Safety first—applying computer-based safety systems in the offshore environment

by C. J. Goring

August Systems

In this article the techniques used to design and apply triple modular redundant (TMR) computer-based safety control systems for the North Sea offshore production industry are described. The effects on the increasing safety awareness of the industry are discussed. In particular, references to the implications of the HSE guidelines and emerging international standards with respect to design, validation and training methods that should be applied through the life cycle of the project are reviewed.

Introduction

The already high awareness of safety in the North Sea production industry was given further impetus by the major disaster in 1988 when the Piper Alpha production platform was destroyed. The subsequent inquiry headed by Lord Callen defined a number of areas where significant safety improvement could be realised. Many of these improvements were being adopted even before the formal release of the Report. This article provides a description of the current direction of safety application philosophy with particular emphasis on the importance of training and documentation with respect to safety protection systems.

It is hoped that by use of this article and the various references, engineers working in the offshore environment may achieve a better understanding of current techniques and philosophies related to computer-based safety systems.

For those unfamiliar with offshore safety systems it is necessary to describe the type and function of the primary electronic safety systems found on a production platform.

Offshore safety systems

On an offshore production platform there are many types of safety and safety-related systems. Most package plant equipment will be provided with integral safety interlocks, an example of which is the safety interlocks associated with a gas compressor unit. In this article we concentrate on the primary electronic safety systems which are normally classified as the emergency shutdown system and the fire and gas detection and protection systems.

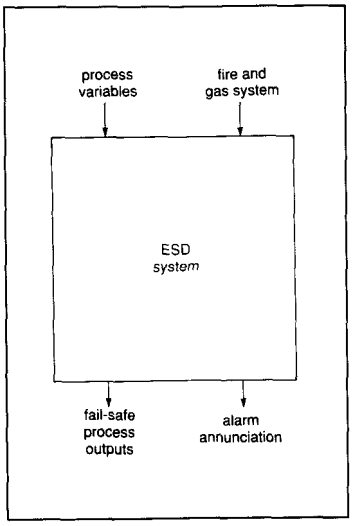

Fig. 1 Fire and gas system block diagram

The emergency shutdown system (or ESD system) has a primary function to protect the process plant from operating out of bounds and to protect life by shutting down the process plant if other localised hazards occur. Where package plant includes localised safety interlocks, the ESD system would normally be in addition to or as a controlling input to these localised packages. To enable the system to operate, process variables in the form of analogue and digital input signals, representing pressure flow, temperature level and position, are scanned on a regular basis and, according to predefined logic shutdown, valves are operated usually via an electrical solenoid. The basic input output interactions are shown for an ESD system in Fig. 1.

The emergency shutdown system is utilised to operate the various levels of shutdown of the process plant. These are normally classified from 0 to 3, where 0 would be the highest level of shutdown. Between different platform operators, minor variations occur as to which class of shutdown an item of plant falls into but, in principle, the following structures are followed.

The Class 3 process shutdown (PSD) is usually associated with an item of equipment or equipment package and would be utilised

during the normal startup and shutdown of the plant; additionally it would be brought into operation when local alarm situations occur causing a package to trip. Class 3 shutdowns are normally registered under operator control and the system additionally provides the safety interlocks to ensure that the correct sequence of process startup is followed when the plant is brought back online.

The Class 2 shutdown is also normally classed as a PSD or process shutdown. This class of shutdown is used to completely shutdown major plant functions or the complete process shutdown and would normally be called on during a localised out of bounds operation (e.g. local low-level gas leaks). The process can normally be restarted by the operators, immediately the hazard is cleared.

When a major hazard occurs, e.g. a major release of gas or a localised fire hazard, a Class 1 emergency shutdown would be initiated. This encompasses the organised complete shutdown of the platform processing facility, and is called on when the major hazard is detected. A complete black start of the process would now be required after a Class 1 shutdown, i.e. blowdown has occurred.

If the severity or extent of the hazard is such, a Class 0 or true emergency shutdown or alarm shutdown will be required; all power would now have been removed from the de-energise to trip valves instantaneously. Shutdown is complete for all utilities except for the emergency generators; fire pumps are normally automatically started when this level of shutdown occurs.

