Accelerator Controls Section
Alternating Gradient Synchrotron Department
BROOKHAVEN NATIONAL LABORATORY
Associated Universities, Inc.
Upton, New York 11973

Accelerator Division Technical Note

AGS/AD/Tech. Note No. 328

PULSE-TO-PULSE MODULATION AT THE AGS COMPLEX

Operational Aspects and Controls

by D.S. Barton

September 1, 1989

1. Introduction

This note describes the scope and method chosen for providing controls to support operation of parts of the AGS complex in a multi-user mode, referred to as PPM. The first section introduces concepts and terminology, and then discusses some important operational aspects of this method. The second section describes the new timing system which is being constructed for the Booster and which will, at the same time, manage the overall intermachine scheduling in the AGS complex. Finally, the third section discusses briefly the principal changes which must be made in the various parts of the AGS Distributed Control System (AGSDCS) to accommodate this new requirement. This document does not address any possible future need for the Tandem Van de Graaff to utilize PPM controls. Since this note discusses some details of control system modifications, it unavoidably contains terminology and references to parts of the system architecture with which the reader is assumed to have a passing acquaintance.¹ This subject was addressed previously, in a more preliminary fashion, in an internal Control Section note by Stevens.²

1.1 Definitions

The name Pulse-to-Pulse Modulation (PPM) was first applied at the CERN PS to the operation of that facility for multiple end users, with characteristics of beam intensity, species, main magnet cycle, and certain other machine parameters which are variable on successive pulses. It is planned to introduce PPM into the controls for the AGS complex in a phased manner, to support the increasing demands for flexibility and multiple use. In this context the term Supercycle

refers to the overall repetitive sequence of the facility as a whole, while a Cycle is defined as the single, complete period of operation of a particular accelerator, usually associated with the sequence of a main magnet and/or rf system. For that machine, the cycle is the smallest unit of time which is assigned to a particular User. The expression, user, refers to a specific application (client) of a set of cycles of one or more accelerators of the complex, usually a recipient of beam.

The implementation of PPM described below utilizes a central timing system, which generates and distributes timing signals to the subsystems of a particular accelerator via a Time Line, a cable carrying information in the form of encoded pulse trains concerning both the time (in answer to the question, When?) and the function (in answer to the question, What?) of the signals. A crucial set of signals called Resets inform the various accelerator subsystems of the user assignment for the next cycle or group of cycles.

1.2 Purpose

The objective of the creation of PPM controls for the AGS complex is to provide a flexible, centrally managed facility which supports multiple activities during a super-cycle. These activities might include one or more physics programs and studies or other secondary uses. In each case, the environment in which a particular user operates should provide an exclusive view of just those matters concerning the cycles associated with his task. Control and data acquisition should be independent of other users, except in cases where there is unavoidable contention for some instrumentation. Resolution of these cases would typically be handled between the device controller (front end microprocessor) and console program. The ultimate goal is to provide as close an approximation as possible to a separate "virtual accelerator" environment for each user.

The first instance of PPM at the AGS complex is already in place at the new Linac pre-injector. This implementation has been preliminary and has not incorporated many of the database and control protocol changes discussed in Section 3 of this note. The first requirement for multiple users, beyond the Linac local usage, is for Linac operation in support of Booster commissioning. This start-up is expected to occur during AGS HEP running with the present (HEBT) injection. During the first several months of Booster operation, PPM will not be utilized for Booster equipment. On the other hand, as soon as the AGS begins operating full-time for the HEP program with injection from the Booster, PPM on that machine should greatly enhance access to study time for further development of proton operation and for commissioning of heavy ion operation.

In the longer term, the same considerations apply to the AGS, both to support rapid changes of primary client, as for RHIC, and to give enhanced productivity in the performance of machine studies in a new era of high intensity proton and high-A heavy ion operation. Experience at the CERN PS has shown that more than 70 percent of scheduled machine studies are carried out using PPM cycles as a component of normal schedules. Whether or not a specific accelerator subsystem is built or converted to support PPM is considered to be a separate issue from the general plan to implement PPM-compatible controls. There will, of course, be systems like vacuum equipment which will definitely not have PPM controls.

While implementing this new system, there has been a conscious effort to constrain its complexity and generality. The experience at CERN with a more general system, which allows more complex combinatorics in the modes of operation was extremely painful.

