WWW FOR INFORMATION ON CERN ENERGY CONSUMPTION

H. Laeger and S. Lechner
CERN, Geneva, Switzerland

Abstract

The electrical energy consumption at CERN and the related costs are of concern for several reasons, not least for their impact on the budget. It is believed that the availability of information on the consumption and the cost of distinct sectors to those who can influence them, i.e. a large part of the CERN population, will stimulate individual initiatives for reductions in consumption and cost. A method has been developed whereby logged data is retrieved from a database and presented in comprehensible histograms. Data can be viewed either as variations over standard periods (day, week, month) or as integration over standard periods. Specific solutions have been devised giving a rapid response time, as well as for generic scaling with optimal resolution for data which may vary over several orders of magnitude. Access to global data is available only to the management and is password-protected, whereas access to all the distinct sectors of consumers is granted to everybody on the CERN site. The implementation is made using the hypertext technique and a WWW server. This paper describes the main graphical tool, the concept of data preparation for instant reply, implemented reliability procedures and some initial experience with its use.

I. INTRODUCTION

By the very nature of its installations CERN is a big consumer of electrical energy, mainly for the accelerators, but also for the basic technical infrastructure. The energy cost is high and represents close to 10% of the CERN material budget.

Since the 1973 oil crisis particular efforts have been made with good results to reduce the energy consumption of someof the larger installations by the introduction of energy-saving operation modes such as pulsing beam transfer lines and tightly adjusting the length of the SPS beam energy flat top to actual beam requirements.

Information on electrical energy consumption, be it for the whole of CERN, for the respective share of its two national suppliers (French EDF and Swiss EOS), or for individual accelerators, technical services or laboratory areas, has in the past been available only to a restricted number of people and only as integrated monthly figures.

A first step towards wider and deeper information for improved energy management, and in particular to ensure the correct response to contractual restrictions on consumption on some critical days in winter time was made a few years ago by putting the actual consumption of larger users at the disposal of all CERN control rooms. This information was then rapidly accessible via the LAN and could thus be consulted by all PC users connected to it.

Following suggestions from the community of physicists, the CERN Energy Management Panel expressed the desirability of extended information such as:

- access to information on energy consumption for the entire CERN population (so far, users of Macintoshes and X-terminals outside the control rooms have been excluded)
- presentation of consumption in the form of histograms and integrated values over time
- presentation of cost, again in the form of histograms and integrated values
- a structured overview of the actual consumption of all larger installations.

It is expected that the availability of useful real-time information on energy consumption and cost, sufficiently detailed that the contribution of individual systems can be spotted, will lead to adaptations and finally to reductions in consumption and in cost.

II. REQUIREMENTS AND CHOICES

A. Software Development Strategy

The procedure adopted for the creation and implementation of this tool was based on interactive progressive prototyping. The reason for this is that the subject is of high sensitivity and rapid reactions from a test group of people were thought vital to verify acceptance of procedure, form and content.

B. Use of World-Wide Web

The need for access on any platform pointed right away to the use of a WWW server since WWW browsers exist for any platform. However, as this information is only intended for the CERN population, the server has to be invisible from outside CERN.

Moreover, as it was requested that some global data concerning contracts with the two national suppliers of electrical energy to CERN should be made available only to a limited number of management staff, a mechanism of password control had to be implemented.

C. Data Retrieval

The required data for these displays can readily be retrieved from a database that is filled by an existing data-logging system [1]. For reasons of response time and reliability, the real-time data is also retrieved from the database and not directly accessed from the equipment. In the present context it is considered acceptable that data which may be up to a few minutes old should be declared 'real-time'.

D. Response Time

A fundamental requirement for any human-computer interface for rapid response time led to the adoption of pre-processed and stored graphs, ready for instant display upon a user call. The price to be paid for this essential advantage is disk space, nowadays not expensive.

III. DESCRIPTION AND ORGANIZATION

A. Reference Data for Configuration

For correct configuration of the system, reference information is needed on electrical energy consumers, such as technical identification and addresses which uniquely identify the measurements in the data-logging system, is needed. Furthermore, the energy supplier (EDF or EOS) to which each consumer is connected as well as consumer tariffs are relevant for cost plotting.

