

Implementation of a Hardware/Software Platform for Real-Time Data-Intensive Applications in Hazardous Environments

Jesús F. Conde, Alberto García-Martínez and Angel Viña

CESAT, Open Real-Time Systems Center
Universidad de La Coruña
La Coruña (Spain)
E-mail: [jconde, alberto, avc]@cesat.es

Abstract

In real-time data-intensive applications, the simultaneous achievement of the required performance and determinism is a difficult issue to address, mainly due to the time needed to perform I/O operations, more significant than CPU processing time. Additional features need to be considered if these applications are intended to perform in hostile environments. In this paper we address the implementation of a hardware/software platform designed to acquire, transfer, process and store massive amounts of information at sustained rates of several MBytes/sec, capable to support real-time applications with stringent throughput requirements under hazardous environmental conditions. A real-world system devoted to the inspection of nuclear power plants is presented as an illustrative example.

1. Introduction

The achievement of the required time response and I/O throughput with the simultaneous satisfaction of strict timeliness requirements are essential issues to address in the implementation of hardware/software platforms capable to support real-time data-intensive applications. These applications can be characterized by the acquisition, transferring, processing and storage of large amounts of information in a predictable way.

Information acquisition and evaluation systems in fields such as aerospace, military or process control have necessarily to be designed and implemented according to the increasing demands of emerging real-time data-intensive applications. The continuous appearance of real-time data-intensive applications is being spawned by performance improvements in processors and peripherals, specially in technology fields demanding high-speed

information transmission rates and extremely large storage capacity under hard timing constraints.

Computer technology advances are fostering the continuous increase in CPU speed and the exploitation of parallelism, whereas improvements in the I/O subsystem are not occurring at a similar pace [6]. This fact is making data-intensive applications very I/O-limited, becoming the I/O subsystem the global performance bottleneck in the vast majority of cases. Large and sequential I/O operations are typical in these applications. Achievement of the desired performance and determinism demands the use of innovative hardware/software platforms.

Additional features need to be considered if these data-intensive applications are intended to perform under hazardous environments, where the presence of environmental factors (electrical noise, presence of radiation, extreme temperatures or variations of temperature, etc.) are of key importance for the correct operation of the system. While enhanced or redundant components (CPUs, disks, etc.) can be designed for the achievement of a correct behavior in such adverse circumstances, making the whole system immune to hostile environmental conditions would be too expensive or even completely unfeasible. Massive storage devices and human interaction interfaces are elements specially sensitive to the aforementioned threats: information can be easily corrupted in hazardous environments, and human presence is very restricted or completely forbidden in dangerous areas.

This paper focuses on the implementation of a hardware/software platform designed to achieve the high I/O, deterministic performance required in real-time data-intensive applications that should satisfy data flows at sustained rates of several MBytes/sec under hazardous environmental conditions.

The following section addresses the generalization of

the requirements found on real-time data-intensive applications obliged to perform in hazardous environments. The description of a hardware/software platform for real-time data-intensive applications designed to perform in hazardous environments is the subject of the third section. Section 4 reports our laboratory work on the implementation of a system devoted to the inspection of nuclear power plants. We will end with the conclusions and further work.

2. Requirements for real-time data-intensive applications in hazardous environments

The vast majority of emerging data-intensive applications are intended to acquire, transfer, process and store extremely large information volumes (in the Terabyte range) at rates of hundreds of Megabits/sec. The main features requested from a hardware/software platform aimed to support real-time data-intensive applications are high global I/O throughput, deterministic temporal behavior and extremely large storage capacity.

The necessity to perform under hazardous environments implies the physical separation of the platform in different nodes, in order to offer the intended safety and reliability. Consequently, the platform design must incorporate the adequate technology to provide the high-bandwidth, low-latency communications needed between the different nodes of the platform.

In order to protect human life and those parts of the platform that could be affected by hard environmental conditions, a practical and cost-effective solution can be achieved by locating the appropriate elements far enough from the hazardous zone, in a safe zone, communicating both areas by means of a point-to-point digital link.

These requirements force the incorporation of some kind of intelligence in the hazardous zone, like controlling cards, or even processing elements to cope with the application communication requirements. The distance to be covered should range from tens of meters to several kilometers, or even more. The issues related to data transferring must be carefully addressed in order to avoid that the communications link become the system's global bottleneck.

