Use of Modern Software Technology in the Redesign of the OPAL Trigger Software

K. Ackerstaff[1], A. Khodabandeh[1], T. Shears[2]

[1]CERN, CH1211 Geneva 23, Switzerland

[2]University of Manchester, Manchester, UK

Abstract

The OPAL experiment at LEP has a central trigger system which is modular and VME based. An upgrade of the control software was necessary to increase the flexibility and maintainability of the system. The software has been completely re-engineered in a project which lasted from January until August 1997. The use and benefits of a customised software development environment as an instrumentation of the design process, and of the project infrastructure, are described. The importance of this approach for the success of the project and its suitability for use in real-time systems is demonstrated.

I. INTRODUCTION

The LEP particle accelerator at CERN has been in operation since 1989. With the advent of the LEP2 era in 1996, the accelerator is continuously being upgraded to gradually reach a maximum centre-of-mass energy of approximately 200 GeV.

The four experiments, which detect particle collisions at LEP, must adapt to the changing accelerator conditions. This has led to substantial increases in background in the OPAL experiment. The planned increase in both the beam energy and intensity is expected to degrade background conditions further.

The OPAL trigger system [1] must be able to reject this background without losing physics events. The original system contains considerable flexibility in programming the fast logic used in trigger decisions. However, during the (long) lifetime of the experiment the trigger control software has had to adapt to changes in the detector, the trigger hardware and the accelerator operating modes, which were not foreseen in the original design. This resulted in a control code which was excessively complex and difficult to maintain.

The difficulties experienced in adapting the trigger control software to these unforeseen changes motivated the redesign of the software. The redesign and implementation was carried out by a small team from the OPAL experiment in close collaboration with one member of the CERN/ECP Information and Programming Technology (IPT) group. The project started in January 1997 and finished successfully in August 1997 in time for OPAL data taking.

The use of modern software methodologies in the high energy physics environment, especially at CERN, is important in view of the software demand of the Large Hadron Collider experimental program. The OPAL trigger redesign project shows that this is a key factor in the software development process. The adoption of suitable methodologies ensured the successful outcome of the project, despite the heterogeneous and real-time environment and the strict time constraints of a running experiment.

II. THE OPAL TRIGGER

The OPAL trigger system selects events according to logic implemented in several custom made VME compatible modules. Programmable combinations of 120 trigger signals from 32 subdetectors are used to form the trigger decision, made every $22\mu s$. A positive trigger decision is sent on a dedicated bus to the subdetector data acquisition systems. After complete readout the subdetector and trigger data are collected by the event builder. The complete events are sent through the filter [2] to data storage. The entire data acquisition system[1] is controlled by RunControl [3] using finite state machines.

The original software ran on two VME CPUs and controlled and monitored the trigger hardware and signals. It was written in real-time Fortran [4], consisted of approximately 60000 lines (11000 statements), and was implemented under the OS9 operating system. The main tasks of the trigger software are :

- to load the user defined trigger logic into the hardware;

- to read out the trigger hardware on each event upon a positive trigger decision;

- to trigger and synchronise subdetector readout;

- to read out monitoring information asynchronously to the event loop and send the data across the network;

- to handle exceptions generated by hardware and software.

The user interface to the trigger monitoring was a commercial histogram presenter which ran on a Macintosh II [5] connected via a VME interface card to the trigger crate.

III. THE REDESIGN

A team of six people worked full- or part-time on the redesign of the trigger software, in a project which lasted from January until August 1997. The project was led by a physicist and supervised by the OPAL online coordinators. The redesign was split into three sub-projects: the core trigger code; the histogramming presenter; the trigger monitoring software. The redesign of the core trigger code is used here as an example to demonstrate the importance of the approach to software engineering.

---

[1]Further details of the OPAL data acquisition system are given in [3]

A. Motivation

The experience of running the trigger system in 1996 when LEP2 started motivated the decision to redesign the software and streamline the hardware. The main motivations were maintainability and flexibility.

Several hardware upgrades carried out in the past could not be fitted into the original design of the old code, which resulted in excessive code complexity. The documentation did not reflect many of these code changes. Further hardware upgrades are foreseen during the remaining lifetime of the experiment, and the manpower required for operations must be reduced. Maintainability can be improved by moving the non real-time functionality to a UNIX platform and thus simplifying the VME hardware. Moving the histogram presentation to the common UNIX platform also removes dependence on specific hardware, and reduces the number of operating systems used.

B. Requirements & Constraints

The new system should cover the full functionality of the old trigger system, and satisfy additional requirements in order to implement the improvements mentioned above.

- The core trigger and the monitoring code have to be separated. The remaining core trigger code has to run on a single VME CPU with no loss of performance.

- The monitoring code has to be moved outside the real-time system.

- The histogram presenter has to be replaced and run on a Unix platform.

- All software must be documented, and designed to be maintainable.

- The design has to facilitate future hardware changes.

