EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

CERN - PS DIVISION

EE

CERN LIBRARIES, GENEVA

CERN/PS 93-49 (CO)

8x9348

9

CM-P00059543

THE NEW CERN PS TIMING SYSTEM

J. Lewis, V. Sikolenko

**International Conference on Accelerator and Large Experimental Physics Control
Systems
Berlin, Germany, October 18-22, 1993**

**Geneva, Switzerland
15/11/93**

THE NEW CERN PS TIMING SYSTEM

Julian Lewis, Vitali Sikolenko
CERN, Geneva, Switzerland

Abstract

The PS complex consists of nine interacting accelerators working together, which, from cycle to cycle, produce beams varying in end user, particle type, energy, time structure and beam-geometry. Since the introduction of the new timing system, the sequencing of the PS accelerators now depends dynamically on their status, so that sequence changes in real time are now provoked automatically. This greatly improves the complex's response time to changing end user requests, and simplifies the task of the machine operators who no longer need to program it manually. Coordinating this intricate time sharing particle factory is the MTG (Master Timing Generator), which broadcasts messages around the complex containing summary information on what each part must do next, and the timing needed to carry it out. These messages are received by Tg8 VME timing modules which then provide nearby equipment with timing pulses, and the VME host processors with task synchronization events and summary information.

1. Sequencing concepts used in the CERN PS complex

The CERN PS complex is a factory which produces particle beams. Providing the complex's coarse timing ¹⁾ is achieved initially through specifying its operational requirements. The operational requirements describe what beams are to be manufactured; from which the coarse timing, can be algorithmically derived. In order to achieve this, we make three definitions which characterize the workings of the complex, namely Cycles, Beams, and, for want of a better term, Beam Coordination Diagram (BCD). Any factory can be described as a directed graph, the nodes representing the various manufacturing processes, the edges carrying the intermediate products from node to node until finally an end product exits on some edge. In our case the nodes of the graph represent the various particle accelerators making up the PS complex, and the transfer lines, the graph edges, which transport the intermediate products, particle beams, from accelerator to accelerator until they are finally ejected from the network. A Cycle is the activity carried out by a node in the process of manufacturing a particular beam, and a Beam can be considered as a path ²⁾ through the graph with a corresponding

Cycle to be sequentially executed on each of the path's nodes. The set of all Beams to be manufactured for the end users is the same as the operational requirements for the complex, and a BCD is a strategy for carrying out these requirements in real time. As a Beam is a linear list of nodes with a corresponding Cycle to be executed on each of them, and each of these cycles occupies a small integer number of basic periods ³⁾ of 1.2 seconds, we can now represent the flight path ⁴⁾ of the beam through the network in time. As each beam in the operational requirements set can be represented in this way, our problem is simply one in which we arrange their departure times in such a way that we do not exceed the node capacity of any node. The node capacity is always one, as at PS, we never have more than one beam in a machine at the same time. Imagine a deck of beam cards, each card representing a beams' flight path through the network in time as follows:

combine beams together in the process of manufacturing a beam, so we can simplify by using a linear tree (PATH), i.e. a list of nodes.

-
- 1) Coarse timing. Interesting instants in time within machine cycles, with a 1 ms resolution and 500 ns precision.
 - 2) Path. In the general case, we need a tree to represent and end product, at PS we do not

- 3) Basic period: This is the LINAC repetition rate, i.e. the rate at which raw materials enter the factory.
- 4) Flight path. We can represent this in time as a beam card: AA beam = { Node, Period }, { ...

Figure 1 : Beam Card graph

A deck of such cards represents the operational requirements of the complex, and starting at say the top of the deck, we can slide lower cards to the right, one period at a time, until no node occurs more than once in any vertical column, I.E. time period. In this way we are able to generate a set of departure times, if we shuffle the deck we can produce another set of departure times. The initial order of the cards in the deck thus determines the outcome of the algorithm, and hence any behavioral mechanisms we may wish to incorporate must order the cards accordingly before presenting the deck to the departure time algorithm. Notice that adding or deleting or replacing one beam by another is very simple in this domain.

2. Dynamic changes in the accelerator super-cycles.

The PS complex operators work with individual accelerator super-cycles⁵⁾ which are determined by the BCD, and they expect an intuitive behavior from the user interface with respect to inserting, deleting and replacing Beams. It is this behavior which, in

our case, gives the criteria for arranging the order of the Beams in the card deck, other criteria could have been based on total beam throughput, or machine physics constraints.