In addition, there are several features that the final platform should possess. Among them, the following two are of particular relevance:

- **Flexibility and Scalability.** The overall performance of the platform should be improved in parallel with future technological advances in any of its integrating components.
- **Openness.** The design of the hardware/software

platform should comply with the major applicable de jure and de facto standards in hardware, communications, real-time, application development, etc.

3. Description of the hardware/software platform

In order to meet the stringent requirements of real-time data-intensive applications subject to hard environmental conditions, heterogeneous components selected by considering their limits regarding I/O throughput, real-time capabilities, storage capacity, computing performance, communication efficiency and cost-effectiveness, must be integrated into the final platform.

The elements of a platform destined to support real-time data-intensive applications under hazardous environments are listed below. Figure 1 shows a graphical representation of a two-node platform.

- **Communications link.** A high speed link provides an efficient communication channel between the nodes. Fiber optic technology and robust data-link transmission protocols are a must to achieve high bandwidth and reliability. Upper level network protocols should not be used in order to obtain more efficiency. Simple, point-to-point, application specific communication protocols should be incorporated instead.
- **High capacity storage media.** The most important feature required from the storage subsystem is a high sustained I/O throughput, in order not to constitute the global bottleneck of the system. In addition, an extremely high capacity is needed. Current trends in high capacity storage systems are confirming that RAID's are the most suitable storage devices for real-time data-intensive applications. High performance standard interfaces with storage devices (SCSI-2, Fiber Channel) are of particular relevance in order to obtain the required storage throughput.
- **Real-Time Operating System (RTOS).** In each zone, the RTOS must provide facilities to achieve the required application performance, meeting the stringent data handling requirements in a predictable way. The RTOS will be obliged to perform I/O intensive time-constrained operations, in order to:
 1. acquire and transfer large amounts of data within specified time intervals;
 2. carry out the required CPU-bound information processing (i.e. visualization algorithms); and

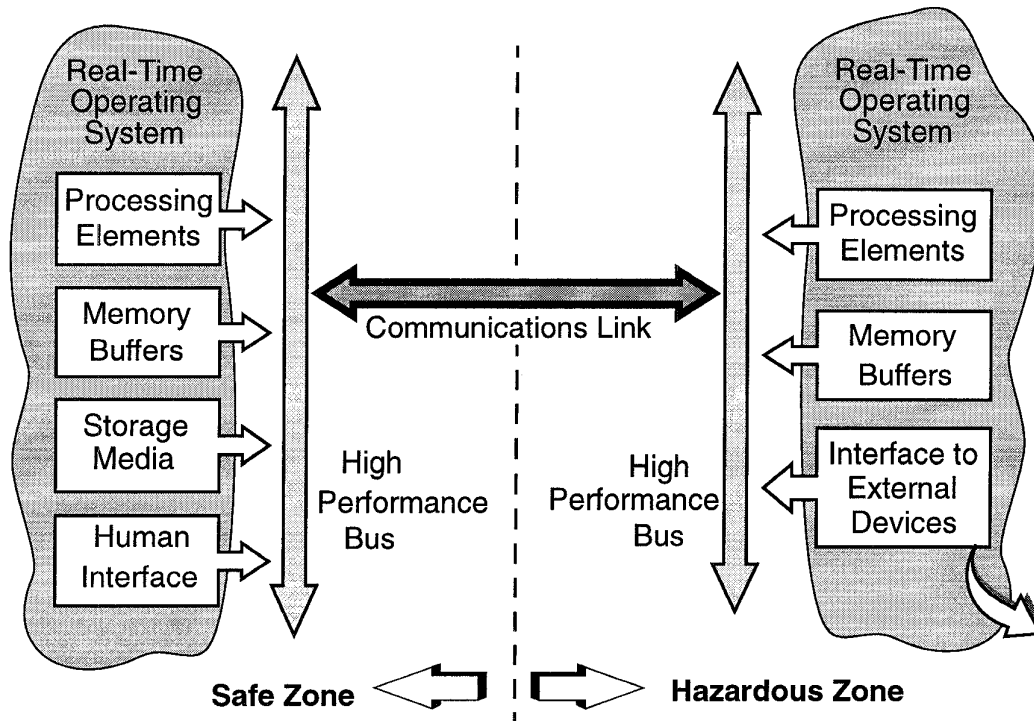


Figure 1. Hardware/software platform for real-time data-intensive applications in hazardous environments

3. store the processed information in high capacity storage media.