The password for full information display is treated as part of the configuration.

As all the dynamic data to be retrieved are stored in an ORACLE database, the obvious way of storing static reference information was in ORACLE tables. For this purpose four tables were created: one describing consumer information, one each for EDF and EOS tariffs and a fourth one for valid passwords. In addition, one SQL-forms interface was built for easy updating of the tariff and supplier tables.

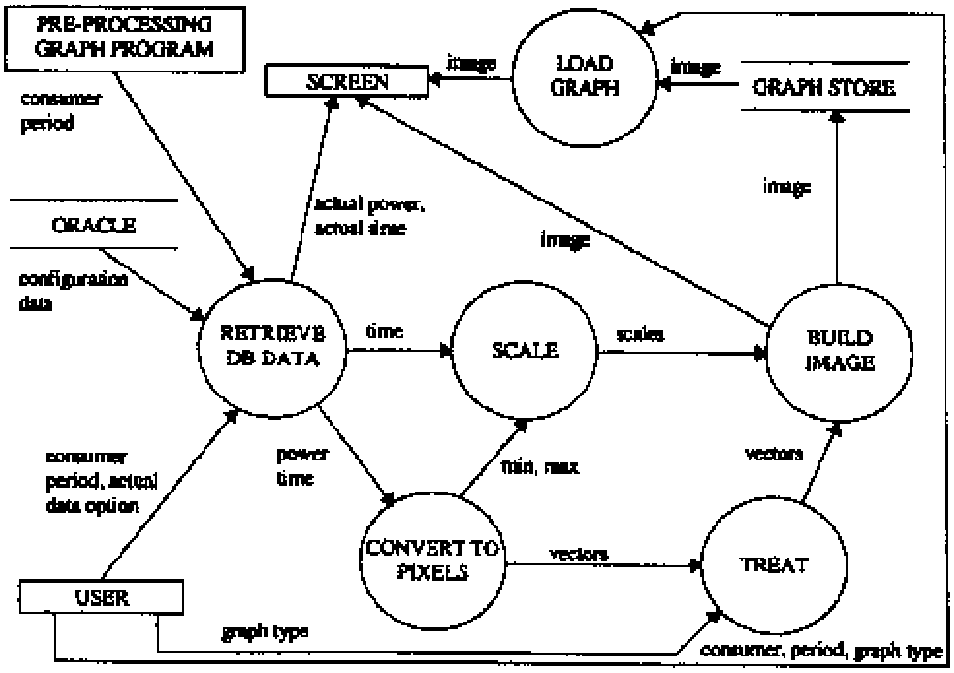

Figure. 1. Data flow diagram for graph selection, production and display

B. Data Management

The graph data (time and value) are retrieved from an ORACLE database (Fig. 1) already filled by an existing data-logging system. A set of routines written in Pro-C (SQL-like syntax) permits the retrieval of data for a particular consumer whose description was read before in the configuration tables.

For cases where data acquisition problems occured at data-logging, with consequent gaps in the database, the data retrieval routines will fill these gaps with the last-available set of data.

Owing to the limited size of the graphic display and the large number of data, some data reduction needs to be performed. The treatment consists in replacing all individual data with the same x-axis values (pixel granularity for time) with their average. The resulting set of points constitutes the final vectors to be plotted.

C. Graphics

The selected graphical library used to plot data is Thomas Boutell's library 'gd' [2]. This library provides low-level functions such as pixel drawing, line drawing, string writing on a graph, as well as graph saving in gif format. It had to be enhanced by a few specific routines (vector plots, legend display, automatic scaling).

D. Scaling

One of the main problems to tackle when displaying data which may vary over several orders of magnitude is that of scaling, all the more since in our case an automatic scaling scheme was desirable.

For horizontal automatic scaling a limited number of natural periods have been selected, i.e. one day, one week, one month. Major subdivisions are then chosen: two-hour periods for a day, six-hour periods for a week, and one-day periods for a month. The C routine built for this scaling task is straightforward.