1.3 Operational Aspects

In order to simplify the supercycle scheduling task and limit the generality of the PPM implementation, all accelerators will cycle on a synchronous frequency of 7.5 Hz (or a subsequently specified Booster standard) and its sub-harmonics. If the Booster is not rapid cycling for protons, or is only accumulating polarized protons, it would be possible to build the supercycle with some other Linac periodicity as its basis (e.g., 5 or 10 Hz). The supercycle will be built with an AGS cycle determining its least duration, with more than one AGS cycle as required by program complexity.

Instead of viewing the manifestation of PPM as comprising several main magnet cycles with very different characteristics, one expects a more likely scenario for the AGS complex will embody cycles with differences primarily in other equipment (rf, correctors, etc.). A trivial example, but one which represents the first operations needed for PPM, is that of the Linac with four or more users (AGS(HEBT), Booster, BLIP, REF, ...). Fig. 1 shows a hypothetical sequence of Linac cycles as a function of time, as indicated by a Linac prepulse and an rf turn-on pulse. Other details of the Linac timing sequence are omitted. Some characteristic Linac parameters are listed, which are PPM variables in the pre-injector area. This means that, in the front-end microprocessor device controllers, the setpoint and command variables are stored separately for each user. They are sent to the equipment according to the receipt of a command signal for the next cycle. Analog and digital acquisitions from the equipment are also separately managed for each user. At present, the Linac rf controls do not support PPM, but they might become PPM variables one day.

Figure 1.

Representative Sequence of Linac Events

Table 1.

Sample PPM User Assignments

Accelerator	Max # modes	User Names
Linac	7	AGS Physics AGS Studies Booster Studies [spare] BLIP REF Linac studies
Booster	4	AGS Physics AGS Studies Booster Studies [spare]
AGS	4	AGS Physics AGS Studies [reserved] [spare]

the effects of magnet hysteresis are minimized. This restriction does not mean, however, that acquisitions of equipment readings will not be available with full details about each instance of a cycle; they will be. At the Linac, and for proton injection cycles at the Booster (7.5 Hz), the data acquisitions by the device controllers will be synchronized at that rate. The reporting of data to Stations on the network will be carried out as now done at the Linac at approximately a 1 Hz rate. This will mean, for instance, that data reports from Linac and Booster devices will be buffered at the controller for several cycles.

It will be unusual, if not impossible, for the principal On-Off control and interlock status of a device to be a PPM variable, with separate user access. Devices will be controlled through a set of "logical devices", some of which may be non-PPM and others (e.g., setpoint control) PPM. There will clearly be cases in which a device one would normally think of as "OFF" or "not used" for a particular user process will need to be "ON" at zero or other appropriate value (e.g., to correct for hysteresis).

A potentially painful aspect of the early operation of the Linac and Booster in a PPM mode results from the lack of direct computer control of the Siemens motor generator set.

The current operational procedures for changing the Siemens cycle are exceedingly manual in nature. Moreover, it is customary for the length of the AGS repetition period to be changed routinely in order to stay within limits on average power usage. On a less frequent basis, the length of the flattop is changed to match program requirements. In PPM operation, the length of the AGS period forms the basis of the duration of the supercycle. The lack of computer controls for Siemens means that a completely computer-synchronized change of supercycle and Siemens cycle length is impossible. This problem will have to be addressed initially by procedural means. The shortening of the AGS repetition period (within limits derived from the AGS cycle length and Booster supercycle definition) would be straightforward. In this case, the requisite number of Linac local user cycles would be removed from the period between AGS cycles. If the flattop length were being shortened, this supercycle revision would be scheduled to follow the actual cycle alteration at Siemens. If the repetition period were to be lengthened, the supercycle duration would be increased first by the addition of Linac local user cycles, followed by any Siemens cycle alteration. It may well become standard operating practice to make all routine changes to the supercycle from pre-established files, which could be identified with certain Siemens conditions and Booster program choices. Even newly defined supercycles will have to be verified in a manner which must be managed much like a file restoration.