In the unlikely event that the hazard has led to a situation where the production platform has to be abandoned, a post abandon platform (PAP) shutdown is initiated from localised hardwired pushbuttons situated at the evacuation points. These pushbuttons remove the power to all process and utilities systems, including the emergency generators, but leave the battery-supported systems functioning for their required support design time, e.g. 3-5 hours minimum.

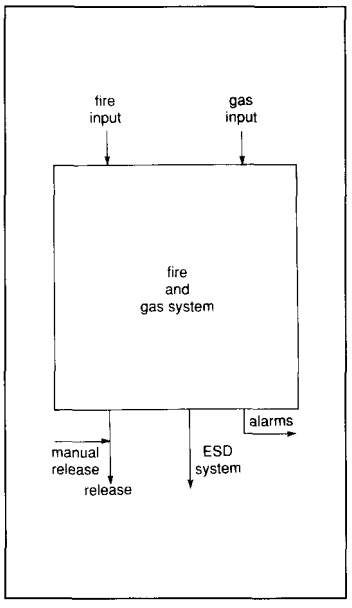

Fig. 2 ESD system block diagram

The fire and gas detection and protection systems form the principal systems for monitoring fire, smoke, combustible gas and toxic gas alarm conditions and producing automatic protection actions as appropriate. Their primary function is to protect the life of those personnel on the operating platform. Separate systems would normally be provided for the process areas and the accommodation areas with similar functions being accomplished by each. The production system cannot operate unless there is an approved fire and gas safety system in operation.

Dependent on the hazard and location, the fire and gas system would demand the ESD system to shutdown plant and, where appropriate, take control of HVAC dampers and fans to prevent smoke or gas hazard spreading. Damper and fan control would also be used after the hazard is cleared to remove smoke from ducts and vents.

Protection would be provided by automatic release of deluge or inert gas extinguishant systems (in the less environmentally friendly days this would have been Halon). The basic input/output interactions are shown for a fire and gas system in Fig. 2.

The detection and protection provided by a fire and gas system is usually instigated on an area basis. Typically a potentially hazardous area will be provided with an array of gas and fire detectors. Each provides an input into the fire an gas system. The system will provide voting logic to determine whether the detected hazard can be confirmed by two or more of the detectors. The actions taken by the fire and gas system will depend significantly on the extent of the hazard and its situation.

If a single detector in an area goes into alarm, this would normally only prompt the fire and gas system to provide the appropriate audible and visual alarm enabling the safety operators to investigate the possible hazard. When coincident detection occurs by more than one detector in an area, then a number of automatic actions may occur, mostly dependent on the location and the contents of the area in question.

The types of action taken by the fire and gas system include the opening and closing of dampers to control smoke, the release of extinguishant to control flames, the starting of fire pumps to allow deluge extinguishant control and providing outputs to the ESD system to ensure that plant and machinery are placed in a safe state.

Alarm annunciation by the fire and gas system is usually provided by either a hard-wired mimic and matrix panel or by a combination of VDU screen mimics and hard-wired area matrix control.

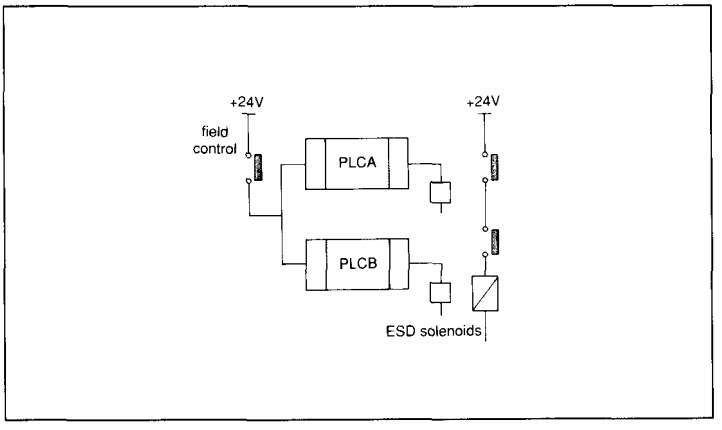

Fig. 3 Dual redundancy

Fault tolerance

For these critical types of safety system, *it is mandatory that if computer-based systems are*