The software application developed in each area should be POSIX [4] compliant, being implemented on top of the corresponding RTOS. The software application will be composed of several threads (data-intensive applications are very appropriate for the use of multithreaded schemes, since the performance concerning time response and I/O throughput are highly improved), which will make use of the real-time resources provided by the RTOS facilities [3]: parallel execution, real-time fixed-priority scheduling, memory sharing, asynchronous event notification, high resolution timers, synchronization elements and asynchronous I/O. It is essential to use asynchronous I/O in order to guarantee high global I/O throughput and deterministic temporal behaviour simultaneously.

- **Processing elements.** The workload in every area is shared among its different processors, which make use of multiprocessing. In this way, the global

performance of the platform is enhanced through the use of parallelism. Interprocessor communication and synchronization are two essential issues. The RTOS should provide the adequate facilities to benefit from the use of several processors, transparently for the application.

- **High performance buses.** Provide a fast, reliable interconnection between processors and peripheral devices in each area. The bus architecture should be modular, in order to supply high flexibility and scalability. The combined use of high performance buses with RTOS makes feasible the processing of large flows of information under time-critical conditions. The integration of different buses in the same node of a platform generally implies a global throughput reduction and should be avoided if possible.
- **Memory buffers.** Used as a temporary storage for data. They must be correctly dimensioned to avoid data losing due to the filling of buffers. The size of the buffers is a parameter of utmost importance,

considering that the global bottleneck of the platform can be caused by an excessive number of data transfer operations. Memory buffers should be large enough to keep the number of data transfer operations below a limit, in order to reach the required global I/O throughput.

- **Interface to external devices.** Dedicated hardware/software (basically A/D and D/A converters) that connects the platform to external sensors and actuators for data acquisition or external device controlling.
- **Human interface.** It performs the necessary functions to display any useful information demanded by the operator and to send control orders to the other node of the platform.

4. A real world application: inspection of nuclear power plants

A specific industrial field where the requirements listed in section two are to be fully satisfied is the comprehensive inspection of large, metallic elements used in high-risk industrial processes (e.g. production of electric energy through nuclear or thermic processes).

Present and future changes in international regulations concerning periodical inspections of high risk industrial elements in nuclear power plants, demanding higher levels of accuracy and strictness, are leading to the development of a new generation of inspection systems. New regulations will oblige to acquire and store all the digitalized information resulting from the signals used in diverse sensor-based exploration techniques.

This fact implies a significant increase in the performance and storage capacity required from these systems. Inspection systems developed in the past few years were too specific and heterogeneous, not capable to cope with current and future needs. An important research effort is being done in this area, leading to the development of high performance inspection systems based on innovative hardware/software platforms [2].

Specifically, inspection systems are intended to detect, analyze and classify the early defects that might appear occasionally in the welds of large containers working in hazardous environments located in nuclear power plants, by means of advanced ultrasonic exploration techniques. Data must be obtained within precise time intervals in a predictable way, accomplishing in parallel the digitalization, visualization and storage of the whole information acquired.

Inspections of high-risk containers are currently carried out by stopping the industrial processes that take place

inside any container to be inspected, and submerging a big mobile robot in them. The robot arm incorporates a set of ultrasonic pulsers, that emit ultrasonic waves, and receivers, which detect the corresponding echoes. This arm moves along the welds' surface inside the container inspected. The inspection system evaluates the state of the welds through the digitalized information issued from the echo signals.

The need to inspect the welds by means of different ultrasonic frequencies determines an extremely large volume of information to be computed. In addition, the requirement of keeping the inspection time below a limit for economical reasons -industrial processes inside the container have to be stopped during the inspection- determines the hard timing constraints imposed by the application.

The hazardous nature of the environment is derived from the radioactive processes carried out inside the containers and the electrical noise generated by the robot arms, that incorporate high powered electrical components.

