of short, pithy comments with scope for expansion on request. The review window may also be scrolled up and down, to view the complete time-spectrum of events.

The user interface is totally mouse-driven, with menus activated and options selected by mouse and pointer operation. It is envisaged that at a later stage there will be a requirement for text input to supplement the knowledge base; this, too, will be simplified by the use of dialog boxes (such boxes are already extensively used in simulation exercises on this system - see section 4).

4. The Test and Development Environment

Clearly the Expert System in isolation cannot give any meaningful illustration of its performance in real-world conditions. In order for the system to function at all, there is a need for continuous streamed input of data relating to own-ship and other vessels in the vicinity. The former will be obtained from sensors attached to rudder, engines etc., the latter from radar (or ARPA). As indicated in section 1, such inputs are not a very practical proposition for ongoing system development. Hence a sophisticated simulated test environment has been evolved to match the needs of an increasingly sophisticated IKBS.

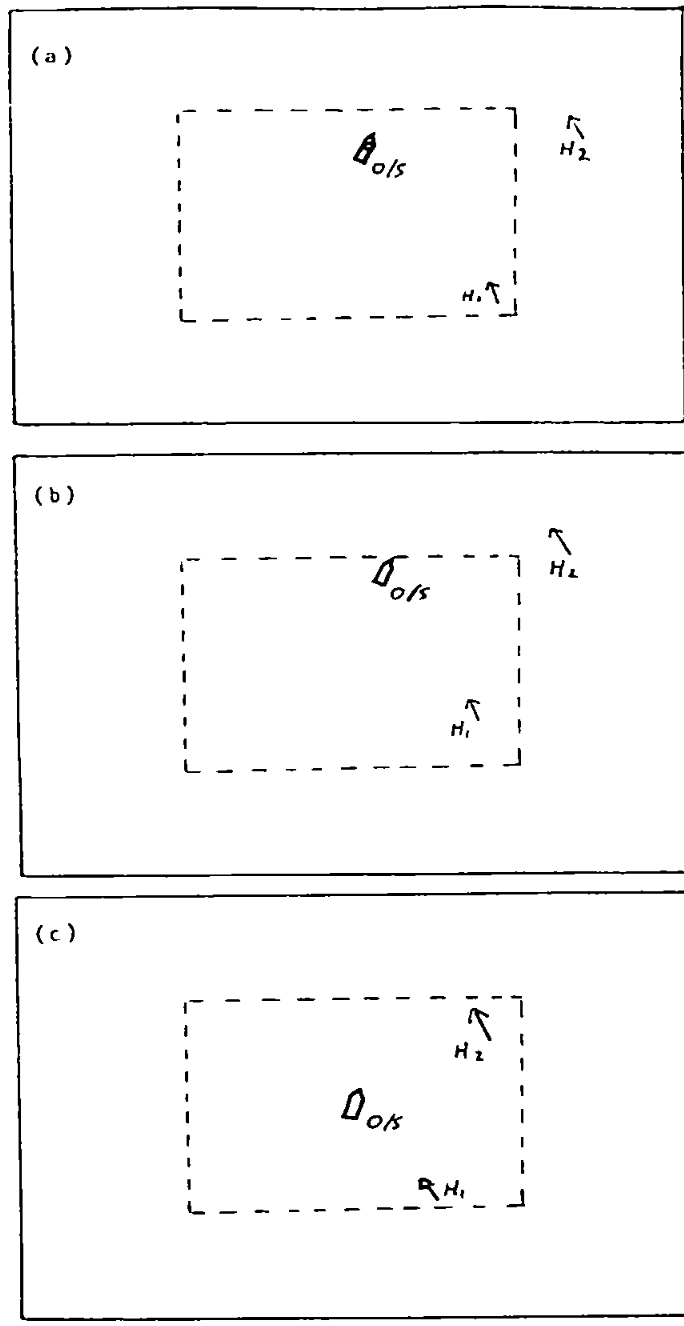

Figure 5. 'Rolling Sea' display format, keeping own-ship in centre section.
(dotted line is not shown on screen)
(a),(b),(c) at 2-minute intervals

e Acorn Archimedes RISC-OS is a multitasking vironment, permitting a number of tasks to nction pseudo-simultaneously. Messaging otocols enable communication between tasks, ving scope for interdependence between parate applications. Time-initiated 'events' y be triggered on a one-off or periodic sis.

e nature of the task handled by this ES quires that it interacts with other telligent entities - in the final analysis, e masters of other ships. It is therefore propriate that the other ships in the mulated environment behave 'intelligently' - accordance with some reasoned strategy. ere is also a requirement for a simulated wn-ship', which implements the decisions of e ES and in return provides the own-ship sitional and behavioural data expected by e ES.

e target application of this simulated vironment - the Expert Collision Avoidance stem - functions as one task under RISC-OS. simple own-ship simulator runs as a separate plication, 'slaved' to the ES by control ssages, and returning status information as quired. Another application, termed 'hazard ip', may be replicated any number of times ubject to memory limitations) to give a mber of potential hazards. Each task is itialised by selection of the appropriate on, and provided with all necessary rameters (position, course and speed for ips, domain size and performance aracteristics for ES) through dialog boxes ilored to this purpose.

tivities of the various vessels are nchronised by regular messages broadcast ound the system by a separate 'ARPA' module response to regular triggering by the terval timer. 20-second intervals are rrently simulated by 1-second steps, i.e. xreal-time; this speed-up factor may be duced as increased sophistication of the ES ts greater demands on the processor.

e main function of the ARPA application dule is to receive 'radar' data messaged om each hazard vessel, and pass on to the ES ch data for all vessels within radar range. e. the ES will not be conversant with the tails of all hazards currently active, only ose within specified radar range (e.g. 15 utical miles). 'Radar' data on own-ship is kewise returned to all ships within range.

s previously observed, it is necessary that zard vessels act in accordance with reasoned rategies. With this in mind, the 'hazard ip' application incorporates a simplified rsion of the ES, enabling it to function cording to the collision avoidance gulations. In order to test emergency noeuvring by the ES, each hazard ncorporates an optional manual override, lowing the operator to control the vessel nto a required situation.

e difficulty in developing an environment th an intelligent response is to know where stop. For example, should the ES take account of the possible reaction of Hazard 3 to the likelihood of avoidance manoeuvres by Hazard 2 around Hazard 1?? It has been decided not to try to second-guess the likely response of one hazard to another (as opposed to their response to own-ship); in consequence, the simplified ES in each hazard only responds to the threat from own-ship, not from other hazards. This highlights a basic rationale of the test environment: there is no merit in making it more sophisticated than the target ES can make use of for test purposes.

A new hazard may be introduced at any time by the operator, either within or outside radar range. If it is within range, data on that vessel is automatically passed to the ES by the ARPA module. If it is out of range, the ARPA module records its presence but does not communicate data on this vessel to the ES. If a hazard passes from within radar range to outside that range, or is recognised by the ARPA as increasing its separation from own-ship (whilst out of range), then that hazard application module is automatically closed down, and takes no further part in the simulation exercise.

The simulation environment, then, comprises an own-ship simulator and an irregular stream of hazard vessels. As each hazard comes into radar range the ES 'considers' it as a potential threat, and adopts an appropriate strategy if necessary; one strategy may handle just a single threat, or encompass a number of hazard vessels. As each hazard ceases to be a potential threat and disappears from radar 'vision', it is deleted from the ES knowledge base, and closed down as an application in the simulation, leaving room for further hazards. Naturally, only hazards within radar range are plotted on the ES graphic display, and then only if they pose a present or future potential threat (see section 5).

5. Strategy Evaluation and Optimisation

In a previous paper [2] the authors describe a technique for predetermining a safe time to manoeuvre in order to ensure required clearance of a hazard vessel. This Predetermined Safe Manoeuvring Time (PSMT) has been further refined to derive an optimal manoeuvring time (POMT) based on a cost function which takes into account deviation from track, loss of time, loss of speed, etc. An additional refinement is the capacity for re-evaluation of strategy on receiving data on a new hazard. As the ES is informed by the ARPA module of a new hazard, it first identifies whether that hazard is an immediate or potential future threat. This is done by consideration of a 'super-domain' about own-ship, based on the Davis model [op. cit. 7] (see figure 6).