To achieve this an object oriented approach was chosen, using the C++ programming language.

The redesign of a central software system within a running experiment imposes several major constraints.

- The existing interfaces between the trigger and outside software systems must be retained, for example databases, event structures and RunControl interfaces.

- The new system must run with the old trigger hardware.

- To ensure a smooth transition, the new system must initially be backward compatible to avoid any loss of data during commissioning.

- Finally the new trigger system must be complete and operational for 1997 data taking, using the manpower available.
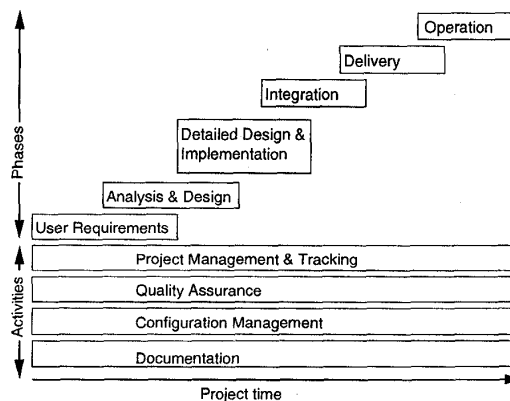

Fig. 1: The software life cycle used in this project. Activities such as documentation and management continue throughout the lifetime of the project, whereas project phases are active only during a specific time.

C. The Software Development Process

The first step in the project is the choice of methodologies in the different activities and phases shown in figure 1. The user requirement document (URD) and the project management follow ESA's Software Engineering Standards (PSS05) [6] insofar as was appropriate to the project. For design and implementation the object modelling technique (OMT) [7] was used. Quality assurance was carried out by peer review and regular analysis of the code quality according to the ISO9126 [8] standards. The project was reviewed on a weekly basis in project meetings. The OPAL online and operations coordinators took part in the review process to ensure that requirements were met.

D. The Software Development Environment

In order to instrument the software development process a software development environment (SDE) was set up at the beginning of the project. Tools for each activity and phase of the process were selected and customised to the specific requirements of the project. Documentation is written using FrameMaker [9] with templates provided by the CERN/IPT group. These documents are hyperised using WebMaker [10] and linked to the OPAL online web. The source code is hyperised using LIGHT [11]. For configuration management the CVS [12] package is used. All relevant documents, utilities and source code are managed under the CVS repository, and changes to the repository are logged automatically on the web.

Code is developed on Unix workstations, whereas compilation, testing and operation take place on OS9 systems. For file and directory transfer across these platforms a tool has been written specifically for the project. The quality of the code has been analysed along ISO9126 quality standards using Logiscope [13], which has been integrated to work with the GNU-make [14] program. Thus GNU-make allows compilation and linking on OS9 systems and automated quality analysis on Unix platforms using the same utility.

As an instrumentation of OMT, Rose [15] is used to produce the C++ class and inheritance as well as message trace

and state diagrams. The ATLAS coding conventions [16] were customised for the project. SNiFF+ [17] is used as the code development environment.

The knowledge and experience of the software engineer from the IPT group has been essential in setting up the SDE and choosing appropriate methodologies and tools. The project team benefited from this in terms of learning and applying the methodologies and tools. The setup of the SDE and the integration of the tools described above took approximately a month at the start of the project. Analysis, design and implementation took altogether three months, and the integration and testing phase two months. This time could only partly be used due to the constraints of the commissioning of the OPAL experiment and the LEP schedule, when the OPAL data acquisition system was often unavailable for testing the trigger.

E. Experience with the Software Development Environment

A detailed documentation of the project infrastructure includes a description of all technical actions in the software development process. This allowed all members of the project team to rapidly learn standard actions and adapt to the SDE. Therefore all tools could be used from the outset of the project.

A substantial amount of information on the details of the trigger system, for example exception handling and interfaces, had to be retrieved from the old trigger code. SNiFF+ has been used to analyse the old software, and FrameMaker used to document the software and hardware, which facilitated the re-engineering process. A detailed quality analysis using Logiscope helped determine the critical parts of the trigger software. SNiFF+ also allows simultaneous navigation through both the old and new software during the analysis phase.

The OMT design enables fast and efficient coding. The models produced with Rose can be tested against the functionality of the old code. Hence the implementation of the models is separated from design and can be carried out by any member of the project team. Furthermore, inheritance allows partial testing. For example, the trigger hardware is reflected in hardware classes inherited from a VME module base class. This base class hides all VME specific memory access and error handling. This class allows either direct access to hardware registers, or else the registers are simulated in local memory if desired. This allows the software to be run on a test setup with only a subset of the hardware modules present. Hardware which cannot be accessed is automatically simulated in software.

The members of the project team were able to work simultaneously on the code using CVS. Following the coding conventions results in a uniform code structure which facilitates peer review and collaborative coding. The majority of version conflicts are resolved by the CVS merging feature. Version tracking and logging are important also during the implementation phase.