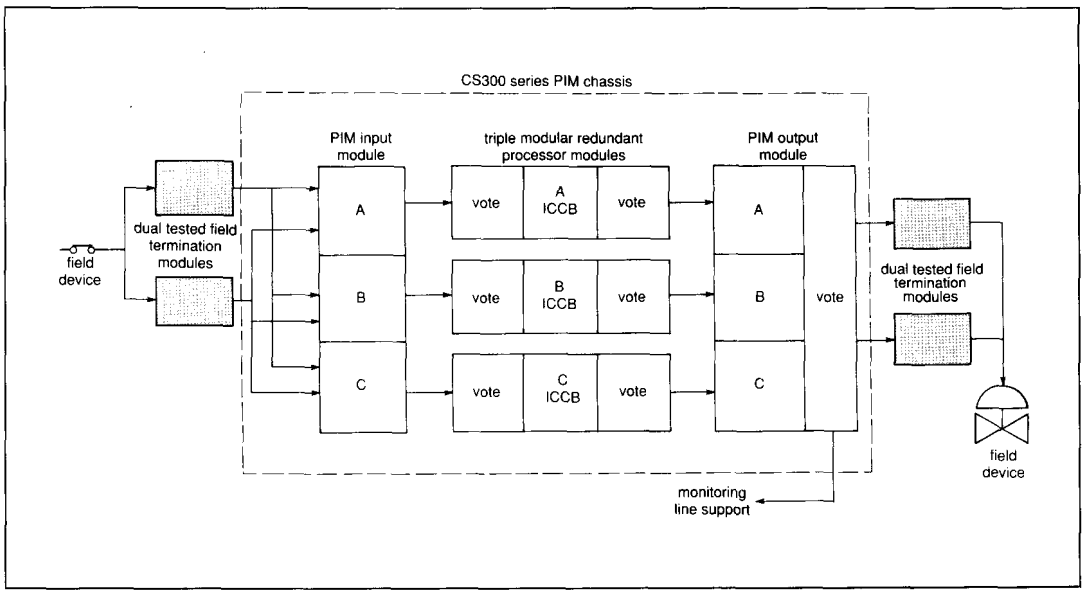

Fig. 4 TMR architecture

used then the application of fault-tolerant computer-based systems is utilised.

There are many mechanisms of achieving fault-tolerant systems, but to provide true fault tolerance, a minimum level of redundancy is required. The mimimum level for true fault tolerance is three. This can be explained as follows.

If we review a dual redundant system as shown in Fig. 3, it can be seen that all works well when both systems are operating correctly. However, when a fault occurs in one system, vital questions must now be asked before it is possible to determine the correct action to take. Typically these questions are:

1 Can the fault be determined or is there just a disagreement between the systems on what action to take?
2 Can the system take the correct safe action?
3 Is it possible to determine which of the two machines is correct?

As can be seen a number of difficulties with dual redundancy can certainly arise. Normal safe practice in ESD systems is to shutdown and in a fire and gas system to alarm but not to action. This will lead to more false trips in the case of an ESD system and to more false alarms in the case of a fire and gas system.

The answer has been to apply triple modular redundancy (TMR) and vote on a two out of three basis for all single faults. In the past this approach has never been cost effective but, with relatively low hardware prices and well engineered TMR systems, it is currently the most efficient approach to achieving high-availability safety systems.

True fault tolerance is not just tolerating a fault, it is the combination of being able to tolerate a fault, diagnose a fault to the replaceable part level and the ability to replace the faulty item online without shutting down or degrading the protection the system is providing.

The combination of TMR and software implemented fault tolerance (SIFT) uniquely provides these facilities. Referring to Fig. 4, we can see that each individual input from the field is taken to three separate input channels.

The three asynchronous processors read the associated input channels and then utilising a unique voting algorithm, through read-only communication links, validates the data between each processor. This data validation is performed every time an input is read, at certain times during the application logic and prior to all data output writes. Additionally a standard six element hardware output voter (Fig. 5) is fitted to each output card ensuring that the correct output action occurs.