2. Timing System

Central to the implementation of PPM for the AGS complex is the creation of a new master timing system. The present system of timing control at the AGS comprises a mixture of local manual control and computer controlled delays of varying vintages. Two real time clock trains and the Gauss clock are distributed where needed for use by these local timing units. The new timing system will be first extended to the areas directly related to the Booster. The general method borrows heavily from the Fermilab approach. Figure 2 contains a block diagram of the central portion of the new timing system, as applied to the Booster. The major timing signals related to a particular machine cycle will be generated centrally by programmed selection of timed events on a master real time clock, and, where appropriate, on the machine's Gauss clock. These event times will be user-specific PPM variables. Each programmed event will generate a serial coded, 8-bit pattern transmitted on the time line. Each code will designate a specific function, and equipment will be configured to recognize and act on selected codes. The central distribution system will comprise three distinct time lines. One line will distribute the master schedule events from a supercycle controller, while the other two lines will distribute Booster real time and Gauss clock events.

2.1 Supercycle Control

The timing master for the complex will be the Supercycle Generator (controller). This device will determine the scheduling of users and of certain fundamental events related to cycle initiation. The master clock of the Supercycle generator will be a power line synchronized, 60 Hz clock, or a clock operating on a harmonic of that clock. The supercycle time line (SCTL) will primarily feed the individual cycle generators for the complex with the user reset codes to indicate which cycle is next to be produced. In addition, each real time cycle generator will distribute the user reset codes along its own time line to the equipment associated with that accelerator.

The supercycle generator will be capable of generating up to 15 codes in rapid succession on a single tick of its clock (60 Hz). The sequence of events produced by the supercycle generator will also include basic prepulse events for each accelerator, and where possible, the actual start-of cycle event (T_0). Exceptions would have to be made in the case of any events which cannot be explicitly scheduled (only armed by prepulse) because of re-synchronization requirements (e.g., shaft synchronization of the Siemens motor generator).

1
00
1

Figure 2

Table 2 gives a preliminary assignment of codes generated by the supercycle generator. These decimal codes, which all have values less than or equal to 30, have been chosen so that they may be rerouted onto the individual machine time lines by direct encoding circuitry in the time line generators. The code convention assigns the tens digit to the machine (i.e., 0=Linac, 1=Booster, 2=AGS), and uses the units digit to signify the specific function. An exception occurs in the case of code 0, which is reserved for Supercycle Start.

Table 2

Supercycle Codes

<u>Code</u>	<u>Function</u>	<u>Mnemonic</u>
0	Supercycle Start	SCS
1	Linac User Reset #1	LR1
2	Linac User Reset #2	LR2
3	Linac User Reset #3	LR3
4	Linac User Reset #4	LR4
5	Linac User Reset #5 *	LR5
9	Linac Prepulse	LPP
10	Booster Cycle Start	BT0
11	Booster User Reset #1	BR1
12	Booster User Reset #2	BR2
13	Booster User Reset #3	BR3
14	Booster User Reset #4	BR4
19	Booster Prepulse	BPP
20	AGS Cycle Start	AT0
21	AGS User Reset #1	AR1
22	AGS User Reset #2	AR2
23	AGS User Reset #3	AR3
24	AGS User Reset #4	AR4
29	AGS Prepulse	APP

* Linac local users (see below)

The supercycle definition will be managed by a special set of applications programs which will be accessible with appropriate restrictions. These programs will provide a graphical environment in which to view, modify, or create the supercycle. The changing of a supercycle will be carried out in a buffered fashion between the high level code and the supercycle generator, with appropriate checks and handshakes. The actual changeover to the new cycle will be smooth and continuous from the previous cycle in the same manner as a replay of the same one. There is no provision for shifting tables in mid supercycle, since this was judged to create too much complexity in the inter-machine scheduling.

Using the supercycle code assignments given above, several examples of possible PPM scenarios are presented in the Appendix. Two examples show only a portion of a possible supercycle time line, while the others have been developed to show complete supercycles. A constraint has been imposed on the supercycle architecture to simplify end effects associated with a change to a new (and different) supercycle. Sufficient local Linac cycles are used to accommodate required initial user resets for the Booster and the AGS. The occurrence of a user reset will generally be scheduled as soon after the end of a previous user as possible, as a signal for completion of buffered acquisitions. At the beginning of the supercycle, another user reset must be scheduled. This reset will override the previous, now obsolete, assignment in the case when a change is made to a different supercycle.