F. Product Assessment

Figure 2 shows the overall quality of the new code against time compared to the old code. The percentage of code rated as fair or poor according to the ISO9126 standards is below 2%, compared to above 30% for the old code. The improved code quality is maintained throughout the implementation and integration phases. The new code contains approximately 6000 statements, as opposed to about 11000 statements in the old code. This is due to the increased reusability of the object oriented code and the removal of the monitoring to outside the core code.

The improved quality is reflected in the performance of the new code. It has been operational since the beginning of August
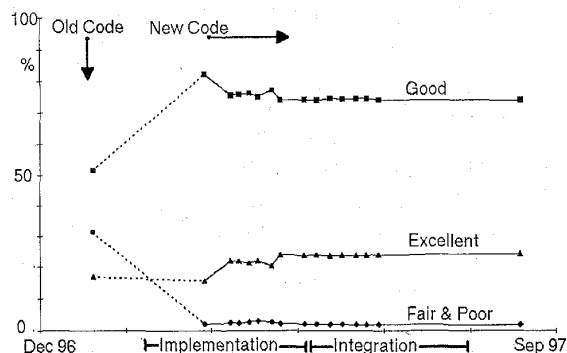

Fig. 2: ISO9126 quality of the old and new code. The percentage of the code in three quality categories is shown. The values for the new code are shown over the duration of the project.

1997, and satisfies all requirements. Unlike the old system no expert intervention has been required during data taking.

The control cycle and exception handler synchronisation is carried out by OS9 signals rather than by introducing artificial delays, as was the case beforehand. As a consequence the software is much more stable. A recent hardware upgrade and associated software changes were completed in one afternoon rather than several weeks needed for the previous upgrade. This is due to the built in flexibility inherent in the design. Moving the monitoring software outside the real-time system protects the core control code from frequent intrusion. Monitoring tasks are automated and data driven from configuration files, thus allowing changes by the user without modifying the code. Due to this and the full documentation maintainability is improved.

IV. CONCLUSIONS

The OPAL trigger software has been redesigned. The project was completed on schedule and met all requirements, despite the constraints imposed by the timelines and existing interfaces of a running experiment. The successful outcome is a direct consequence of using modern software engineering methodologies.

V. ACKNOWLEDGEMENTS

VI. REFERENCES

[1] M.Arignon et al. "The trigger system of the OPAL experiment at LEP," *Nucl. Instr. and Meth. A313 (1992)* pp. 103-125.

[2] D.G.Charlton et al. "The on-line event filter of the OPAL experiment at LEP," *Nucl. Instr. and Meth. A325 (1993)* pp. 129-141.

[3] J.T.M.Baines et al. "The data acquisition system of the OPAL detector at LEP," *Nucl. Instr. and Meth. A325 (1993)* pp. 271-293.

[4] "RTF; Real-Time Fortran-77 for 68K Processors", H. von der Schmitt, available from the OPAL Secretariat, CERN, 1211 Geneva, Switzerland.

[5] Macintosh II and AppleTalk are trademarks of Apple Computer Inc., Cupertino, CA 95014, California, USA.

[6] "Software Engineering Guides", C.Mazza et al. Prentice Hall, London 1996.

[7] "Object Oriented Modelling and Design", J. Rumbaugh et al., Prentice Hall 1991.

[8] ISO/IEC 9126, "Information technology - Software product evaluation - Quality characteristics and guidelines for their use," International Organization for Standardization and International Electrotechnical Commission, December 1991.

[9] FrameMaker, available from Adobe Systems, Adobe House, West One Business Park, 5 Mid New Cultins, Edinburgh EH11 4DU, Scotland, United Kingdom.

[10] WebMaker, available from Harlequin Limited, Barrington Hall, Barrington, Cambridge (GB), CB2 5RG.

[11] LIGHT, available from CERN/ECP/IT, *http://www.cern.ch/LIGHT/.*

[12] "Version management with CVS, CVS 1.9", Per Cederqvist et al., available from Signum Support SA, Box 2044, S-580 02 Linkoping, Sweden.

[13] Logiscope, available from Verilog SA, 150 rue Nicolas Vauquelin, P.O. Box 1310, F-31106 Toulouse Cedex, France.

[14] "GNU make, A program for directing recompilation, edition 0.5 for V. 3.75 Beta" R. M. Stallmann and R. McGrath, Free Software Foundation 1996.

[15] Rose, available from Rational Software Corporation, Immeuble de la gare, 1,place Charles de Gaulle, f-78180 Montigny le Bretonneux, France.

[16] "C++ coding standards for ATLAS", S. M. Fisher and L. Tuura, CERN Atlas Software Note 29 (1996). *http://atlasinfo.cern.ch/Atlas/documentation/notes/SOFTWARE/note29/cxx-rules.html*

[17] SNiFF+, available from TakeFive Software GmbH, Jakob-Haringer-Strasse 8, A-5020 Salzburg, Austria.