The Database Systems for LEP Control

J. Poole,
CERN, CH-1211 Geneva 23, Switzerland

Abstract LEP data are maintained in a central relational database which is accessible from most of the CERN computer systems via the network. The data and the people responsible for it are spread across many disciplines – equipment specialists, machine optics designers and control system engineers for example. Systems have been implemented which interface the database to the control system and to the various users. Data required for controlling LEP are extracted from the central database and loaded into data structures in the control system where they may be accessed by the application programs. Utilities have been built to assemble the data in a convenient form and load it into the control system, to generate alarms in the control room when a change is detected in sensitive data in the database, to archive measurements and settings of LEP and to maintain a relational index to the archived information. Modern methodologies for information analysis and database design have been employed and application software has been analysed and designed using structured techniques. This paper describes the systems and utilities and their evolution from conception to operation.

1 Implantation of the Database

Hardware and Networks for Accelerator Control

The main database for the control system is installed on a VAX cluster which is situated in the computer centre and is connected to the control system via an ethernet link and gateways. The control system networks (shown schematically in Figure 1) use token-passing rings to connect Apollos (consoles), IBM PC_RT's (file and alarm servers), and IBM PC_AT's (process controllers (PCA's) and local consoles). The equipment micro-processors are connected to the PCA's using mil 1553 buses. The networks[1][2] are TCP/IP based and there is connectivity from the level of the process controllers up to the mainframes. There are many database installations on some 20 platforms which include: mainframe IBM, Macintosh, IBM PC_RT and VAX workstations.

Oracle [3] is used as the Database Management System for most of CERN's engineering data including the control system data.

Software

The programs which make up the database systems for LEP control have been implemented across all of the platforms excluding the IBM PC_AT's. Database administration and data maintenance are executed on the central database machines (VAX). In general, users are discouraged from using SQL and have been provided with SQL*Forms to maintain the data. The facilities were designed before SQL*Menu was usable and consequently are based on an in-house menu facility. Associated with each of the major components of the system are forms to maintain the data and reporting facilities.

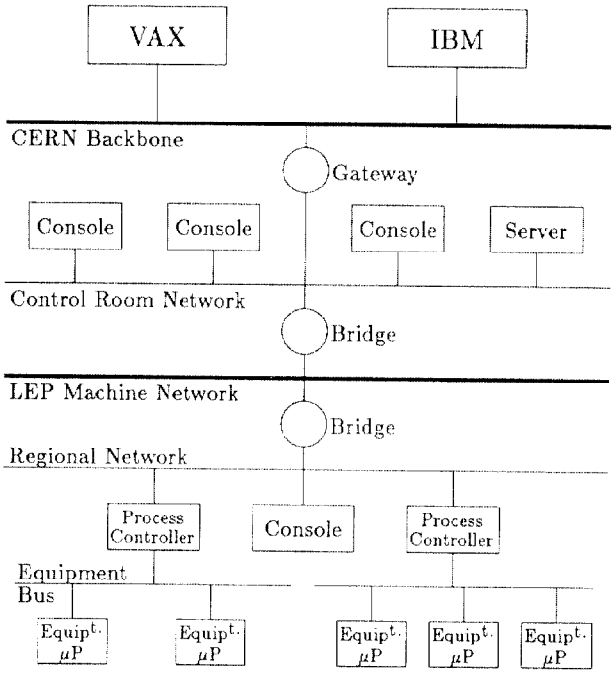

Figure 1: LEP Control System Networks

In the control system itself, are utilities for extracting data from the machine and archiving it and for loading and surveying data.

Typically, control system application programs need to get their data in about 1 second, a timescale which is not compatible with an RDBMS. Access to the data in the control system is therefore via an in-house data management system which makes the distribution of data in a heterogeneous system transparent to the user. This package, Distributed Lep DataSets (DLDS)[4], accesses the data files through re-direction tables on the local hosts. The data are held in self-describing files and are manipulated with a package known as the Table File System (TFS)[5].