If the new hazard is within this region, or its track will carry it into this region, then it is registered in the knowledge base as a potential threat, and the 'strategy optimiser' reviews its decisions in the light of this new data. This involves consideration of a large number of alternative manoeuvring times and rudder settings, identifying which combination ives' successful clearance with minimum cost; f a secondary manoeuvre is required to clear wo or more hazards, this is likewise onsidered from an optimisation point of view.

learly, it is not possible to evaluate a ultiplicity of options at each stage (from he initial viewpoint) without exponentially aising the number of look-ahead simulations. onsequently, the look-ahead optimisation onsiders only the first-stage manoeuvre in etail, at the same time ensuring that a way hrough to completion is feasible: further tages are evaluated in detail as they are pproached.

t is considered that this pattern closely ollows the technique adopted by human ship's asters in such situations: devise an fficient strategy for the immediate future, hich will also leave own-ship in a position o successfully negotiate through further azards; as the extended situation develops, volve further optimal strategies to cover ach new eventuality.

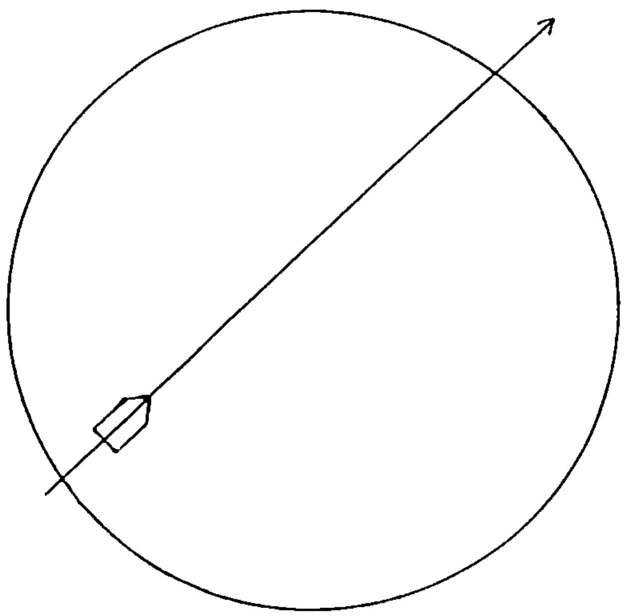

Figure 6. Super-Domain used to identify likely hazards. Own-ship to port, and well to rear, of centre. Typical radius: 5 nautical miles.

. Conclusions

he task of grouping/separating hazard vessels nto discrete 'encounters' is not susceptible o a simple rule of thumb. Any multi-ship ituation is fairly fluid, having the capacity o move from one grouping to another in a hort time, with regard to the problem of ormulating avoidance strategies.

The approach adopted by the authors has met with marked success in a large number of simulated encounters; namely formulation of an optimal strategy based on the most threatening hazard, with due regard in any strategy formulation (both in feasibility and costing) to other hazards which are, or may become, involved in the encounter. Such an approach avoids the question of separating or bunching vessels, the optimisation technique inherently operating an indeterminate (or 'fuzzy') boundary between encounters with successive hazards.

It is anticipated that the principles outlined in this paper may be extended without modification to encompass fixed hazards - notably coastal features - as well as other vessels. Any combination of such fixed and mobile hazards will be managed by an extended rule base under the optimisation logic as currently implemented.

7. References

1. Blackwell, G.K. and Stockel, C.T., 1988
   Simulation of Ship Encounters Using Multiple Linked Intelligent Systems.
   Proceedings SCSC89, Austin, Texas pp766-771

2. Blackwell, G.K. and Stockel, C.T., 1990
   Strategic Planning for Intelligent Real-Time Management.
   Proceedings SCSC90, Calgary, Alberta, Canada, pp694-699

3. International Marine Organisation, 1981
   International Regulations for Preventing Collisions at Sea.

4. Colley, B.A., 1985
   Computer Simulation of Marine Traffic Systems.
   Ph.D. Thesis, Plymouth Polytechnic (CNAA)

5. Blackwell, G.K., Colley, B.A., and Stockel, C.T., 1988
   An Intelligent Knowledge-Based System for Marine Collision Avoidance.
   Proceedings, Conference on Maritime Communications and Control, Institute of Marine Engineers, London. pp167-174

6. Goodwin, E.M., 1975
   A Statistical Study of Ship Domains.
   Journal of Navigation. Vol. 28, pp328-344

7. Davis P.V., Dove M.J. & Stockel C.T. 1983
   'A Computer Simulation of Marine Traffic Using Domains and Arenas'.
   Journal of Navigation. Vol. 33, pp215-222

SCS The Society for Computer Simulation
PO Box 17900, San Diego, CA 92177
Phone: (619) 277-3888, FAX: (619) 277-3930 SCS

Paper 056

Enhancing Simulation Utility
1991 Summer Computer Simulation Conference (SCSC91)
Hyatt Regency Hotel (Inner Harbor), Baltimore Maryland
July 22-24, 1991

SCSC 1991
Baltimore, Maryland
July 22-24, 1991

AN INTELLIGENT INTERACTIVE ENVIRONMENT FOR A MARITIME REAL-TIME EXPERT SYSTEM

Blackwell G.K., Rangachari J., Stockel C.T.
Department of Computing
Ship Control Group
Polytechnic South West
Plymouth, Devon, ENGLAND

. Introduction

he task of marine collision avoidance equires a combination of various technologies nd competences. Initially there is, of ourse, the need for appropriate control echanisms and technical skills to steer one's wn vessel (own-ship) along a predetermined rack, initiating and terminating turning anoeuvres at such points as may be required: his further presupposes an in-depth knowledge f the handling characteristics of own-ship. econd, the ability to detect the presence, peed and course of any potential hazard is ssential, implying the use of radar plus the xpertise to interpret incoming information nd apply it as required - an Advanced Radar 'lotting Aid (ARPA) may be used to perform art of this function. Third, the requisite easoning skills and experience, coupled with knowledge of the statutory regulations, hould lead to a strategy for avoiding any hreatened collision. Fourth, continued igilance whilst implementing that strategy is equired to ensure successful completion, ncluding modification of plans if any hazard essel executes a threatening manoeuvre.

t follows that any computer-based expert ystem for marine collision avoidance must ncorporate all of these features, including acilities for linking in and interpreting ata from any relevant sensor (e.g. radar). It ust also be capable of presenting its onclusions or recommendations in a meaningful nd easily-assimilated form. Any change in ircumstances (e.g. a hazard changing course) hould immediately precipitate a review of the lanned strategy, and if necessary clear and rompt notification of the revised plan.

Such a system specification implies a concomitant requirement for a staged development process, and a thorough testing regime at each stage. This is not a trivial consideration, since the finished system (and hence its progenitors) will carry major responsibility for safety of marine vessels. The degree of testing required, if conducted in a 'live' environment, would prove extremely costly and probably highly inconvenient. A seagoing vessel of substantial size (perform-ance characteristics of small craft are notoriously unreliable) would need to be fitted out with instrumentation and sensors, and held available for those times when test runs were required. At such times, the expert system would need to be transferred on board the ship, and the vessel taken out to some point where the traffic was such as to provide a realistic test of the new features (the alternative, a sea-going software house, raises different, but equally daunting, problems). Furthermore, suitable encounter situations would have to be contrived to test the system, either with other vessels which just happen to be in the vicinity, or with vessels specifically chartered for the purpose; the former could be less than popular (especially if a test of emergency manoeuvring at close quarters is required), the latter simply escalates the cost of the exercise.

Clearly, there is a requirement for a simulated test environment, to facilitate land-based development of the system to an advanced stage. To be in any way effective, that environment must offer all of the features, and challenges, which may be faced by the system in a real-life situation. It must also be susceptible to combining controlled or guided elements (to set up specific situations) with random elements which test details of system functioning. Interfacing of the test environment to the expert system should correspond as closely as possible to the 'live' situation, and permit monitoring of the test environment alongside full interaction with the system itself.