The implementation of a high performance inspection system has been addressed in our laboratory. There are two different nodes in the system, for the reasons detailed in the previous sections. Figure 2 shows the overall picture of the system. The hazardous zone node is in charge of acquiring and sending the digitalized information organized in frames to the safe side 2 Km. away from the inspected container, which for security reasons is placed in a specific area with restricted access. As the robot arm moves along, this node digitalizes and sends the information through a fiber optic link to the safe zone node, where it is processed, visualized and stored.

The reflective memory technology allows different processors to transparently share a region of memory [5]. In this application, a high speed point-to-point link is implemented with two reflective memory boards and a fiber optic cable. A simple interrupt-based protocol has been developed for this concrete application in order to achieve high efficiency.

A detailed description of the specific functions of the two nodes can be found below:

Hazardous Zone Node. The Hazardous Zone Node governs the set of ultrasonic sensors that acquire the data. This node must be physically situated very near the container that is being inspected, to prevent the distortion -caused by noise- of the information acquired by the ultrasonic sensors. It is composed of a VMEbus with several boards plugged in, such as multiplexers, digitizers and special-purpose hardware. A CPU target running the OS-9 real-time operating system is in charge of controlling the rest of boards.

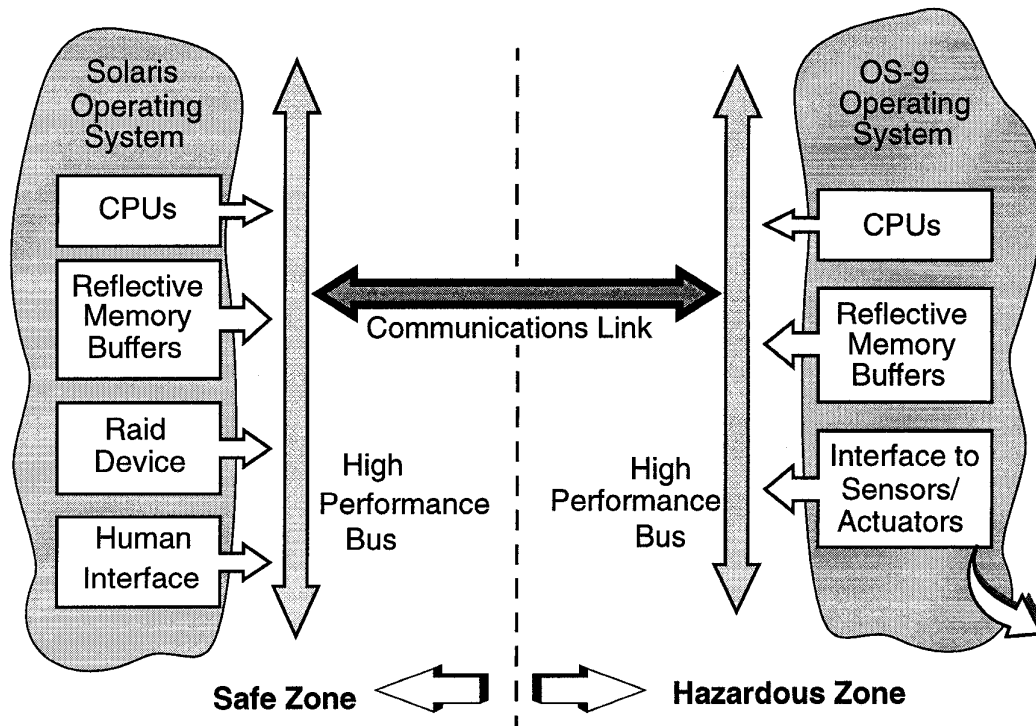


Figure 2. Platform for inspection of nuclear power plants

Safe Zone Node. The movement of the robot submerged in the container is controlled by a specific software controller task in this node. This task traces the physical coordinates in which the sensors are located at each point of the scanning. The controller task is in charge of sending the corresponding acquisition orders to the Hazardous Zone Node whenever the robot arm reaches an acquisition point.

This node receives the information corresponding to each frequency in the reflective memory board plugged in its VMEbus, being temporarily constricted in the sense that the frames have to be read from the buffers located in its reflective memory board before they fill up completely.

The information buffered in the Safe Zone Node memory buffers is processed, graphically displayed, and stored in a high capacity RAID device.