We can transpose the Beam Card graph in Figure 1 into a Super-cycle graph similar to that of Figure 2. The Normal and Spare super-cycles of each machine are laid out on the work station screen as shown in the small extract above. The PSB Normally executes an AA Cycle after which the beam is transferred to the CPS, where a two basic period AA Cycle is then executed; this constitutes the AA Beam. In the event that the LEAR requests beam, the Spare Beam sequence is used to replace the Normal sequence in real time. Notice in Spare that the PSB sends its Beam to the ISOLD machine instead of to the CPS which is now executing a LEAR Cycle, later on the PSB will send a TeST beam to the CPS. This kind of manipulation which is driven by accelerator hardware status in real time, would be impossible to manage without the Beam concept.

⁵⁾ Supercycle. A quasi-stable sequence of cycle repeatedly carried out by an accelerator which is determined from the BCD and the hardware status conditions of the machines in the complex.

Figure 2 : Super-cycle graph

3. The Master Timing Generators (MTG)

Once the BCD has been constructed with the Graphical User Interface, it may be sent to an archive containing the current hardware settings which is also seen by the MTG. The primary function of the MTG is to coordinate and sequence the activities of the complex in real time based on the BCD and the hardware status of each accelerator.

The MTG reads the hardware status of the accelerators under its control, and tests each Cycle of the next Normal Beam against the corresponding accelerators' decision procedure ⁶⁾. If the result of any test fails for any Cycle in the Normal Beam, then that Beam can not be executed, so the Cycles of the Spare Beams are next tested. Any failure at this stage causes the Dump decision to be taken, in which case the Beams' injector-destination is modified to a safe beam dump.

Figure 3 : MTG Decisions

⁶⁾ Procedure. Interpreted texts, easy to modify at run-time, as demanded by changing operational circumstances.

Having decided on a beam production strategy, the MTG builds the Telegram⁷⁾ messages describing the next Cycle to be executed in each machine, and determines which coarse timing events must be produced accordingly. From the resulting schedule, the MTG builds an event list describing each millisecond in the next basic period for each accelerator and broadcasts it over the timing drop net. These events are then received locally by Tg8's [1], which produce local timing pulses for nearby equipment, and provide local application programs with timing and sequencing information.

4. Timing reception concepts

All timings in the PS control system are assigned a unique member number in a timing super class. This class is composed of two sub sets, namely the MTG events which are transmitted over the timing drop net, and locally produced timings which may be related to the MTG events by means of, an optional clock, a delay specified in this clock, and a coincident telegram condition.

These local timings exit from the Tg8s via cables which connect them to the nearby equipment. Effectively these cables and the MTG drop net carry instances of timings to their destination, and the set of all instances of a timing constitutes a timing class member. A task running in a host processor controlling a number of Tg8s can be connected to any instance of any MTG event or to any instance of any local timing, whereas the local equipment can only be connected to a local timing member.

The telegrams which are also transmitted over the timing drop net, are used by the Tg8s to control the generation of local timing instances, the so called Pulse to Pulse (IE Cycle to Cycle) Modulation (PPM) of the timing delay values. The telegrams are also available to the host processor for controlling Pulse to Pulse Modulation of other parameters such as power supply settings etc.

5. The Tg8

Each Tg8 system consists of a host processor device driver, and the hardware and firmware of each Tg8 module in the hosts VME crate. They receive information transmitted from the MTG over the drop net, and provide the functionality described above.

Each Tg8 module contains 8 programmable counters, the drop net reception logic, and a MC68332 processor. The firmware simulates 256 quasi parallel timing engines which are defined and may be redefined at run-time by the host VME. A timing engine which is called an action is activated by the incoming MTG event associated with it, it then checks the specified telegram condition, chooses the specified clock and delay and produces either an output signal to local hardware using one of the 8 Tg8 output channels or a VME interrupt which is translated then by the associated host software to a signal submitted to connected tasks. The translation of the description of local timing to the table of actions is also performed by the host VME software. The Tg8 also provides a lot of status and diagnostic information related to the timing system for the host VME processor.

Acknowledgements

We would like to thank G. Daems and G. Beetham for their help in specifying the basic functions of the Tg8 system: B. Puccio for the design and implementation of the Tg8 hardware; Ph. Nouchi for the MTG PC card hardware and firmware, and the Tg8 XILINK firmware.

Reference

- [1] A new VME timing module: TG8 ICALEPCS 1991, KEKJapan.ver the 88 tiC, G.Beetham, G.Daems, J.Lewis, B.Puccio .

⁷⁾ Telegram: a message sent out to each accelerator at a precisely defined time just before the next cycle is to begin. The message contains a set of parameters which both describe and identify the beam for the accelerator's control system. These messages drive the accelerator's Pulse to Pulse Modulation (PPM) mechanism.