In the vertical axis, for ease of readability, only a limited number of pre-defined scales are used. These scales are based on multiples of 10^n of the rounded linear $\sqrt{2}$ progression (1, 1.4, 2, 3, 4, 6, 8). The extremities of the scales can only have these values; furthermore if the minimum value is below 25% of the scale maximum, the scale minimum will be set to zero. For scale subdivisions, the same scheme has been adopted with multiples of 10^n of 0.5, 1, 2, and 4 for major units, with a minimum number of 4 and a maximum of 10 major units presented.

The implementation of the above algorithm is simple. Once minimum and maximum values of the vectors to be plotted have been found, the scales are selected from the set of pre-defined authorized values. The same routine is applied to the consumption and to the cost scales.

E. Graph Layout

The general layout chosen for all graphs has been fixed to a 750×400 pixel display, which fits most computer screens and WWW browsers. Pre-defined areas are dedicated to scales and to legends. Graph borders and time-scales are drawn in black, energy consumption scales and plots in red and cost scales and plots in blue. In order to make the plots also readable on black and white screens as well as on printouts, consumption and cost are distinguished by different line styles. The graph title is set with consumer name and graph dates.

The whole graph appears with a transparent background on a WWW browser.

F. Plotting

The resulting vector after data retrieval and treatment is directly plotted as a histogram using a step-function drawing routine. For integrated values, the same drawing function is used but the vector is integrated beforehand by a C routine. Scaling and plotting are treated separately.

G. Pre-Processed Graphics

Initial prototyping showed that the response time was unacceptably long when the user has to wait for the time it takes to access the database, to retrieve the data and to build the graph. Typical response time for a 'day' graph with about 10 000 data couplets was 5 seconds, 15 seconds for a 'week' with about 70 000 couplets and 45 seconds for a 'month' with 250 000 couplets. We therefore opted for a procedure with pre-processed graphs, which reduces the response time to that required for the connection to the server and for the transfer, decoding and display of the graph. Normally, the user gets the requested display within a few seconds. The 'real-time' data, and the graph of the actual day are not prepared in advance and there the response time is higher (an average of 30 seconds and 10 seconds, respectively).

The global strategy is to store graphs of the last five weeks, of the last 31 days excluding the current day and of all past months including the last 31-day period.

The program written to create pre-processed graphs begins by reading configuration tables and tries to process all 60 graphs in sequence (last day, current week and last 31-day period for every defined consumer). However, in case of problems (ORACLE unavailability essentially) no graph is stored at all. Graphs are processed at a given time at night; afterwards verifications are done periodically to look for missing graphs in which case processing is reinitiated. Assuming ORACLE data retrieval problems remain for a whole day, only the graphs for the last day will be missing and current week and month graphs will show incomplete data.

Since the defined policy is to keep the last 31 days and the last 5 weeks, older 'day' and 'week' graphs are deleted in order to minimize the disk space used.

H. User Interface Entry Page

The entry page (Fig. 2) into the system called ‘Energy Consumption and Cost Estimation for CERN Accelerators and Experiments’ displays all selection points as hypertext links that lead to the corresponding graph page for the display of histograms.

It was initially intended to display the real-time data as well as the hypertext selections for the individual users. For reasons of response time, we opted for a more rapid display of the selection points which now contain an additional button for access to the actual data.