The application of TMR/SIFT architectures can now provide extremely high levels of hardware availability.

Safety systems strive for the highest availability and this can only be achieved by providing true fault tolerance. To increase availability on all but the simplest systems, redundancy is required. It is by redundancy management that the highest level of system availability can be attained. System availability can be defined as the ability of the

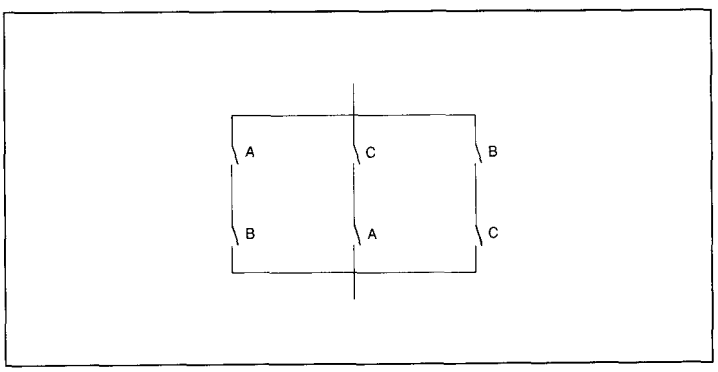

Fig. 5 Six element voter

Fig. 6 Safety matrix

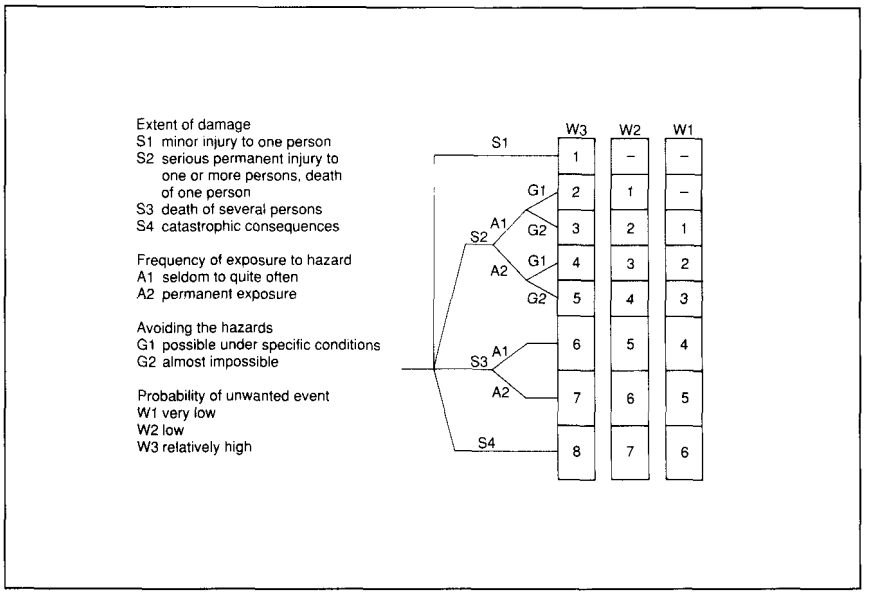

systems to respond correctly to demands at all times. It is calculated from the simple equation:

$$\frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}} \%$$

MTBF: mean time between failure (hours)
MTTR: mean time to repair (hours)

Certain fundamentals affect the MTTR significantly. With respect to system design, the time taken to find a fault is critical. Additionally the accuracy of the fault diagnostics is a primary function in keeping the MTTR to a minimum. The importance of proper training and thorough documentation in influencing the MTTR is covered in later sections of this article. In many instances unrealistically short times are quoted for MTTR. Considering the need to locate, obtain and then fit, a spare unit, figures of less than four hours would, in general, be inappropriate for a process plant.

As computer-based safety systems contain software, when producing the system availability, software, of course, cannot be ignored. Most estimates of software correctness (software does not fail; it may, however, contain errors) are not fully quantitative. There is a tendency to judge software harshly when providing system availability figures.

