The Parliamentary and Scientific Committee

INFORMATION FOR MEMBERS

From the

Parliamentary Office of Science and Technology

Briefing Note

20

January 1991

SAFETY CRITICAL SYSTEMS

From motor cars to aeroplanes, medical equipment to nuclear power stations, computers are increasingly used to control the highly complex systems on which we all rely. As systems become more sophisticated, so the importance of making computer control systems acceptably safe becomes paramount.

This briefing note reviews the science of safety-critical systems and related issues of potential interest to Parliamentarians.

COMPUTER CONTROLLED SYSTEMS AND SAFETY

Domestic gas bills for a million pounds are one thing, aeroplane crashes or runaway nuclear reactors are another. Yet such mishaps may all originate from a computer error, independent of the human operator and where proper system maintenance has been carried out. Such faults may arise because some of the modern-day tasks for computer systems demand such complex software (the set of written instructions that directs the computer's actions) that they are incapable of being fully tested.

The implications have been considered by Government and expert bodies. In 1986, the Advisory Council for Applied Research and Development (ACARD)[1] emphasised the importance of software in safety-related systems. The Health and Safety Executive has issued a number of generic guidelines on safety-related applications of computers. In 1987/88, the DTI commissioned a report from the Institution of Electrical Engineers (IEE) and the British Computer Society (BCS). Most recently, an Interdepartmental Committee on Software Engineering (ICSE) has led to two consultative documents being prepared in June 1990 under the 'SafeIT' initiative.

Nevertheless, there are no mandatory requirements which apply to all production or use of safety-critical computer systems, and there is a danger that, in the worst case, a potentially lethal product may be freely sold or used. Examples where computer errors caused or contributed to serious malfunctions are shown in Box 1 overleaf.

APPLICATIONS

The use of computers in safety-related applications is growing, particularly in the following areas:

(1) increasingly complex systems for recording data - e.g. medical records which contain thousands of millions of pieces of information;

(2) increasingly complex safety-related control systems, e.g. fly-by-wire aeroplanes such as Airbus, and the primary emergency shut-down system of the Sizewell B nuclear power station;

(3) in many simpler but more widespread applications, e.g. anti-lock braking systems for vehicles;

(4) in opening up new application areas - current developments range from improved commuter transport, such as the driver-less Docklands Light Railway, to 'robotic surgeons' to perform surgical operations at high precision and low cost.

The importance of getting software right is illustrated by estimates that UK industry suffers cost penalties of £500m p.a. on failures in software produced by a third party, with much higher losses if software developed in-house is included. Sales of software in safety-related applications in the UK are estimated to be around £500m per annum, growing at 20% p.a.

WHY DO ERRORS OCCUR?

The Complete System. The safety of a computer-controlled system must be considered on a systems-wide basis, i.e. the overall effect of the reliability of the hardware (e.g. of the silicon chips that constitute the 'brain' of the computer), the software or computer programs, and the way in which these work together and interact with the human operator. Error-free operation also depends on system designers incorporating some of the complexity of the 'real world' in their design. Given the complexity of the many tasks a computer must control (e.g. flight behaviour), a completely accurate specification may be difficult - let alone trying to predict every possible eventuality. Nonetheless, it is true that present computer-controlled systems are **safer** than the less sophisticated systems (whether controlled by earlier computers or human operator) they are replacing. Many industries (e.g. aviation, nuclear power) make extreme efforts to minimise risks.

[1] ACARD Report: 'Software: a Vital Key to UK Competitiveness', 1986.

Computer Hardware. The main difficulty lies in detecting any latent errors in the design of the silicon chip. Errors are not simple to find since **exhaustive** testing of a sophisticated chip is impossible. The chips at the heart of a computer consist of many thousands of components (transistors) that are wired together. Fully testing these would need every possible combination of every transistor to be tried out to ensure that the design was without fault - a task which would take countless years. In practice, chips are subjected to an extensive series of tests chosen to detect errors. Potentially serious errors can nevertheless remain hidden. In operation, computer hardware has a reliability that is predictable in the same sense that other every-day articles, such as motor cars, have a predictable reliability.