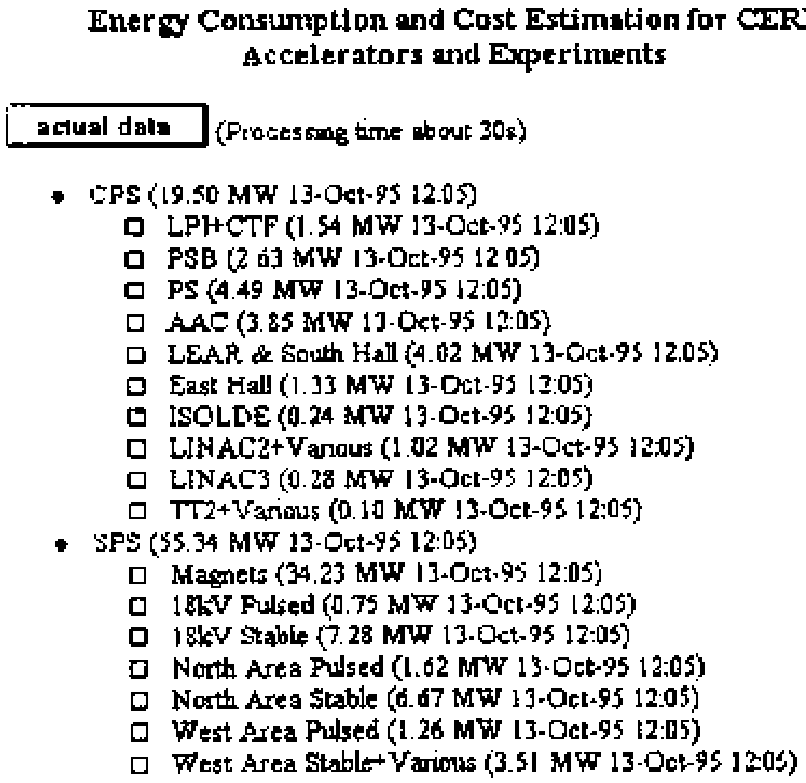

Figure. 2. Part of the entry page with real-time data

Moreover, as access to information on the global energy consumption of CERN was not to be displayed to everybody, a mechanism of password access for a limited number of management staff had to be implemented in the form of an additional button on the entry page.

Finally, a hypertext link called ‘Comments’ at the bottom of the page allows users to address mail easily and directly to the project administrators.

I. Graph Page

The graph page (Fig. 3) used for different selection points and graph types (histogram or integration) is the same except for the page title.

Users can choose between one of the last 31 days, one of the last 5 weeks, or any month including the last 31-day period. Once an option is selected, the requested graph is displayed on the same page just below the menu options. Every graph page presents a direct link to the entry page and to the other graph type (histogram or integration) for the same point. If another option for the same point is requested, the resulting graph replaces the previous one.

Whatever options are chosen it is always possible to get information on energy cost and cost estimation since the graph page title contains an hypertext link called ‘Cost Estimation’.

From any position, a direct link to the CERN home page is present for new navigations starting at this point.

J. Implementation of the User Interface

Entry and graph pages treat data input from the user (button pushed, value selected) and have therefore been implemented as Common Gateway Interface (CGI) [3] C scripts. Much effort was put into obtaining an application compatible with most browsers especially NSCA Mosaic and Netscape and their different versions.