Fundamental Constraints and Principles

- Oracle tables are by definition the unique source for LEP data and each group is responsible for the validity of its own data.
- The data in the control system are stored in files which are generically known as the Reference Dataset (RDS)[6], organised in hierarchical directories (UNIX files) structured according to the accelerator system to which they belong.
- A watchdog surveys the Oracle data and warns the operator when data in the control system are no longer valid, in this way the live data are guaranteed.

- Non-relational data from the control system are not stored in Oracle tables.
- Whenever possible, data is thrown away. Data is gathered, condensed, saved and indexed.
- Maintenance of data is done interactively on the database mainframe, not across the network from the control system.
- It was decided to use temporary staff to do much of the coding of the application software and database systems. As a consequence of this, an effective method of analysis and of requirements definition was required.
- The database service is not a round-the-clock service. LEP must be able to continue running without a database service during backup and system maintenance etc.

2 Project Description

Boundaries of the Project

The Database section was responsible for delivery of 'guaranteed' data to the RDS and the access and use of the RDS by applications was taken care of by the Applications section. The task of data distribution and the access methods, implemented by the Applications section, are not described here.

Tools Used

Requirements for control system application software were defined by a working group which used Structured Analysis and Structured Design (SASD) to formulate and record their ideas. The Database section also analysed application software using SASD but the analysis on the database side (Information Analysis) was done using NIAM techniques and then designs were produced from the analyses.

The structured analysis was recorded using the Yourdon Analyst/Designer Toolkit [7] on IBM PC_AT machines. This allowed the construction of Data Flow Diagrams, Yourdon Entity Relationship Diagrams, State Transition Diagrams and an integrated project dictionary.

The information analysis was done using NIAM but at the time, no suitable tool for the platforms in use at CERN was available. The analyses were therefore laboriously made using pencil and paper. The NIAM Information Structure Diagrams (ISD's) were used to produce Entity Relationship Diagrams which were used for the database design. This approach reflects our view that NIAM is essentially an analysis method and Entity Relationship Modelling is a design tool.

Oracle tools were used extensively for the implementation with SQL*Forms for data maintenance, SQL*Net to communicate across the network and SQL*Loader to put data into the VAX from a number of remote hosts. Applications software has been written in the C language and the whole of the control system runs under UNIX System V.

Major Components

Dictionary System The Dictionary System is basically a directory of all objects known to the control system and their aliases. Its primary function is to deliver network addresses to the control system for subsequent loading in the process controllers.

There are several tables and applications for maintaining these data. Family/hierarchical relationships are recorded in these tables and are often used when building reference data files from user tables.

Data Loader The control system Data Loader comprises tables containing SQL statements, used to gather and extract data for the RDS, and software for making the transfer.

There is a one-to-one mapping between the hierarchical directories/files of the reference dataset and the systems/sub-systems in the loader database.

The transfer software uses the system/sub-system name to find the SQL, which it retrieves and then executes and finally the data is written to the reference dataset.

LEP Machine Description The machine description is based on many tables which describe the individual components of LEP, their layout, their physical properties, the optical configurations and run specifications.

Different versions of LEP and even of other accelerators are recorded in these tables and the data are used for installation, maintenance and design as well as control.

Watchdog Rapid access by application programs to their data has been achieved through the RDS, TFS and DLDS. The RDS contains only reference data (read only) and therefore the main problem is that of data consistency. Several instances of the watchdog are scheduled at different frequencies to check different systems, the sensitive data being checked more frequently. The operator is warned via the alarm system when a change is detected.

Archiver The Archive utility is intended primarily as a tool to aid machine development by taking snapshots of the machine and relating its performance to settings and operator requirements. Data gathering is done on request or automatically during a run.

The non-relational data (like settings) are saved in bulk storage resident in the control system and pointers to the files are put with the relational data in the Archive Index on Oracle. The relational data consists of the context information (run number, date, comments etc.), operator requirements and environment parameters. It is therefore possible to make relational searches in the archive index and find pointers to the relevant settings data which can then be re-loaded.