The software structure of this real-time data-intensive application can be modelled with a producer-consumer tasks scheme. Specifically, in order to perform the tasks indicated above, the application software is running on top of the Safe Zone RTOS. The RTOS supporting the application is UNIX System V (including the necessary real-time extensions), under the Solaris development environment [7]. The application threads execute concurrently in the workstation processors. The application

software is in charge of the following actions:

- **control** the physical movement of the robot arm
- **send** the acquisition orders in a deterministic way to the Hazardous Node
- **read** the digitalized data from the hazardous side (special dedicated I/O hardware will carry out all the sensor-related issues, A/D conversion, and other additional operations),
- **transfer** the data into the safe side memory buffers,
- **process** the whole information acquired,
- **store** the information acquired in the storage device, and
- **display** the useful information graphically.

The utilization of sophisticated hardware solutions have to be complemented by the use of an adequate software structure. In order to guarantee both required global I/O throughput and determinism, conveniently modified Rate Monotonic algorithms have been applied to the tasks involved [8].

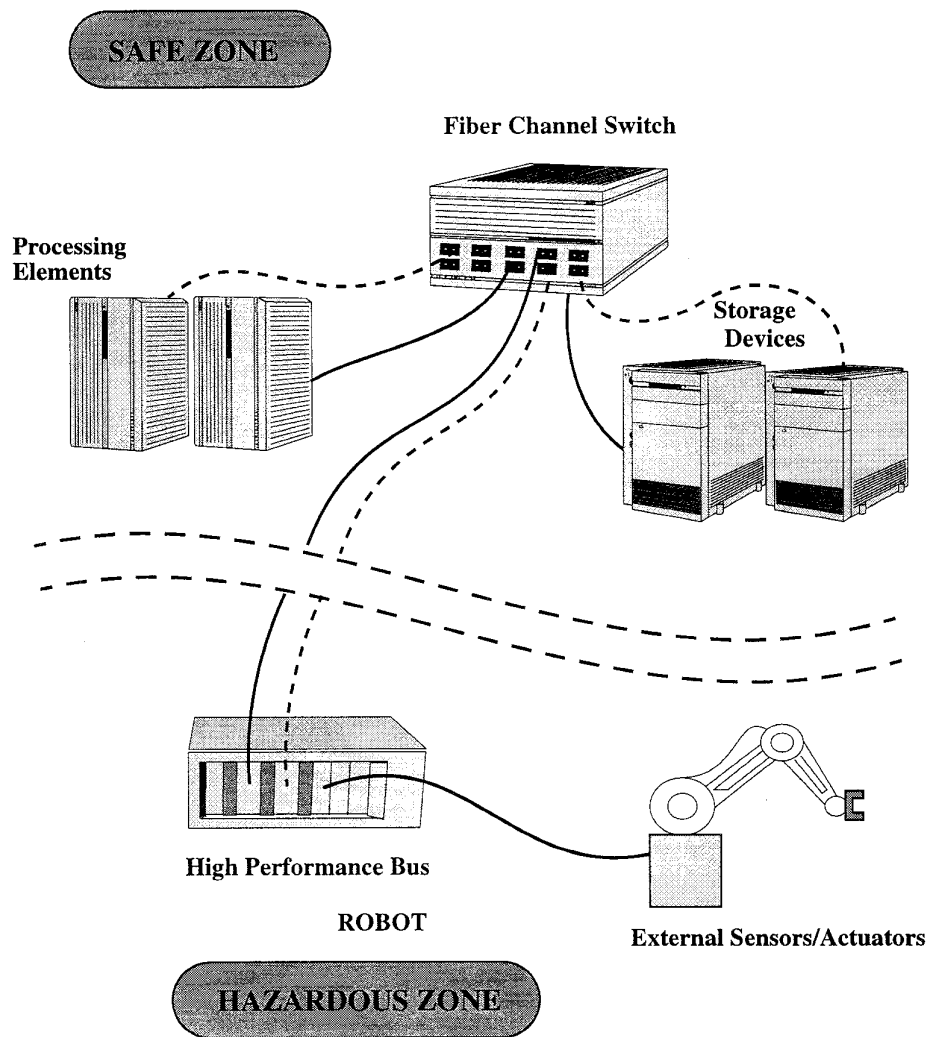


Figure 3. Future architecture of the system

Finally, we report the performance of the system, obtained in a very extensive testing under typical working conditions:

- Precise acquisition intervals, acquisition orders sent to Hazardous Zone Node with deterministic temporal behaviour
- Required global I/O throughput reached (1.9 MBytes/sec.), performing in parallel the acquisition, visualization and storage of the whole information captured (without any data losing)

Extensive laboratory testing of the system has set up the

real I/O throughput limit close to 6 MBytes/sec. of sustained transfer rate. The global bottleneck of the system is found in the communications link. The model of reflective memory used in these inspection system is not capable to achieve data transfer rates over this limit.