Computer Software. In contrast, errors in software are not predictable, making software unlike other common commodities[2]. Computer programs often run to millions of separate statements and the effects of a simple error in one of these may be far reaching. Again, the complexity of contemporary software does not permit complete testing. For instance, a modern avionic control system where the total (hardware and software) system failure rate must be no higher than 1 in a billion hours would require at least 10 billion hours of testing.

The Human Factor. Human error may also affect computer controlled systems in their design, construction, use and maintenance. The man-machine interface is a key to the safe operation of the system (e.g. system failure warnings may be ignored or misunderstood), and careful design is important here to ensure that the human operator (e.g. the pilot in the cockpit, the nurse operating medical equipment) will make the right decisions when faced with a potentially dangerous situation.

ELIMINATING ERRORS

Since the basic components of safety-critical systems cannot be made fault-free, a number of techniques are used to achieve the required levels of safety in the complete system. In operation, multiple control systems working in parallel are often used (Figure 1). A failure in one controller will not cause the entire system to fail as it is 'outvoted' by the other, correctly working, controllers. (In practice, other safety features would be incorporated, e.g. reversion to manual control). Sometimes, in very complex cases, separate design teams use different computers and different languages to design the system. The operator then uses all the systems in parallel, as in Figure 1. Other design aids include analysis of the potential hazards with which systems must cope in the real world, the use of formal mathematical proofs ('Formal Methods') of software, and pragmatic solutions based on sound engineering principles. The use of Formal Methods is one area where the UK has a considerable lead on other countries (including the USA and Japan). However, there is disagreement amongst experts over the use of Formal Methods in systems design. Some, whilst accepting that Formal Methods are not a panacea, claim that if all possible behaviour of a system cannot be expressed mathematically, the system may be potentially unsafe. Others consider that the methodology is not sufficiently mature for general use.

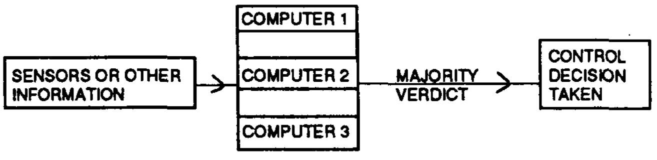

Figure 1 Enhanced Safety through Multiple Controls

As well as illustrating the lack of consensus among experts in safety critical systems, the range of techniques used implies that no single approach is likely to eliminate all errors. Several attempts have been made to establish **safety criteria** on an industry-by-industry basis, reflecting specific safety concerns. Even within industrial sectors, such criteria may be different. For instance in the transport sector, the control system of an aircraft has to provide continuous control over a range of complex functions (from adjusting fuel flow to protecting against fire). In contrast, the setting and verification of a line signal is a key safety feature in rail transport.

GOVERNMENT INITIATIVES

In 1990, the DTI issued the 'SafeIT' consultation document, with the aims of facilitating the development of international standards (generic and sector/application-specific), fostering the use of high-integrity computer systems in safety-related applications and heightening awareness throughout industry.

With regard to R&D, a 4-year Safety Critical Systems Research Programme is due to commence in 1991 with Government funds of £9m (DTI) plus £3m (SERC) available to support up to 50% of project costs. This collaborative programme aims to develop, formalise and disseminate enhanced understanding of software techniques used in various applications. Participation from industry, universities and regulatory bodies is expected. The research will be complementary to European collaborations, such as ESPRIT, already underway. The HSE is actively involved in developing international standards and research is also supported by the DTI through the British Standards Institute and at the National Physical Laboratory.

[2] **Typically, software is beset by 1000 mistakes per million opportunities. To put this figure in context, baggage handling at airports has a fault rate of 3000 per million opportunities, restaurant bills 10,300 faults per million, and airline fatalities, 4.3 per million.**

Box 1 Examples of Computer-Related Accidents

The autopilot of a passenger aircraft tried to hold an altitude of 41,000 ft after engine failure. As a result, the aircraft stalled causing the other engines to die. The aircraft lost 32,000 ft before control was regained by the pilots.