The authors have been concerned with the development of an expert system for marine collision avoidance. A substantial part of this work has involved the provision of a suitable test environment. Early work, based on an Atari ST microcomputer, used multiple processors to simulate vessel action [1]. This paper describes the use of an Acorn Archimedes RISC (Reduced Instruction Set Computer) system to emulate the hazards encountered on a sea passage. The multitasking capability of the RISC-OS operating system is used to model the actions of a variable number of vessels simultaneously. Control functions for each vessel permit the construction of a virtually unlimited variety of encounter situations, whilst inherent 'intelligence' in each hazard vessel ensures its conformity with collision avoidance protocols (unless manually directed otherwise).

2. Maritime Collision Avoidance

The task of marine collision avoidance is, in the first instance, encompassed by the International Regulations for Avoiding Collisions at Sea [2]. These regulations cite specific rules to cover specific types of encounter situations at sea. Such situations divide broadly into three categories:

(1) head-on;
(2) crossing;
(3) overtaking.

The nature of the encounter is identified by the position of one vessel with respect to another, vis-a-vis the navigation lights of each vessel. These navigation lights cover three separate sectors around a vessel, as shown in figure 1. Figure 2 illustrates a variety of encounter situations, as defined by the lights.

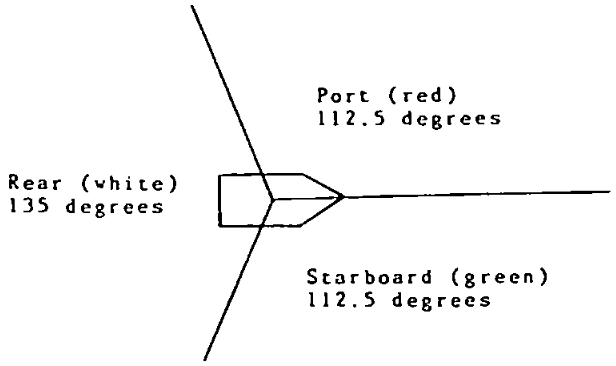

Figure 1. Sectors subtended by a ship's navigation lights.

The regulations are adhered to by every responsible mariner, and have the force of law in most countries. However, they are not totally definitive in terms of decisions and actions, and in many cases must be interpreted by mariners in the light of experience and discretion. For example, the master of a vessel on a collision/ near-miss course is expected to take action 'in good time' to avert disaster: this may be interpreted as anything from 10 to 30 minutes before closest point of approach (CPA). Likewise, a vessel manoeuvring to avoid another should indicate its intention by making a 'clear turn' (10-30 degrees), and maintaining a 'safe distance' from the other vessel (how long is a piece of string??).

Such flexibility (or vagueness) leads to differing responses by different masters, or even by the same master under different circumstances: a 'safe distance' when carrying a cargo of tractor parts may not be quite so safe with a hold full of volatile chemicals; the luxury of a 20 degree 'clear turn' in the open sea may not be so practical (or necessary to make the point) in a crowded seaway such as the Dover Strait. Clearly decisions on such factors as time to take avoiding action depend crucially on the performance characteristics of the vessel concerned: a cross-channel ferry has a substantially faster response, and rather smaller turning circle than (for example) a Very Large Crude Carrier (VLCC).

These differences in interpretation may prove significant in particular situations, since it is possible for the nature of an encounter to change with time, giving latitude for masters of the vessels in contention to interpret the

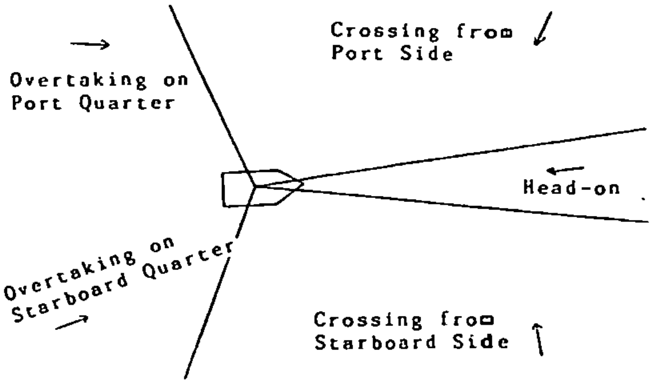

Figure 2. Types of ship encounters

same situation in two different ways. For example, figure 3 illustrates a situation which, evaluated 20 minutes before CPA would be seen as an overtaking encounter; 10 minutes before CPA it would be recognised as a crossing encounter. According to the regulations, responsibility for avoidance would rest with either one vessel or the other, depending on the interpretation applied. The regulations also direct that, once evaluated, the nature of a situation should be regarded as invariant for manoeuvring purposes (even if a changed scenario evolves, as in fig. 3); hence two masters could each deem the other responsible for avoiding, leading to the necessity for emergency manoeuvring at a late stage.

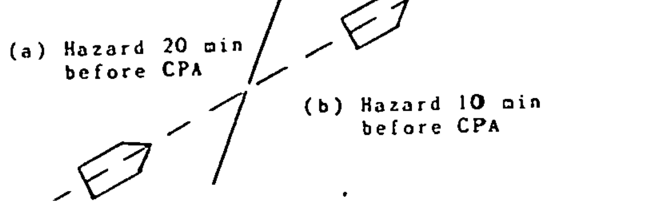

Figure 3. Encounter situation changing from hazard overtaking (hazard give-way) to crossing (own-ship give-way).

The regulations also allow for an early course change, not subject to fixed protocols, if a decision to manoeuvre is made well before any anticipated incident. This 'well before' can likewise be very loosely interpreted, leading to non-standard manoeuvres fairly late on in an encounter. Figure 4 illustrates the classic 'radar-assisted collision', in which master B has chosen a 'course change' to port (being somewhat to port of the other vessel's track) as opposed to the regulation 'avoidance manoeuvre' to starboard. The combination is a

Figure 4. Ship A: regulation turn to starboard
Ship B: 'course change' to port
Head-on near-miss becomes collision.

recipe for disaster, illustrating the need for clear adherence to well-defined rules. It has even been observed that some mariners will feel free to make an unrestricted 'course change' no matter how late on, if the vessels are not actually on a direct collision course - thus in some cases converting a near-miss into a disaster.

Previous research [3] has identified the major variants of the three standard encounter situations, and the strategies employed by mariners in those situations. These strategies have been formalised into a set of 'rules' which form the basis of an Intelligent Knowledge-Based System (IKBS) for marine collision avoidance [4]. Criteria relating to time for action, safe clearances and adequate turns are system parameters provided by the user, according to taste and circumstance. The knowledge base for the system comprises a combination of:

(1) 'static information', such as that above plus length, draught and performance characteristics of the vessel;

(2) 'dynamic information' such as current speed and course of own-ship plus radar data on any potential hazards in the vicinity; anticipated future developments include input of satellite position-fixing data and information on fixed navigational hazards from electronic chart digital information systems (ECDIS).

3. The Collision Avoidance Expert System

The system is built on the anti-collision regulations, filled out by observed good practice among mariners. Successful functioning of the system depends on two fundamental concepts relating to zones around a ship: the domain and the arena.

The domain is that area around own-ship which the master would wish to keep free of any potential hazards. Early research in this field identified the domain as disparate circle sectors [5] or alternatively a circle with own-ship off-centre [6], reflecting:

(a) the need for more sea room to starboard (the normal manoeuvring side); and
(b) comparatively less risk from hazards to the rear (closing speeds being generally less from behind).

For the purposes of this system, a circle centred on own-ship provides an adequate model for the domain; the domain radius, selected by the operator, is of course the representation of the 'clear distance' required by the regulations to be kept from other vessels.

The arena is that boundary around own-ship within which any hazard should be evaluated as a potential threat, and avoiding action initiated if appropriate [6]. This appraisal of a 'threat boundary' has been found to be over-simplistic, as it takes no account of different closing rates by various hazards. A better discriminator is a time-based boundary, triggered by time to domain infringement [3]. This RDRR (Range-to-Domain/Range-Rate) acts as follows: if a 10-minute RDRR is active (for example), a ship with a closing speed of 12 knots would come within this 10-minute threshold at a distance of 2 nautical miles from own-ship's domain boundary; a vessel closing at 6 knots would cross the RDRR threshold only 1 nautical mile from own-ship's domain.

The RDRR criterion has itself been found to have limitations, since every encounter situation is unique in terms of juxtapositions and relative velocities of vessels involved. Section 4 deals with the provision of a look-ahead facility which optimises selection of the appropriate RDRR for a particular encounter.