5. Conclusions and further work

In this paper we have described an architectural model intended to satisfy stringent real-time performance requirements (response time and I/O throughput) under hostile working conditions. The approach taken suggests that two different locations should be considered in order to

protect human life and the most sensitive elements of the platform. In such real-time data-intensive applications the communication channel plays a key role for achieving the maximum performance.

A real-world application has been presented, illustrating the concepts previously exposed. The problem reported, the inspection of large metallic welds, encompasses several difficulties: real-time response requirements, high I/O throughput, and unsafe area operation. The adopted solutions (RTOS support, RAID storage, reflective memory communications) are technologically advanced, and their integration has outcome as a innovative inspection system.

The development stage is still in progress. Continuous technological advances in external devices (sensors, electromechanical devices), and increasing demands from the users are encouraging the upgrading of the system in order to provide enhanced overall performance and behavior. Required I/O throughput will become higher to keep the system working time at a minimum. Flexibility and scalability, two of the requirements stated in the second section, are of capital importance to cope with current and future demands.

The basic platform architecture for hazardous environments will undoubtedly benefit from emerging technologies. We are currently designing a novel implementation of the system, incorporating Fiber Channel technology [1] to provide scalable and extremely fast communications. Figure 3 shows the estimated final architecture of the system. Fiber Channel is a switched protocol that allows high speed (transmitting rates exceeding 1 gigabit/sec) data communications among processors, storage devices and other peripherals, providing distances up to 10 km. Its intelligent interconnection scheme ensures to provide the needed performance and scalability.

The performance of the hazardous zone cards is being increased, and a platform incorporating duplicated elements is likely to be adopted in the near future. Logical

independent communication channels can be established from the generation of the data to its final processing, surpassing the global bottleneck of the system in a scalable fashion.

This modular approach will enhance a feature of utmost importance for this kind of applications: reliability. While the current implementation addresses this issue with the use of robust components designed to provide high availability (VMEbus, reflective memory error check protocol, RAID storage), it would be desirable to improve this significant virtue by means of redundancy.

6. References

- 1 Fibre Channel Association. *Fibre Channel: Connection to the Future*. HighText Publications, CA, 1994.
- 2 Conde, J. and Viña, A. "A Distributed Real-Time Architecture to Satisfy Hard I/O Throughput Requirements". Proceedings of the 21st IEEE Euromicro Conference on Design of Hardware/Software Systems. Como (Italy), 1995
- 3 Furht, B., Grostick, D., Gluch, D., Rabbat, G., Parker, J. and McRoberts, Mc. *Real-time UNIX systems*. Kluwer Academic Publishers Norwell, MA, 1991.
- 4 Institute of Electrical and Electronics Engineers. *Standard for Information Technology - Portable Operating System Interface (POSIX) - Amendment 1: Real-time extension, P1003.1b*. IEEE Computer Society Press, 1993.
- 5 López, J., del Val, D., Conde, J. and Viña, A. "An Heterogeneous Distributed Architecture based on a Reflective Memory Interconnection Network". Proceedings of the 12th IFAC Workshop on Distributed Computer Control Systems. Toledo (Spain), 1994.
- 6 Ousterhout, J. *Beating the I/O Bottleneck. A Case for Log-Structured File Systems*. UCB/CSD-88-467, October 30, 1988.
- 7 SunSoft. *Guide to Solaris*. Sun Microsystems Inc., 1993.
- 8 del Val, D. and Viña, A. "Applying RMA to Improve a High Speed, Real-Time Data Acquisition System". Proceedings of the 14th IEEE Symposium on Real-Time Systems. Puerto Rico, 1994.