3 Project Evolution

Analysis and Design

Logically, the databases and data structures in the control system should be determined by the requirements of the software but the groups had to develop their systems in parallel because of manpower availability. As a result, a conceptual schema defining broad outlines for the control software and data structures was established and this was used as a starting point for both groups.

Subsequent software analyses at CERN have been made using a different tool, team*work* [8], but this was only available at a late stage in the analysis of the LEP application software and was therefore not used. The process of transferring the analyses to team*work* has begun and maintenance and future work for the LEP control system will be based on this tool.

The input to the information analysis came partly from the detailed data dictionaries of the structured analysis and partly from interviews with 'customers'. Following the analysis phase came the design phases for software and databases resulting in entity relationship diagrams, structure charts and mini-specifications. Finally the designs were implemented on the various platforms using the appropriate tools.

Implementation

The database project really started late in 1987 following an initial analysis. Early in the project it was planned to use an in-house Remote Procedure Call (RPC) mechanism to communicate across the networks but Oracle announced SQL*Net before this was fully implemented. After a period of testing in our environment it was decided to opt for the SQL*Net solution. Apart from this change to the original ideas, little has changed from the initial concepts.

Problems

The multi-platform environment at CERN allows great flexibility but brings an extra overhead in software maintenance. We always have to maintain compatibility across the system between the different versions of Oracle and this can delay upgrades when appropriate versions are not available for all of our platforms.

A severe limitation in the analysis and design phases was the absence of a suitable graphics tool for information analysis and database design. The products which are available in 1990 are an improvement, but there is still no tool capable of supporting rigorous software analysis (e.g. Yourdon Techniques) and information analysis (e.g. NIAM). The tools based on entity relationship modeling are useful design tools, but they are not analysis tools.

Volume of the Data

The volume of data in the database is relatively small: there are some 80 Mbytes currently in tables in the control database. More than half of this is derivable data which will be removed as soon as some new software is available. When the data is loaded into the RDS, ~65 Mbytes are used up in ASCII files. The control database does not contain all of the data relating to LEP, in fact it is only a small part of the total of about 1 Gigabyte. The difference is accounted for by the planning, installation, purchasing, inventory and maintenance data which is stored there.

4 Conclusions

The Database systems for LEP control have been successful: during the running period in 1989 there was only 1 failure which affected operation. The Database section invested approximately 5 man-years (by 7 people) in the construction of these systems.

The stability of the data has been surprising – there has been very little updating although as the system evolves more and more data is required. This stability means that it has been possible to run LEP for significant periods without having to load or re-load data in the live datasets (RDS).

The structured methods have proved particularly effective in view of the high turnover in software technicians. Even without the constraint of high turnover, the benefits from well defined and documented systems are evident.

The concept of feeding tuned data structures from a central database has been proved and is being extended to other accelerator systems at CERN.

References

[1] P. Lienard, *The SPS and LEP Control Network Architecture*, CERN SPS/89-44 (ACC), November 1989

[2] F. Hemmer, *SQL*Net at CERN*, Proceedings of the 7th EOUG, Brussels, 1989

[3] Oracle RDBMS, Oracle Corporation, Belmont, California, USA

[4] Ph. Hofmann, *Distributed LEP Data Sets*, SL-CO Internal Note, CERN, February 1990

[5] Ph. Hofmann, *Table File System*, SL-CO Internal Note, CERN, February 1990

[6] J-P. Koutchouk and J. Poole, *Data Management for LEP Control*, Proceedings of the IEEE Particle Accelerator Conference, Washington, D.C., March 1987

[7] Yourdon Analyst/Designer Toolkit, Yourdon International Limited, London WC1E 7AH, England.

[8] Cadre team*work*, Cadre Technologies Inc., Providence, RI 02903, USA.