The expert system (ES) consists of an inference engine, or logical framework, applied to a rule base consisting of directives on preferred/alternative actions at various stages of different types of encounter. The rules are structured into a highly flexible state table, identifying different sequences of manoeuvring stages appropriate to different encounter types; more complex manoeuvres, including various emergency options, may be added to the rules by simply extending the 'binary tree' data structure which holds them. The rules function by reference to the knowledge base referred to previously. Each rule incorporates a decision function, usually involving extensive trigonometrical calculations based on velocity and displacement vectors, plus possibly intersections of turning arcs and domains. The ES operates on a 20-second time interval, evaluating the current scenario at each step and advising on preferred action.

The user interface is based on a Windows-Icons-Mouse-Pointer System (WIMPS) environment, the output display comprising three windows:

(1) a graphics window showing tracks of own-ship and any relevant hazards in the vicinity, in different colours; sequentially-numbered markers are used on each track to show corresponding positions in time;

The user interface is totally mouse-driven, with menus activated and options selected by mouse and pointer operation. It is envisaged that at a later stage there will be a requirement for text input to supplement the knowledge base; this, too, will be simplified by the use of dialog boxes (such boxes are already extensively used in simulation exercises on this system - see section 5).

4. Look-Ahead for Strategy Optimisation

In a previous paper [7] the authors describe a technique for predetermining a safe time to manoeuvre in order to ensure required clearance of a hazard vessel. This Predetermined Safe Manoeuvring Time (PSMT) has been further refined to derive an optimal manoeuvring time, based on a cost function which takes into account deviation from track, loss of time, loss of speed, etc. An additional refinement is the capacity for re-evaluation of strategy on receiving data on a new hazard. As the ES is informed by the ARPA module of a new hazard, it first identifies whether that hazard is an immediate or potential future threat. This is done by consideration of a 'super-domain' about own-ship, based on the Davis model [op. cit. 6] (see figure 6).

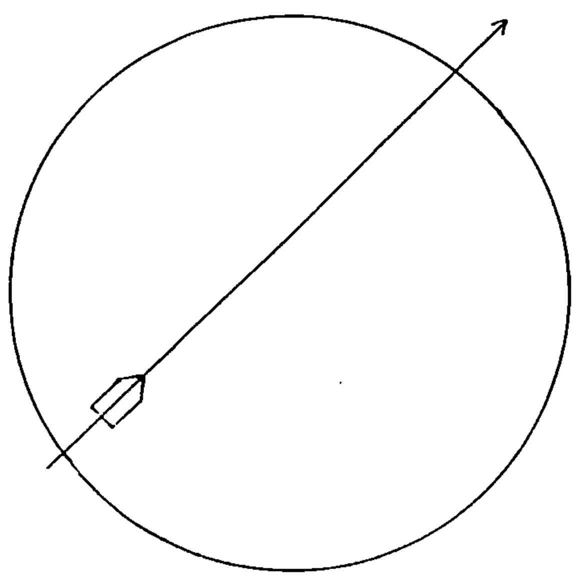

Figure 6. Super-Domain used to identify likely hazards. Own-ship to port, and well to rear, of centre. Typical radius: 5 nautical miles.

If the new hazard is within this region, or its track will carry it into this region, then it is registered in the knowledge base as a potential threat, and the 'strategy optimiser' reviews its decisions in the light of this new data. This involves consideration of a large number of alternative manoeuvring times and rudder settings, identifying which combination gives successful clearance with minimum cost; if a secondary manoeuvre is required to clear two or more hazards, this is likewise considered from an optimisation point of view.

Clearly, it is not possible to evaluate a multiplicity of options at each stage (from the initial viewpoint) without exponentially raising the number of look-ahead simulations. Consequently, the look-ahead optimisation considers only the first-stage manoeuvre in detail, at the same time ensuring that a way through to completion is feasible; further stages are evaluated in detail as they are approached.

It is considered that this pattern closely follows the technique adopted by human ship's masters in such situations: devise an efficient strategy for the immediate future, which will also leave own-ship in a position to successfully negotiate through further hazards; as the extended situation develops, evolve further optimal strategies to cover each new eventuality.

5. The Test and Development Environment

Clearly the Expert System in isolation cannot give any meaningful illustration of its performance in real-world conditions. In order for the system to function at all, there is a need for continuous streamed input of data relating to own-ship and other vessels in the vicinity. The former will be obtained from sensors attached to rudder, engines etc., the latter from radar (or ARPA). As indicated in section 1, such inputs are not a very practical proposition for ongoing system development. Hence a sophisticated simulated test environment has been evolved to match the needs of an increasingly sophisticated IKBS.

The Acorn Archimedes RISC-OS is a multitasking environment, permitting a number of tasks to function pseudo-simultaneously. Messaging protocols enable communication between tasks, giving scope for interdependence between separate applications. Time-initiated 'events' may be triggered on a one-off or periodic basis.

The nature of the task handled by this ES requires that it interacts with other intelligent entities - in the final analysis, the masters of other ships. It is therefore appropriate that the other ships in the simulated environment behave 'intelligently' - in accordance with some reasoned strategy. There is also a requirement for a simulated 'own-ship', which implements the decisions of the ES and in return provides the own-ship positional and behavioural data expected by the ES.

The target application of this simulated environment - the Expert Collision Avoidance System - functions as one task under RISC-OS. A simple own-ship simulator runs as a separate application, 'slaved' to the ES by control messages, and returning status information as required. Another application, termed 'hazard ship', may be replicated any number of times (subject to memory limitations) to give a number of potential hazards. Each task is initialised by selection of the appropriate icon, and provided with all necessary parameters (position, course and speed for ships, domain size and performance

2) a text window giving (on request) a summary of circumstances and recommendations to date, cross-referenced by sequential lettering to points on own-ship's track in the graphics window;

(3) an optional 'status' window, displaying current information on speed and course of own-ship, perceived speed and course of most threatening hazard vessel, separation and time to domain infringement, and present encounter state ('clear', 'ready' or stage of current manoeuvre).

Window (1) incorporates a 'pop-up' menu which permits: selection of large/small scale; 1-mile grid on/off (superimposed on display); provision/update of review information in window (2). The contents of window (2) are initially limited to brief statements of fact and advice; this window may be expanded to full screen width (covering window (1)) to allow fuller detail to be given on each stage.

Much consideration has been given to the best format for the graphics window, in terms of the position and motion of own-ship within this window: should the display be Head-up or North-up; should own-ship be static at the centre and relative motion of all other objects shown? Clearly, Head-up display implies a rotational transformation of all other graphical data as own-ship executes a turn. The question of disorientation requires careful consideration, balancing objective/subjective viewpoints against (for example) distortion of true ship tracks by a display of relative motion. Possibly the most significant factor is the ongoing nature of the task; simple continuous plotting of ship tracks in absolute motion will quickly lead to own-ship being 'lost' off the edge of the display.

The solution proposed, and in the process of implementation, is the plotting of absolute ship tracks (showing true motion) using a variant of the 'rolling road' principle: a boundary (not displayed) is fixed 25% of the way in from each edge of the window (see figure 5); as own-ship reaches this boundary, the whole display is relocated back to a position such that own-ship is at the centre of the display. In this way, at least the 50% of the display around own-ship (horizontally and vertically) at any time is maintained intact, giving continuity and minimising disorientation. It is accepted that periodic 'jumps' in the display may require a degree of 'acclimatisation', but this is considered of minor significance by comparison with either of the two alternatives: relative motion gives a distorted view of manoeuvres, leading to a totally false appraisal of any situation; full window wraparound (off the side - back on the other side, off the top - back on the bottom, and vice versa) could be extremely untidy without being very helpful.

The review text marks a departure from conventional expert systems technology, in that the observations and advice given are time-constrained. The logical analysis of 'this is so and that is so, therefore the following is true' provides an inadequate model for such time-based decision processes.