A rocket launch failure was traced to one program error - the statement read DO I = 1.10 instead of DO I = 1,10.

An error in the tactical software of a Fighter aircraft caused an uncontrollable spin, which made the aircraft crash.

A malfunctioning robot killed a Japanese mechanic.

Wiring harness fires in US motor cars were traced to a fault in the computer-controlled air suspension system.

A woman killed her daughter and tried to kill her son and herself after a computer error led to a false report that all three were suffering from an incurable disease

A patient died when a computer-controlled intravenous drip malfunctioned and supplied an overdose.

The 'SafeIT' document also supported the establishment of an Interdisciplinary Research Centre (IRC) for long term research work. Much of the research in safety-critical systems is at the leading edge involving several disciplines - from mathematics to psychology. However, as discussed below, the proposed IRC is not now to go ahead.

CURRENT ISSUES

Are Current Regulations Adequate?

Specific regulations cover safety-critical computer systems in certain sensitive industries (e.g. aviation and nuclear power). However outside of these special cases, only more general legal obligations apply. Both criminal and civil liabilities apply to software designers, suppliers and users. Applicable statutes are the Health and Safety at Work Act (1974) and the Consumer Protection Act- CPA (1987). The former confers responsibility on the designer and supplier both to ensure safe design and construction and for adequate testing and documentation, whilst the user is obliged to carry out reasonable checks, staff training and so on. The CPA (which also enacts the EC Directive on Product Liability of 1985) covers the general safety of consumer goods, which includes computer software but excludes one area where safety-critical software is beginning to be widely used, that of motor cars.

A major objective of the HSE generic guidelines was to stimulate various industries to produce their own guidance on specific applications. A number of bodies and industries have responded (e.g. the Institute of Gas Engineers and several chemical processing industries) but the burgeoning use of software means that many other areas are not so covered and, for many UK industries, the only means of quality control is via their own internal procedures. Some observers express doubts as to the effectiveness of such regulation.

On **standards**, the IEE/BCS found that those being formulated by the International Electrotechnical Commission (IEC) were well regarded in British industry and are likely therefore to find a wide acceptance. It is however important to ensure harmonisation of standards within and outside of the EC, to avoid cost penalties accruing from diverging UK and international standards. The UK's role within the relevant bodies (IEC internationally and CENELEC[3] in the EC) is thus critical.

The MoD has an obvious concern with safety-critical systems, not simply because weapons are inherently dangerous when out of control, but more in reflection of the growing complexity of many military systems. In 1989, MoD released two interim Defence Standards on safety-critical software and hazard analysis, which include the compulsory use of Formal Methods. Some industrialists fear that the Defence Standards may be applied *ipso facto* to civil work as well, which could be a source of concern for smaller organisations without the necessary human or financial resources. However those familiar with Formal Methods claim that their use can actually save money by preventing costly mistakes throughout software development.

Many in the software industry see the general provisions of existing statutes as sufficient. However, uncertainties remain, particularly over the liability aspects of software used in a process or product, and industries report difficulties in obtaining insurance cover. This can lead to suppliers using less software in their product lines as a defence against possible legal action. Not only is this approach likely to be inherently less safe, it may result in handicaps for an industry in which the UK is still a world leader. Consequently there is a need for an early resolution of these uncertainties - perhaps by drawing up generic, worldwide standards and regulations. This is not without technical difficulty, however. Some experts feel that more progress must first be made in understanding the scientific aspects of safety-critical systems in order that regulations are effective, and to avoid the danger of setting up arbitrary, unworkable or costly procedures.