As with all designs, hardware and software errors can be introduced at various stages of the life cycle. *Unfortunately with software, the number of combination paths that can exist usually preclude the possibility for 100% test coverage*. For safety systems we therefore need to provide confidence factors to reassure the user of the operational integrity. A preferential order of core attributes should include:

1 *n* millions of hours of bug free operation in similar applications (where *n* is greater than 1).
2 Independent authority certification.
3 Quality software maintenance structures. ISO 9000-3 certification.
4 Application software independent verification and validation.
5 Quality tools and languages used in design.

International standards and their direction

Over the last few years the safety industry has been providing systems with full consideration of both the Health and Safety Executive PES I and II Guidelines and the EEMUA Document No. 160. These two documents have led the way internationally on safety guidelines and the principles contained within them are now being incorporated in the emerging international standards.

There are two current IEC working groups producing standards for programmable safety systems, these are:

- IEC 65A WG 9
- IEC 65A WG 10.

Working group 10 is concerned principally with the overall systems and general requirements hardware design, whilst WG 9 is providing guidelines and standards for the development of safe software. Additionally in the USA a committee and working groups under ISA have been generating complimentary standards under SP84. It is intended that these standards will be merged over the next one to three years to form the new ISO international standards for safe programmable systems.

In addition, certain national standards and certification authorities have evolved to become recognised internationally as acceptable bodies to approve safety systems. The most renowned and internationally accepted of these is currently the TUV authority in Germany. Certain of the TUVs have the expertise to certify software and hardware for operation in specified risk requirement classes.

The first major hurdle a user has to cross when adopting these standards is to establish the level of safety classification. This level needs to be defined at the outset in order for the emerging standards to be used effectively. This will always be a difficult area to be decisive in; however, with the use of independent auditing and safety committee decisions erring to the conservative, the correct (or best) classification can be achieved for each safety case. When defining the classification of a safety system it is detrimental to safety to classify all inputs and outputs of the system to the highest level as this will add to

the complexity of the design. To achieve the most effective and, in fact, the most safe approach to safety design, the highest classification areas should be constrained to the minimum which is deemed acceptable by the safety committee in accordance with the guidelines and, of course, at all times erring towards a conservative approach. A typical safety classification matrix is shown in Fig. 6.

When a system or part of a system is classified at the highest level, then all new design requirements should follow the highly recommended (HR) or recommended (R) approaches indicated by the emerging standards. It should be remembered that, for the highest class of safety, diversity is mandatory. The need to minimise the areas that require the highest degree of safety should therefore be recognised to ensure that maintenance risks are kept to a minimum.

It is a fact that a primary cause of hazardous operations occurs during the maintenance of, or change implementation of, a safety-related system. The addition of complex diverse architectures multiplies the possibilities of maintenance errors occurring. It is therefore imperative that the diverse architecture is both kept as simple as possible and the function it needs to cover as restricted as possible, thereby making the diverse section relatively straight forward to maintain.

Diverse implementation is approached differently in the two types of primary safety systems under discussion in this article.

For the emergency shutdown system having achieved an extremely high availability system using TMR and SIFT architectures, analysis needs to be performed on plant parameters to discover if an additional diverse electronics safety system is required. It is quite possible, dependent on the process and its location, that use of pressure relief valves, bursting disc and similar devices can provide the adequate levels of safety and diversity to the primary protection system. However, in certain applications no additional physical device can be applied, in which case a diverse system will be required in the highest safety classification areas. This will often be based on trip amplifiers and hard-wired relays and will be applied as an OR function with the TMR system to the trip valves. A typical system is shown in Fig. 7.

In the case of a fire and gas detection and protection system, the human operator invariably forms part of the diverse leg. This is certainly acceptable in the general case for this type of system, as the speed of reaction required for fire protection is normally measured in seconds and not milliseconds.

To achieve diversity the following additional functions must be provided by the fire and gas system:

1 independent/diverse alarm annunciation
2 hard-wired/independent/diverse protection release.

With the implementation of the above in conjunction with a TMR SIFT logic system the highest level of safety classification can be achieved in a fire and gas protection and detection system.