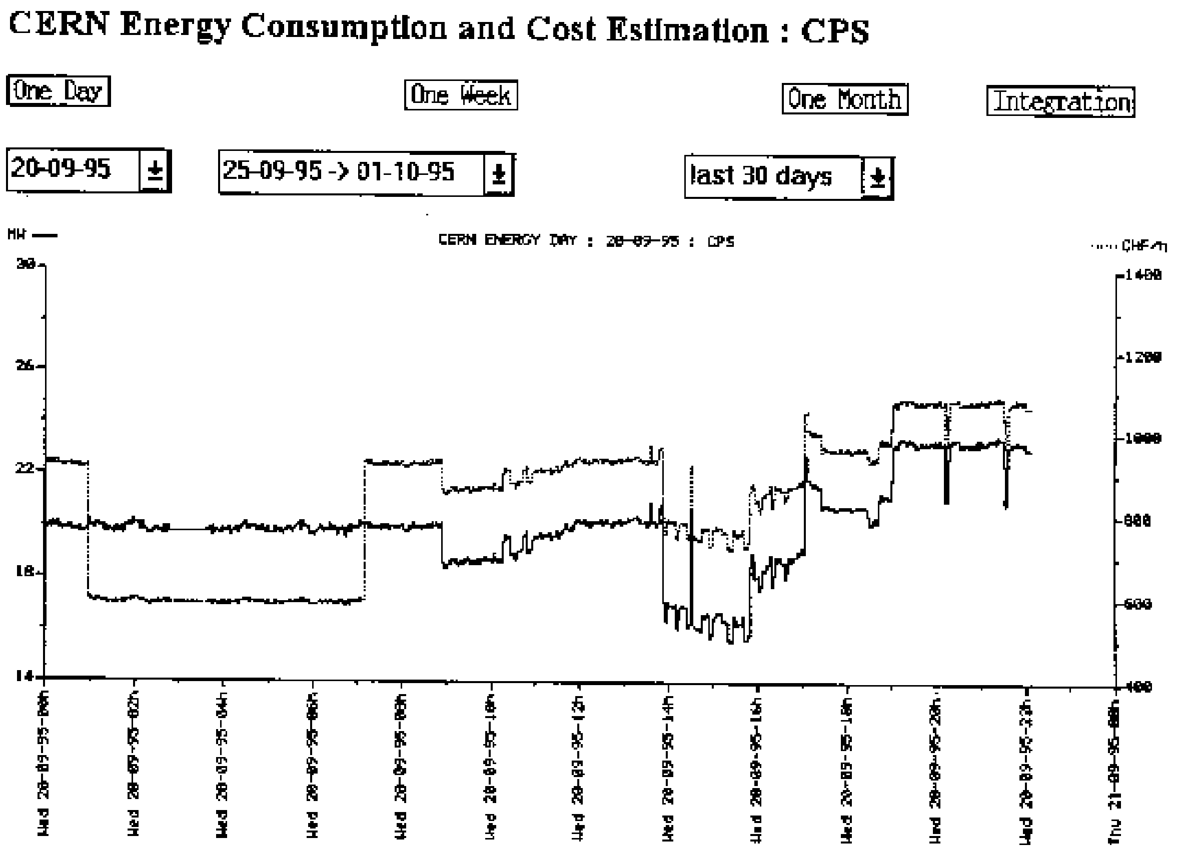

Figure. 3. Example of graph

K. Entry Page

The default call to this page displays all information mentioned above using the C basic print function. When real-time data are requested a call to ORACLE is made for the last logged values for the selection points and the result is displayed on the screen. For access restriction, the entered password is compared with the valid one and, if correct, full information is displayed (Fig. 4).

L. Graph Page

The graph page is called by the entry page which sends information such as the requested point and the graph option to the WWW server. Once the user requests a particular graph, additional information such as day, week, or month and their corresponding dates is available and the graph is plotted. If the requested graph exists as a stored gif file it is just redisplayed, if not, it is built in real time. When several graphs are requested one after the other, the selection point and the type of graph remain as hidden parameters and the other options are changed by the user.

To avoid use of multiple submit buttons on the same page which may not be recognized by all browsers, each button has been replaced by a small gif image representing the button. These images, linked to the HTML command 'input type=image src=...', act as buttons.

IV. EXPERIENCE AND OUTLOOK

The tool is currently in its final testing phase. Availability is very high, limited essentially by the fact that the WWW server is installed on a HP-UX machine which runs a multitude of development programs. This can be improved by installing the application on a more appropriate machine.

Response time is now close to optimum; it remains to be seen if the delay of 'real-time' and 'actual-day' data is accepted by the users.

User comments are so far only available from the limited number of test persons; they are essentially positive and welcome the availability of the information in the proposed form. The application has been built for easy adaptation to new user requirements and it is expected that once it is used by the entire CERN population we will need this facility.