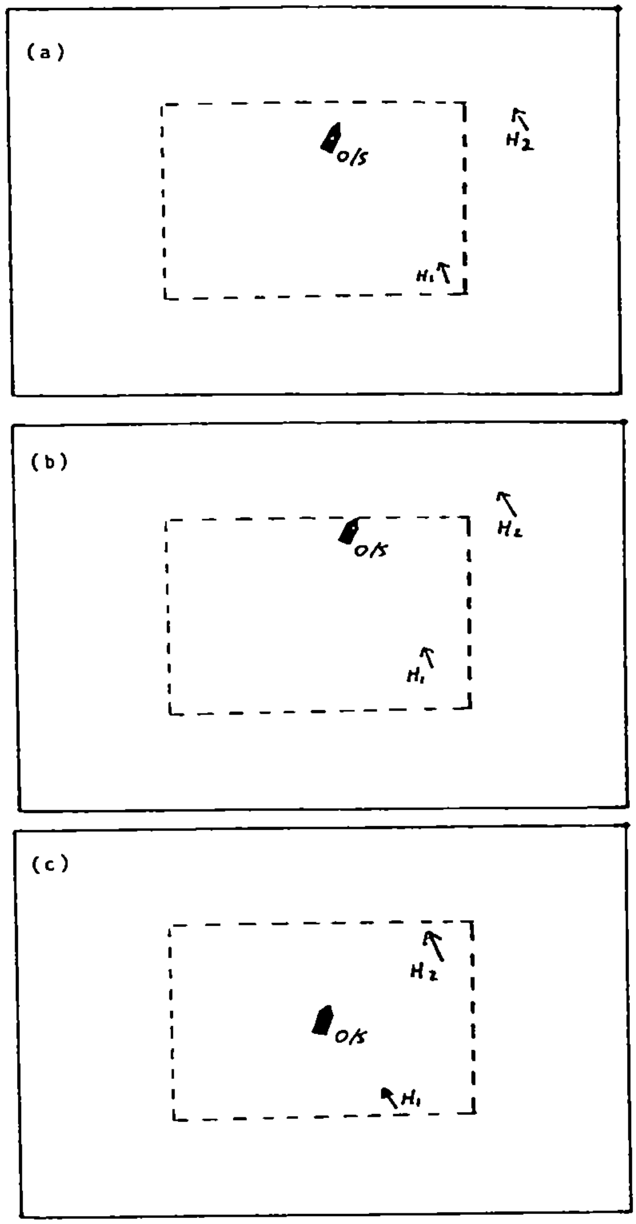

Figure 5. 'Rolling Sea' display format, keeping own-ship in centre section.
(dotted line is not shown on screen)
(a),(b),(c) at 2-minutute intervals

The corresponding logic pattern is of the form 'this was so three minutes ago, so-and-so has occurred since, indicating that the following will be true in twelve minutes - unless such-and-such happens, in which case the consequences will be ...' etc. In short, the added dimension of time broadens the whole decision process, making any clear resumé of that process a very demanding task. This is particularly true in the light of the requirement to present pertinent information in a succinct, easily-digested form. Hence the format of short, pithy comments with scope for expansion on request. The review window may also be scrolled up and down, to view the complete time-spectrum of events.

haracteristics for ES) through dialog boxes ailored to this purpose.

ctivities of the various vessels are ynchronised by regular messages broadcast round the system by a separate 'ARPA' module n response to regular triggering by the nterval timer. 20-second intervals are urrently simulated by 1-second steps, i.e. 20xreal-time; this speed-up factor may be reduced as increased sophistication of the ES puts greater demands on the processor.

he main function of the ARPA application module is to receive 'radar' data messaged from each hazard vessel, and pass on to the ES such data for all vessels within radar range. i.e. the ES will not be conversant with the details of all hazards currently active, only those within specified radar range (e.g. 15 nautical miles). 'Radar' data on own-ship is likewise returned to all ships within range.

As previously observed, it is necessary that hazard vessels act in accordance with reasoned strategies. With this in mind, the 'hazard ship' application incorporates a simplified version of the ES, enabling it to function according to the collision avoidance regulations. In order to test emergency manoeuvring by the ES, each hazard incorporates an optional manual override, allowing the operator to control the vessel into a required situation.

One difficulty in developing an environment with an intelligent response is to know where to stop. For example, should the ES take account of the possible reaction of Hazard 3 to the likelihood of avoidance manoeuvres by Hazard 2 around Hazard 1?? It has been decided not to try to second-guess the likely response of one hazard to another (as opposed to their response to own-ship); in consequence, the simplified ES in each hazard only responds to the threat from own-ship, not from other hazards. This highlights a basic rationale of the test environment: there is no merit in making it more sophisticated than the target ES can make use of for test purposes.

A new hazard may be introduced at any time by the operator, either within or outside radar range. If it is within range, data on that vessel is automatically passed to the ES by the ARPA module. If it is out of range, the ARPA module records its presence but does not communicate data on this vessel to the ES. If a hazard passes from within radar range to outside that range, or is recognised by the ARPA as increasing its separation from own-ship (whilst out of range), then that hazard application module is automatically closed down, and takes no further part in the simulation exercise.

The simulation environment, then, comprises an own-ship simulator and an irregular stream of hazard vessels. As each hazard comes into radar range the ES 'considers' it as a potential threat, and adopts an appropriate strategy if necessary; one strategy may handle just a single threat, or encompass a number of hazard vessels. As each hazard ceases to be a potential threat and disappears from radar 'vision', it is deleted from the ES knowledge base, and closed down as an application in the simulation, leaving room for further hazards. Naturally, only hazards within radar range are plotted on the ES graphic display, and then only if they pose a present or future potential threat (see section 4).

6. Conclusions

The test environment described here provides a thorough basis for validation of the Expert Collision Avoidance System, prior to actual sea trials. A continuous, but irregular, stream of hazards corresponds to real-life circumstances for such decision-making. Each hazard carries the ability for intelligent avoidance in its own right, but with scope for overriding such decisions, thus creating the type of 'rogue ship' or 'vessel in distress' situation encountered in reality.

It is acknowledged that actual sensory input may carry 'noise' or inaccuracies not included in this simulated environment; this is a possible area for enhancement of the multi-tasking test environment developed so far. The rule structure in the Expert System will also be modified to take account of such possible degradations of input data.

7. References

1. Blackwell, G.K. and Stockel, C.T., 1988
   Simulation of Ship Encounters Using Multiple Linked Intelligent Systems.
   Proceedings SCSC89, Austin, Texas pp766-771

2. International Marine Organisation, 1981
   International Regulations for Preventing Collisions at Sea.

3. Colley, B.A., 1985
   Computer Simulation of Marine Traffic Systems.
   Ph.D. Thesis, Plymouth Polytechnic (CNAA)

4. Blackwell, G.K., Colley B.A., and Stockel, C.T., 1988
   An Intelligent Knowledge-Based System for Marine Collision Avoidance.
   Proceedings, Conference on Maritime Communications and Control, Institute of Marine Engineers, London. pp167-174

5. Goodwin, E.M., 1975
   A Statistical Study of Ship Domains.
   Journal of Navigation. Vol. 28, pp328-344

6. Davis P.V., Dove M.J. & Stockel C.T. 1983
   'A Computer Simulation of Marine Traffic Using Domains and Arenas'.
   Journal of Navigation vol.33, pp215-222

7. Blackwell, G.K. and Stockel, C.T., 1990
   Strategic Planning for Intelligent Real-Time Management.
   Proceedings SCSC90, Calgary, Alberta, Canada, pp694-699

APPENDIX D

LISTING OF CURRENT SYSTEM RULES

The following pages contain a computer printout of the rules currently used within the system. The rules are held in an (extendable) array of structures, each structure holding the relevant information for one node within the binary decision tree which comprises the rule base. The nature of each node structure is as follows:

(a) a pointer to a boolean decision function which determines which branch of the rule tree is followed;

(b) a text string to be displayed in the 'status' window whilst the rule is active (blank in the case of 'transient' rules, which immediately action further rules);

(c) a sub-structure for each of the branches ('yes' and 'no'), holding information relating to subsequent action should the boolean function return 'true' or 'false' respectively; each sub-structure consists of