Probably the most common approach to diverse alarm annunciations from a fire and gas system is to bring every single fire zone input and gas detection input into a separate front end independent alarm module. This module normally contains no software and provides for fire inputs, alarm and fault annunciation and for gas inputs, gas level display plus alarm and fault annunciation.

There is, of course, a level of information interpretation needed to determine exactly which areas are in alarm; however, it should be recognised that the high availability systems normally operating the fire and gas system will probably never be out of operation. Nevertheless this is not an excuse for skimping on training of the hard-wired aspect of safety control.

To provide for a diverse mechanism for extinguishant release, a number of hard-wired release pushbuttons are provided on the fire panels to enable the safety operators to take the appropriate safety actions. These would connect direct power to items such as extinguishant release solenoids. To prevent accidental operation, a keyswitch enable and a flapguard on the release pushbuttons would normally be provided.

Life cycle support

At the time of commissioning the safety systems will be in a known supportable state, approved and certified by the appropriate regulating authority and ready to perform their safeguarding role.

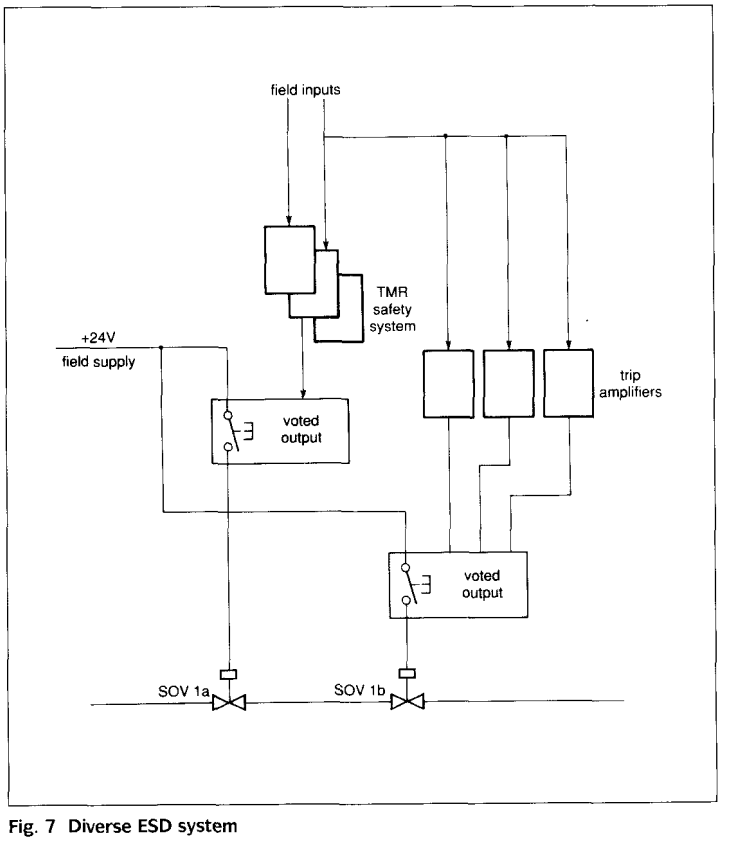

Fig. 7 Diverse ESD system

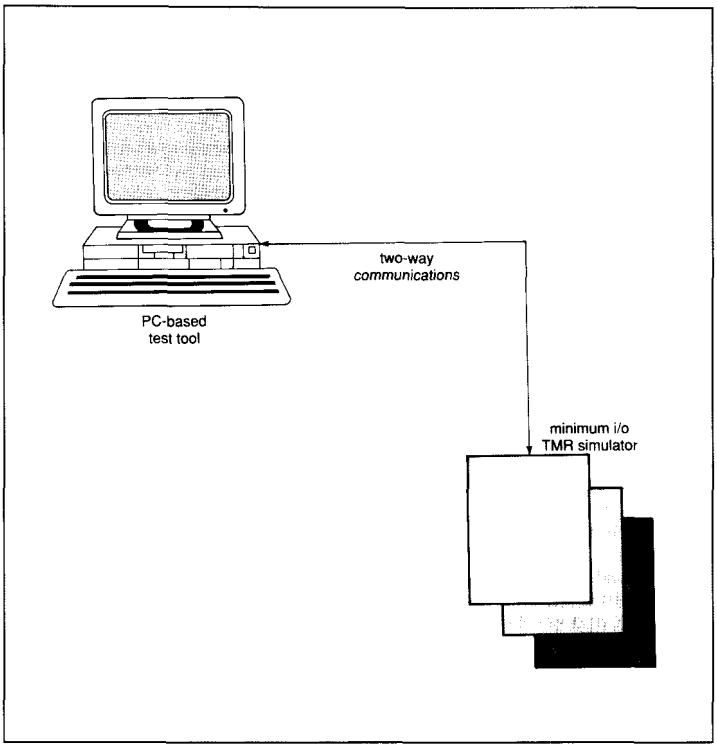

Fig. 8 Software application programming test unit

Without, however, the appropriate level of both operator and maintenance training and support, the efforts of the design engineering staff will have been unproductive. The life cycle support system must be seen to provide:

1 Thorough training to the operators for both normal and adverse operational conditions.
2 Training to the maintainers to develop a thorough understanding of the system and their safety role.
3 Complete, accurate and comprehensible documentation at all relevant levels.
4 Controlled and safe working environment with permit to work procedures under secure and documented control.
5 Full and thorough safety case analysis for all system changes post certfication.
6 Documented re-verification of all changes.
7 Documented re-validation of all system changes.
8 Tools to enable thorough re-validation.

To ensure the correct operation of the safety system during both its normal operation and maintenance condition, as well as during and after configuration changes, a well engineered simple design for the safety systems is a prerequisite. A TMR system provides such a simple to operate and understand system enabling virtually all aspects of the safety systems configuration to be viewed by the user as a simplex system.

During any period when system configuration has to change, due perhaps to plant development or additional safety criteria, the management of change with the view 'safety first' is of essence.