Certification

Some advocate a certification system be applied to all safety-critical systems - either to the individual engineer, the employing organisation, or both. Certification of an individual implies that a degree of competence

[3] CENELEC is the European Committee for Electrotechnical Standardisation (Comite Europeene de Normalisation Electrotechnique).

can be defined and maintained, which may be difficult in a rapidly-advancing field. Nevertheless, the feeling exists that some recognisable status (e.g. 'chartered software engineer') is required. After all, computer-controlled medical equipment would only be used by highly qualified medical staff - in contrast, the software in the equipment may have been assembled by a complete novice. In this respect, the EC may introduce a quality initiative in 1991, which could lead to an EC directive to force software houses to comply with higher professional practice.

Some have noted that rigorous certification of any kind will lead to increased costs and divert scarce expert resources. The fear has also been expressed that the ultimate goal of a system supplier will become the award of a certificate, not the manufacture of a safe system. Ultimately, a certificate does not guarantee safety since the certification process itself may not have included all the important safety-critical issues.

Increasing Awareness

As computer controlled systems become more complex, the limit on safety may become the extent to which the computer conveys information to the human operator - the human-computer interface (e.g. that of pilots on advanced flightdecks). In such cases, there is some concern that very sophisticated systems may hold an unacceptable risk. Awareness of the safety implications of their products or services was shown by the IEE/BCS study to be lacking in some two-thirds of those surveyed. In recognition of this, 'clubs' have recently been set up under 'SafeIT' to increase awareness and foster technology transfer of generic standards between industrial sectors. However, without some form of regulatory framework, such mechanisms may not raise standards in all sectors since those who fail to appreciate the problems may not take part. This is of concern since economic pressures will inevitably increase the complexity of all products and services.

Another suggestion, which would require legislation, is that a formal requirement be introduced to report significant failures or 'near-misses' to a single central authority rather than to industry-specific bodies. The latter include the Civil Aviation Authority and the Institute of Aviation Medicine to whom aircrew can report, anonymously, problems experienced whilst airborne. The anonymity may help to offset the conflict that an employee may feel between an ethical obligation to society and a legal responsibility to an employer. The Fellowship of Engineering has recently issued guidelines for engineers facing such a dilemma.

One obvious difficulty is in defining a 'near miss' in industries other than those of transport. However, such a register would constitute a valuable source of information, over and above the possibility of saving lives by remedial action in the sectors affected (if necessary, worldwide). One illustration of unheeded warnings concerns the near-meltdown of the nuclear reactor at Three Mile Island. The same series of incidents which led to the emergency in March 1979 (and the evacuation of 160,000 people), had occurred twice before in October '77 and May '78.

Is The Research Base Adequate?

The proliferation of complex safety-related systems requires research support - in order to maximise protection of life and the environment, to maintain a strong UK presence in discussions on international standards and to reinforce the UK's capability in software production.

As noted above, the IRC proposed by the SERC is not now to proceed, despite the recommendation of a technical assessment panel to set up a centre based on York, Newcastle and Durham. A smaller activity was to be funded at York, with the possibility of converting to IRC status within two years. However, the current financial difficulties besetting SERC mean that this option may be in jeopardy. One UK company is seeking to set up a centre for dependable systems research, with links to academia. However, such an initiative is necessarily industry-specific and unlikely to sustain the wide-ranging research which the problems of safety-critical systems demand. Widespread dissemination of results may also be restricted on intellectual property grounds. Some therefore feel that the delay in establishing an IRC, which would be the first in the EC, will reduce UK competitiveness in this field.

Others argue that pressing safety issues exist now and that immediate practical solutions are required in the applications area, not longer term research. They point out that some UK industries have already produced elegantly structured solutions to meet their safety needs, based on developing sound engineering practices. In this respect the 'SafeIT' initiative is seen as providing a national focus for the development and dissemination of good practice in safety-critical technology. However, many in industry are concerned that the amount of financial support for the 'SafeIT' projects may be insufficient to meet the needs identified.

FURTHER READING

Additional details and background information are available from POST, 2 Little Smith St., London SW1P 3DL, tel: (071)-222-2688.

The ***PARLIAMENTARY OFFICE OF SCIENCE AND TECHNOLOGY*** has been set up by the Parliamentary and Scientific Committee to inform Parliamentarians on scientific and technological issues.