- (i) a reference pointing to the next rule to be actioned;
- (ii) rudder position (port, starboard, ahead) for next stage;
- (iii) indicator for deferred or immediate action (1 or 0);
- (iv) indicator for review text (1 if text to be added);
- (v) two lines of review text, if appropriate.

```
*
/*  FILE CONTAINING RULE DEFINITIONS FOR SHIPS APPLICATION
*   Created  14/5/90  - Updated 21/5/90, then again 27/5/90
*
*/
£define FULL_DOM TRUE
£define MIN_DOM FALSE
£define NO_ACTION 99
£define USUAL TRUE
£define EMRGNCY FALSE
£define JUMP FALSE

typedef struct               /* Info for each YES or NO branch in rules */
{
  int next_rule;             /* array ref. for next node */
  int change;                /* rudder position for next stage */
  BOOL defer:1;              /* defer move to next rule */
  BOOL label:1;              /* add letter to review, and on track */
  char *review[2];           /* decision text for review */
} Option;


typedef struct               /* One node/state in rule structure */
{
  BOOL    (*fptr)();         /* pointer to decision function */
  char    *stat;             /* status line info while rule active */
  Option  yes;               /* relevant info if TRUE  */
  Option  no;                /* relevant info if FALSE */
} Enc;


static char AllClear1[] = "Course now clear.";
static char AllClear2[] = "Back on initial heading";
static char Blank[] = "";

/**** FOLLOWS DECLARATIONS OF BOOLEAN DECISION FUNCTIONS FOR RULES ****/

BOOL all_clear(PROTAGS);   /* Hazard within BALLPARK? If so, calculate psmt */

BOOL in_arena(PROTAGS);    /* Is hazard within psmt range yet? */

BOOL tst_hdon(PROTAGS);    /* Is this a head-on encounter? */

BOOL chekturn(PROTAGS);    /* Has a clear turn been effected yet? */

BOOL turn_over(PROTAGS);   /* has turn cleared domain ? */

BOOL tst_mtB(PROTAGS);     /* Is hazard abaft of beam, port side? */

BOOL chekover(PROTAGS);    /* Is own-ship back on set course? */

BOOL tst_upass(PROTAGS);   /* Is own-ship in hazard's overtaking sector? */

BOOL tst_pas0(PROTAGS);    /* Are two ships on near-parallel courses? */

BOOL cleared(PROTAGS);     /* Safe to return to course (If haz ON course)? */

BOOL safe_return(PROTAGS); /*Will Turn to return to course infringe domain */
```

```
BOOL no_intersect(PROTAGS, double turn_cir_rad, double turn_cir_posx,
                  double turn_cir_posy, double *x, double *y);
                                /*Does own turning circle cut haz domain? */
BOOL tst_pas1(PROTAGS);    /* (Own-ship) Overtaking from port side? */

BOOL tst_pas2(PROTAGS);    /* OK to parallel-up and overtake ahead, stbd? */

BOOL chekpas2(PROTAGS);    /* Own-ship now on parallel course to hazard? */

BOOL tst_hepass(PROTAGS); /* Is hazard in own-ship's overtaking sector? */

BOOL tst_hedone(PROTAGS); /* Is way clear, and hazard back on course? */

BOOL look_emergency(PROTAGS); /* Decide on time at which emergency action */
                         /* should be taken                              */

BOOL tst_rdrr_emergency(PROTAGS); /* Is it time to check for possible */
                                  /* emergency                        */

BOOL tst_action(PROTAGS);         /* Will emergency action be required ? */

BOOL tst_h_action(PROTAGS);       /* Has hazard taken action ? */

BOOL tst_emergency(PROTAGS);    /* Minimum domain in danger - emergency ? */

BOOL tst_ucros(PROTAGS);          /* Is hazard to starboard of own-ship? */

BOOL chekcross(PROTAGS);  /* Hazard now to port of own-ship? (REVIEW INF?)*/

BOOL tst_crsB(PROTAGS);  /* Hazard 15+ degrees off port bow, & way clear? */

BOOL dizzy(PROTAGS);                    /* TEMPORARY EMERGENCY MANOEUVRE */

/******** FUNCTIONS USED BY DECISION FUNCTIONS ARE DECLARED BELOW ********/

extern void get_psmt(PROTAGS);          /* Simulation to get psmt */

extern void get_t_emergency(PROTAGS);        /* Simulation for emergency */
                                            /* action if necessary */
extern BOOL look_ahead(PROTAGS, cost_struct *cost,BOOL type, int stage,
                       double rudder, double psmt, double req_turn);
                                       /* Common look ahead routine for */
                                       /* psmt and emergency  */

extern BOOL yes_turn(PROTAGS, double rudder); /* Will turn clear domain ? */
                                              /* for get psmt routine */

static double chekangle(double angle);           /* Puts angle in 0->2*Pi */

static double anglegap(double angleA, double angleB);  /* Diff'ce, 0->Pi */

static double fixangle(double xfactor, double yfactor);/* Vector: 0->2*Pi */

static void findttg(PROTAGS, BOOL fulldom);        /* Time to domain breach */

static BOOL oncourse(Own_Ship *own, double track); /* heading = 'track'? */

/*****************------------------------------------------*****************/
```

```
/*** FOLLOWS NODE STRUCTURES FOR RULE BASE ***/
/* NB STATUS LINE COMMENT IS Blank FOR RULE NEVER STOPPED AT */

Enc rule[] =             /* Array of rules */
{
   /* £0             ANY HAZARD WITHIN RANGE YET? IF SO, SET PSMT */
  {all_clear, "clear",
   {1, STET, 0, 1, {"Hazard within range.", "POMT set at @P mins."}},
   {0, STET, 1, 0, {Blank, Blank}}},

   /* £1             WITHIN PSMT LIMIT YET? */
  {in_arena, "ready",
   {2, STET, 0, 1, {"@T mins to anticipated", "Domain infringement."}},
   {1, STET, 1, 0, {Blank, Blank}}},

   /* £2             HEAD-ON ENCOUNTER, OR SOME OTHER TYPE? */
  {tst_hdon, Blank,
   {3, STBD, 1, 0, {"Head-on encounter.", "Turn@A/6->stbd."}},
   {6, STET, 0, 0, {Blank, Blank}}},

   /* £3             COMPLETED HEAD-ON STAGE A? */
  {turn_over, "meet_A",
   {4, AHEAD, 1, 1, {"@C+ deg turn. No threat", "to domain. Hold course."}},
   {3, STET, 1, 0, {Blank, Blank}}},

   /* £4             COMPLETED HEAD-ON STAGE B? */
  {tst_mtB, "meet_B",
   {5, PORT, 1, 1, {"Hazard 90+ degs to port","Alter-back."}},
   {4, STET, 1, 0, {Blank, Blank}}},

   /* £5             COMPLETED HEAD-ON STAGE C? */
  {chekover, "meet_C",
   {0, AHEAD, 1, 1, {AllClear1, AllClear2}},
   {5, STET, 1, 0, {Blank, Blank}}},

   /* £6             OWN-SHIP OVERTAKING, OR HAZARD (OR Xing)? */
  {tst_upass, Blank,
   {7, AHEAD, 0, 0, {Blank, Blank}},
   {22, STET, 0, 0, {Blank, Blank}}},

   /* £7             NEAR-PARALLEL,  OR OTHER, OVERTAKING? */
  {tst_pasO, Blank,
   {8, STBD, 1, 0, {"Overtaking, nr-parallel", "course. Turn@A/6->stbd."}},
   {11, STET, 0, 0, {Blank, Blank}}},

   /* £8             COMPLETED PASSING_O STAGE A? (NR-PARALLEL, ->STBD) */
  {turn_over, "passO_A",
   {9, AHEAD, 1, 1, {"@C+ deg turn. No threat", "to domain. Hold course."}},
   {8, STET, 1, 0, {Blank, Blank}}},

   /* £9             COMPLETED PASSING_O STAGE B? */
  {cleared, "passO_B",
   {10, PORT, 1, 1, {"Safe to resume initial", "heading. Alter-back."}},
   {9, STET, 1, 0, {Blank, Blank}}},

   /* £10            COMPLETED PASSING_O STAGE C? */
  {chekover, "passO_C",
   {0, AHEAD, 1, 1, {AllClear1, AllClear2}},
   {10, STET, 1, 0, {Blank, Blank}}},
```

```
  /* £11           PASS BEHIND FROM PORT, OR PASS FROM STARBOARD? */
 {tst_pas1, Blank,
  {12, STBD, 1, 0, {"Overtake from haz port", "quarter. Turn@A/6->stbd"}},
  {15, STET, 0, 0, {Blank, Blank}}},

  /* £12           COMPLETED PASSING_1 STAGE A? (BEHIND PORT->STBD) */
 {
turn_over, "pass1_A",
  {13, AHEAD, 1, 1, {"@C+ deg turn. No threat", "to domain. Hold course."}},
  {12, STET, 1, 0, {Blank, Blank}}},

  /* £13           COMPLETED PASSING_1 STAGE B? */
 {cleared, "pass1_B",
  {14, PORT, 1, 1, {"Safe to resume initial", "heading. Alter-back."}},
  {13, STET, 1, 0, {Blank, Blank}}},

  /* £14            COMPLETED PASSING_1 STAGE C? */
 {chekover, "pass1_C",
  {0, AHEAD, 1, 1, {AllClear1, AllClear2}},
  {14, STET, 1, 0, {Blank, Blank}}},


  /* £15           PARALLEL-UP & OVERTAKE AHEAD, OR BEHIND STBD->PORT? */
 {tst_pas2, Blank,
  {16, STBD, 1, 0, {"Overtake ahead to stbd.", "@A/6->stbd, parallel-up"}},
  {19, PORT, 1, 0, {"Overtake from haz stbd", "quarter. Turn @A/6->port"}}},

  /* £16           COMPLETED PASSING_2 STAGE A? (PARALLEL-UP, STBD SIDE) */
 {chekpas2, "pass2_A",
    {17, AHEAD, 1, 1, {"Courses now parallel.", "Hold course."}},
    {16,  STET, 1, 0, {Blank, Blank}}},

  /* £17           COMPLETED PASSING_2 STAGE B? */
 {cleared, "pass2_B",
  {18, PORT, 1, 1, {"Safe to resume initial", "heading. Alter-back."}},
  {17, STET, 1, 0, {Blank, Blank}}},

  /* £18           COMPLETED PASSING_2 STAGE C? */
 {chekover, "pass2_C",
  {0, AHEAD, 1, 1, {AllClear1, AllClear2}},
  {18, STET, 1, 0, {Blank, Blank}}},


  /* £19           COMPLETED PASSING_3 STAGE A? (BEHIND STBD->PORT) */
 {turn_over, "pass3_A",
  {20, AHEAD, 1, 1, {"@C+ deg turn. No threat", "to domain. Hold course."}},
  {19, STET, 1, 0, {Blank, Blank}}},

  /* £20           COMPLETED PASSING_3 STAGE B? */
 {cleared, "pass3_B",
  {21, STBD, 1, 1, {"Safe to resume initial", "heading. Alter-back."}},
  {20, STET, 1, 0, {Blank, Blank}}},

  /* £21           COMPLETED PASSING_3 STAGE C? */
 {chekover, "pass3_C",
  {0, AHEAD, 1, 1, {AllClear1, AllClear2}},
  {21, STET, 1, 0, {Blank, Blank}}},
```

```
 /* £22            HAZARD OVERTAKING OWN-SHIP? (ELSE CROSSING ENC.) */
{tst_hepass, Blank,
 {27, AHEAD, 1, 0, {"Overtaking encounter.", "Own-ship stand-on."}},
 {23, STET, 0, 0, {Blank, Blank}}},

 /* £23            CROSSING:- OWN-SHIP GIVE-WAY or STAND-ON? */
{tst_ucros, Blank,
 {24, STBD, 1, 0, {"Crossing - haz to stbd.", "Turn@A/6->stbd"}},
 {27, AHEAD, 1, 0, {"Crossing encounter.", "Own-ship stand-on."}}},

 /* £24            COMPLETED CROSSING STAGE A? */
{chekcross, "cross_A",
 {25, AHEAD, 1, 1, {"@C+ deg turn. No threat", "to domain. Hold course."}},
 {24, STET, 1, 0, {Blank, Blank}}},

 /* £25            COMPLETED CROSSING STAGE B? */
{tst_crsB, "cross_B",
 {26, PORT, 1, 1, {"Haz 15+ degs to port of", "initial hdg. Alter-back"}},
 {25, STET, 1, 0, {Blank, Blank}}},

 /* £26            COMPLETED CROSSING STAGE C? */
{chekover, "cross_C",
 {0, AHEAD, 1, 1, {AllClear1, AllClear2}},
 {26, STET, 1, 0, {Blank, Blank}}},

 /* £27  CALCULATE  TIME TO TAKE EMERGENCY MANOEUVRE */
{look_emergency, "standon",
 {28, AHEAD, 1, 0, {"Emergency RDRR set","at @R minutes"}},


 {27, STET, 1, 0, {Blank, Blank}}},

   /* £28   TIME TO CHECK EMERGENCY */
{tst_rdrr_emergency ,"alert",
 {30, AHEAD, 0, 0, {Blank,Blank}},
 {29, STET, 0, 0, {Blank, Blank}}},

/* £29            HAZARD TAKING ACTION ? */
{tst_h_action, "alert",
 {32, AHEAD, 1, 0, {"Hazard diverted", "No Emergency"}},
 {28, STET, 1, 0, {Blank, Blank}}},


/* £30   IS ACTION NECESSARY */
{tst_action,"emrgncy",
 {31,STBD, 0, 1, {"Emergency situation","Altering Course"}},
 {32, STET, 0, 1, {"Near Miss situation", "Continuing on course"}}},

/*£31   RUN ROUND IN CIRCLES!! */
{dizzy,"emrgncy",
 {0 ,AHEAD, 1, 1, {AllClear1,AllClear2}},
 {31, STET, 1, 0, {Blank, Blank}}},



 /* £32            HAZARD COMPLETED MANOEUVRE? */
 {tst_hedone, "standon",
  {0, AHEAD, 1, 1, {AllClear1, AllClear2}},
  {32, STET, 1, 0, {Blank, Blank}}}
};

/****************-------------------------------------------****************/
```