The safety case for the change must be independently reviewed and approved by the plant safety committee; each member of the safety committee must be fully aware of his or her responsibilities.

The design changes must be modulised and the verification of these modules should be achieved by independent personnel from the designers.

Full validation of the final redesigned safety systems must be completed before the system is brought back online.

To enable the expedient verification and validation of the safety system design changes, the use of independent verification and validation tools provides a significant aid in ensuring that the system is safe to bring back online.

An example of such a tool is shown in Fig. 8. The tool SAPTU (software application programming test unit) enables cause and effect actions to be entered into a database format and automatic pseudorandom testing for both positive and negative actions to occur.

Finally it is imperative that the operators and maintainers are regularly put through refresher courses, particularly concentrating on hazard response. This can be achieved to a level with the proper use of training simulators. However, perhaps the best approach is that adopted by one major offshore operator, who has built a full-scale operating replica of the safety system interface and, utilising a Vax computer, is able to generate hazard scenarios for the operators to respond to during their training and retraining course.

Conclusion

In a short article, it can only be possible to provide an indication of the levels to which safety engineering has now become an integral part of the total design of an offshore installation and in particular the design of the computer-based safety system.

The industry is actively pursuing the release and adoption of new standards and, through the formation of such concerns as the Safety Critical Systems Club, is stimulating the wider dissemination of valuable safety information.

The list of references includes a number of books and articles that the reader may wish to consult to broaden his understanding of safety-related computer systems.

References

1 'Programmable electronic systems in safety related application, Volumes I and II', Health & Safety Executive
2 'Safety related instrument systems for the process industries', Publication No. 160, Engineering Equipment and Materials User Association (EEMUA)
3 'Software for computers in the application of industrial safety related systems', IEC/TC 65A.3.1 WG9
4 'Draft: functional safety of electrical/electronic/programmable electronic systems: Generic aspects; Part 1: General requirements', IEC/TC 65A.9.1 WG10
5 'Fundamental safety aspects to be considered for measurement and control protective equipment', DIN 19250/89
6 'Principles for computerising safety related systems', DIN V VDE0801/90

Chris Goring is Managing Director of August Systems Ltd., 1–5 Kelvin Way, Crawley, W. Sussex RH10 2SE, UK.