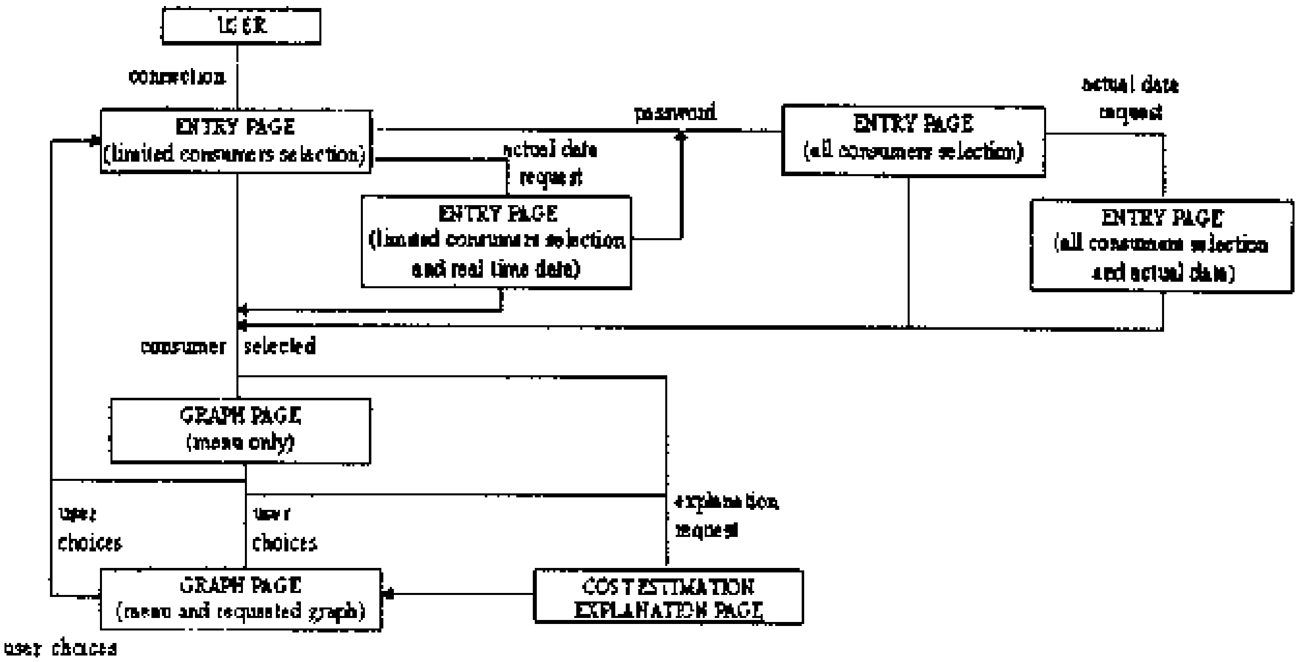

Figure. 4. Navigation overview

The data-logging system which has been in operation for a two years sets different intervals for data-taking according to individual preferences. For energy measurements some are set to one minute, others to five minutes, most to ten minutes. Because of the 'real-time' use of this data by the present application, all intervals for energy measurements have now been set to two minutes. The subsequent increase in the number of equipment calls and in storage capacity required for the database is well within system capacities and so far has not presented particular problems. The diskspace required for the pre-processed graphs for all 'day', 'week', and 'month' graphs for a year amounts to about 20 Mbytes.

The intensive testing during the prototyping phase of the energy displays revealed a relatively high number of gaps of different duration where data was not logged by the data-logging system. This stimulated efforts to increase the reliability of the data-logging system. A Technical Data Server [4] to be used for the supervision and control of the technical infrastructure of CERN has been defined to also serve this purpose.

V. CONCLUSION

The tool described in this paper provides an easy and reliable way to distribute information on CERN energy consumption and cost all over the site, whatever the computing platform used. The final testing stage has shown that it is reliable and that changes can be made easily. All of its constituent modules proved to be well interfaced and, to a large extent, they can be changed independently.

Clearly, the mechanism for the display of information on CERN energy consumption and cost is only a tool. However, it may stimulate individual actions and together with other management incentives may achieve the basic aim of the reduction in energy consumption and cost.

References

[1] R. Martini, H. Laeger, P. Ninin, E. Pfirsch and P. Sollander, "Data Logging for Technical Services at CERN", these Proceedings

[2] gd is copyright 1994, 1995, Quest Protein Database Center, Cold Spring Harbor Labs.

[3] The Common Gateway Interface, https://hoohoo.ncsa.uiuc.edu

[4] P. Ninin, H. Laeger, S. Lechner, R. Martini, D. Sarjantson, P. Sollander and A. Swift, "The Technical Data Server for the Control of 100 000 Points of the Technical Infrastructure at CERN", these Proceedings