APPENDIX E

EXTENDED BIBLIOGRAPHY

In addition to the references previously cited, the following works are of relevance in the context of the developments described in this thesis. These works are listed in various categories, followed by a summary of the relevant content.

[el]	Agogino A M, Russell S and Guha R 1988	Sensor fusion using influence diagrams and reasoning by analogy: application to milling machine monitoring and control. Artificial Intelligence in Engineering: Diagnosis and Learning (Ed. J.S.Gero). Elsevier Computational Mechanics Publications, Amsterdam, 1988, 333-357
[e2]	Zrimec T and Mowforth P 1988	An experiment in generating deep knowledge for robots. Artificial Intelligence in Engineering: Robotics and Processes (Ed. J.S.Gero). Elsevier Computational Mechanics Publications, Amsterdam 1988, 21-33
[e3]	Abdelgalil E M 1978	Shipping casualties and ship's domains. Marine Traffic Service - Proceedings of the Third International Symposium, Liverpool, 1978, 95-107

[e4] Coldwell T G — Encounters in restricted waters. Journal of Navigation, Volume 37, 271-278

[e5] Curtis R G 1977 — An analysis of the danger of ships overtaking. Institute of Mathematics and its Applications. Conference on Mathematical Aspects of Marine Traffic. Pub. Academic Press, 1979, 175-191

[e6] Degré T and Lefèvre X 1981 — The manoeuvring room concept - applications to collision avoidance problems. Symposium on Vessel Traffic Services, Bremen, 1981, Volume 2, 144-162

[e7] Hollingdale S H 1961 — The mathematics of collision avoidance in two dimensions. Journal of Navigation, Volume 14, 243-261

[e8] Hollingdale S H 1977 — Can mathematics help the mariner? Institute of Mathematics and its Applications Conference on Mathematical Aspects of Marine Traffic. Pub. Academic Press, 1979, 23-34

[e9] Lamb W G P — Calculation of the geometry of ship collision zones. Journal of Navigation, Volume 42, 298-305

[e10] Taylor D H 1990 — Uncertainty in collision avoidance manoeuvring. Journal of Navigation, Volume 43, 238-245

[e11] Zhao-lin W 1984 — An alternative system of collision avoidance. Journal of Navigation, Volume 37, 83-89

[e12] Cannell W P 1981 — Collision avoidance as a game of coordination. Journal of Navigation, Volume 34, 220-239

[e13] Cockcroft A N 1984 — Collisions at sea. Safety at Sea, June 1984, 17-19

[e14] Colley B A, Curtis R G and Stockel C T 1984 — A marine traffic flow and collision avoidance computer simulation. Journal of Navigation, Volume 37, 232-250

[e15] Corbet A G 1972 — Collision avoidance at sea, with special reference to radar and radio-telephony. Journal of Navigation, Volume 25, 520-543

[e16] Habberley J S and Taylor D H 1989 — Simulated collision avoidance manoeuvres: a parametric study. Journal of Navigation, Volume 42, 248-254

[e17] Jingsong Z and Fen-cheng W 1991 — Methods of collision avoidance by radar. Journal of Navigation, Volume 44, 205-212

[e18] Jingsong Z, Fen-cheng W and Zhao-lin W 1992 — The development of ship collision avoidance automation. Journal of Navigation, Volume 45, 107-113

[e19] Cockcroft A N and Lameijer J N F 1990 — A guide to the collision avoidance rules (4th edition). Heinemann, Oxford, 1990

[e20] Biancardi C G 1988 — On a simplified mathematical model for an on-board manoeuvring simulator and its potential applications to ship control. International Shipbuilding Progress, Volume 35, no. 401, 1988

[e21] Dand I W 1987 — On modular manoeuvring models. Proceedings of the International Conference on Ship Manoeuvrability, London, R.I.N.A., Paper no. 8, April-May 1987

[e22] Kose K — On a new mathematical model of manoeuvring motions of ships. Proceedings of the Third International Conference on Marine Simulation, Rotterdam, 1984

[e23] Matthews R 1984 — A six degrees of freedom ship model for computer simulation. Proceedings of the Third International Conference on Marine Simulation, Rotterdam, 1984

[e24] Bryant D T, de Bievre A F M and Dyer-Smith M B A 1987 — Investigating human factors in marine casualties. Proceedings of the Nautical Institute Seminar: Navigation and the Human Factor, 10 December 1987, 67-86

[e25] Hopkin V D 1990 — Some human factors aspects of single manning. Journal of Navigation, Volume 43, 343-352

[e26] Pearson M 1987 — Navigational safety and the human factor. Proceedings of the Seminar on Navigational Safety and the Human Factor, NI & RIN, December 1987, 1-7

[e27] Reeve P E 1987 — Ship automation: some implications for manning, personality, performance and safety. Seaways, January 1987, 10

[e28] Broadfoot G, Everett R, Halbesma H and Mulder B 1988 — ECDIS/ECDB operational in the Netherlands. Hydrographic Journal, October 1988, 50, 13-19

[e29] Burns R S, Dove M J, Evison J L, Stockel C T and Tapp N J 1985 — Improving the accuracy of own-ship's position on the Viewnav Electronic Chart. Proceedings of RIN & NI Seminar - Navigating With Electronic Charts, December 1985, 49-57

[e30] Eaton R M 1987 — Electronic charts - minimum data set for safe navigation. Seaways, March 1987, 11

[e31] Kerr A J 1990 — ECDIS: the state of play. Navigation News, September/October 1990, 18-19

[e32] Grabowski M 1990 — Decision support to masters, mates on watch and pilots: the piloting expert system. Journal of Navigation, Volume 43, 1990

[e33] Sugisaki A M, Imazu H, Tsuruta S and Inaishi M 1988 — Expert system for coping with collision avoidance and the ex post facto at sea. Navigation (Japan), Issue 97, 33-38

[e34] Inaishi M, Imazu H and Sugisaki A M 1991 — An approach to solving of collision avoidance problem based on a new concept: circumstantial judgement neural networks. Journal of Japan Institute of Navigation, Volume 85, 1-8

[e35] Coenen F P 1989 — A rule based collision avoidance system. Ph.D. thesis, CNAA (Liverpool Polytechnic)

[e36] Iijima Y, Hagiwara H and Kasai H 1991 — Results of collision avoidance manoeuvre experiments using a knowledge-based autonomous piloting system. Journal of Navigation, Volume 44, 194-204

References e1, e2 are concerned with sensor-based real-time expert systems. The subject matter of this thesis is considered to fit broadly into this area of study.

References e3-e11 relate to mathematical considerations in respect of potential collision situations and avoidance manoeuvring. This includes the concepts of domains and manoeuvring areas.

References e12-e19 deal with the topics of collisions and collision avoidance from a strategic viewpoint.

References e20-e23 consider the subjects of manoeuvring characteristics and manoeuvring mathematical models for various types of vessels.

References e24-e27 give consideration to the human factors involved in maritime collision situations. These are clearly fundamental to the style and content of any information presented to the mariner by a system such as the one described in this thesis.

References e28-e31 give an overview of the situation with respect to electronic charting. As highlighted in this thesis, substantial further developments are still needed in this area before a reliable standard is available for processing by any computerised navigation/guidance system.

In the paper cited as reference e32, Grabowski describes a rule-based system designed to aid pilots and officers on watch aboard large vessels in New York harbour. Substantial consideration is given to the pressures on those on watch, particularly the problem of 'information overload' when entering or exiting a busy harbour. The system is described as a decision support system, intended to present relevant information in a clear and meaningful form. The system uses a combination of frames, which hold information, and production rules which process that information and draw inferences, leading to pertinent advice. This Piloting Expert System was developed on a Symbolics 3670 LISP processor, using the Knowledge Engineering Environment. The user interface uses a mouse for input, with selection from hierarchical menu structures. Output is in the form of brief textual statements, giving recommended action. The system was evaluated by trials on a ship simulator, the subjects being senior cadets at the US Merchant Marine Academy. Results showed a substantial improvement in watch team performance as a consequence of using the system. A similar system is now under development for US tankers transiting the Gulf of Alaska.

In the paper cited as reference e34, Inaishi et al investigate the use of neural nets to address the collision avoidance problem. An expert ship operator provides the training data for the system. The output from the nets has been found to be in agreement with the expert's decisions as circumstances change. The paper also discusses the quality of the supervised training data. This work is a development of that described in the paper cited as reference e33, detailing a rule-based system of some 300 rules.

The Ph.D. thesis cited as reference e35 outlines the development of a rule-based system for marine collision avoidance. Coenen uses Prolog if-then production rules to develop a number of rule bases, for clear visibility, restricted visibility, emergency situations etc. The knowledge held by the system is held in frames. In these respects the system is similar in structure to that of Grabowski (above). The rules are centred around the collision avoidance regulations, interpreted by good practice as observed in the actions of experienced mariners. The rule bases may be extended by the simple addition of further if-then production rules. Clearly, extension of the rule bases leads to a corresponding (linear) increase in the search time for the rule applicable in a particular case. This contrasts with the binary tree form of the rule structure described in this thesis, in which the search time increases only logarithmically with the number of rules; in a real-time system, such a consideration could affect the expansion potential of the rule structure.

Coenen's work includes rules for multi-ship encounters, and allows for hampered vessels. The user interface consists of text string input, with a simple parser to select key words, textual output and a simulated radar display. The system has been evaluated by a number of experienced mariners on a variety of encounters run on a ship simulator, and found to deal satisfactorily with those situations; it has also been shown to deal competently with circumstances which led to two well-documented historical cases of shipping disasters. Some preliminary work is also included on accommodating chart data into the decision process, using a quadtree structure to hold relevant chart information.

Reference e36 cites advanced developments in a shipboard system which has been successfully tested on-board the training ship Shioji Maru, over a period of a month in congested traffic off Tokyo Bay. The system is linked in to an experimental control system, using ARPA for input of hazard data and an optical fibre local area network for sensor input and control systems. The system provides collision avoidance advice, and may optionally be used to directly control the vessel.

The system considers alternative options through a sequence of branches which between them form a complete manoeuvre. By evaluation of a cost function, it then selects the optimum route for hazard avoidance; the main criterion for a safe manoeuvre is preservation of the integrity of a circular domain centered on own-ship. In both of these respects this work mirrors the actions of the system which is the subject of this thesis. The output display includes preferred track superimposed upon an electronic chart, a screen showing current navigation data (speed, heading, information on target ship etc.) and a display of the rules used for the inference process. The system uses the OPS 5 software tool, a high-speed forward inferencing package.