2.2 Booster Time Line Controls

The Booster Real Time Line and Gauss Time Line Generators each contain four tables of event times and codes which may be scheduled for the current cycle according to its user reset code. These tables will be managed by the controller and the value for a particular delay will be controlled rather like a presently-used autotet. One new feature of this controller will be its ability to schedule multiple instances of a particular coded event in one cycle. For example, it will be possible to load an array of instrumentation acquisition times under the control of an application program. The sorting of new values into the executable table of events will be carried out by the time line generator device controller.

At the receiving end, in a device controller or in a timing distribution chassis, a decoder unit capable of recognizing several codes will produce pulses for distribution or for processor interrupts on a separate output line for each code. In the case of the Real Time Line (RTL), the master clock (10 MHz) will be extracted from the carrier for local use. Since some accelerator equipment requires a cascade of delays to establish a significant event (e.g., a momentum - Gauss clock - trigger, followed by a real-time delay), local delay hardware will be provided in those cases. Local delays will often be required to have PPM controls. For the example mentioned, the primary trigger might appropriately be centrally generated and distributed on the Gauss Time Line (GTL).

The GTL generator will generate events corresponding to desired field values and the sign of the derivative of field with respect to time, since it will utilize true up-down counting. Many device controllers will use only the RTL connection, but for those which need both RTL and GTL, the

local receiver-decoder card will permit recognition of up to 24 codes on each of 2 time lines. The 24 choices will be split into 8 interrupts, 8 direct trigger pulses, and 8 locally delayed triggers; where the same code may be used in more than one of the groups. The principal uses of the GTL system will be in the areas of injection (PKR, etc.), extraction and for various instrumentation triggers.

2.3 Local Timing Controls

Requirements for local timing services will be met by the use of delays (autodets) located in local controllers either on receiver-decoder boards or on autodet boards. A Linac timing interface will convert the user reset #5 into a pulse on one of three local lines to designate the local user of Linac beam (e.g., #5-BLIP, #6-REF, or #7-Studies). This feature will allow gating and/or manipulation of local Linac scheduling without involving the supercycle in all the details. PPM equipment at the Linac will be equipped with memory and database support for 7 users, of which 4 will correspond to the 4 users supported in the rest of the complex.

2.4 Backup Signals

An important subject in the area of timing is the need of some accelerator systems for backup triggers. It is currently standard practice at the AGS to provide a backup Gauss clock to simulate a cycle when the main magnet is turned off. Provision should be made to simulate the Booster magnet waveform for use in generating a backup clock. The matter of backup events, especially for those scheduled on the main time lines, is problematic; in fact, the provision of a backup Gauss clock is a way of relieving this need in certain areas. In general, any specific need for a backup trigger - for instance, a safe discharge trigger or function backup stop - would be met by local gating logic which can determine whether the primary trigger arrived. This scheduling cannot easily be done by the central system, which is table driven. The RTL and GTL controllers are distinct and separate units operating on asynchronous clocks, and deterministic relationships between events cannot be guaranteed between the two time lines.

3. Implications for Controls Subsystems

3.1 Database

The accelerator controls database to be used for PPM operation will have to distinguish between particular users. These changes will be reflected in the live data managed by the stations, and in the DDF and the Interbase database.³ This distinction will be made at the device level, for both simple logical devices (SLDs), complex logical devices (CLDs), and parameters by means of a single code which may take the values 0-7 for Linac devices, and the values 0-4 for Booster and AGS devices. Non-PPM devices will be indicated by USER CODE = 0. It is envisioned that certain of the database entries for a "PPM" device are user-dependent quantities, while others which are related to general properties of the control variable are not. The following list gives a preliminary view of some quantities in each category:

User-specific Data

Setpoint/Readback
Tolerances or
Upper/Lower Limits
Watch Masks

Shared Data

Device Class Limits
Allowed Commands
Status Description
Device Description

Various user instances of a particular accelerator control variable will be implemented with the same device (parameter) name. It is envisioned that the user code will be joined with the name to form a "key" to the relevant database relations. The appropriate dissemination of information concerning PPM and non-PPM characteristics must be made between device class and device records. It may be necessary to define a new relation for PPM data which would be pointed to by the device record. Controller data structure pointers will presumably belong to both categories. It would be appropriate for device controller records and station records to contain a code to indicate maximum allowable user codes; a value of zero here would indicate no PPM devices. Linac will have a different maximum number of users than the rest of the complex.

A question which must be resolved in establishing a scenario for the implementation of PPM controls is whether the same DDF should encompass all devices or whether a new separate DDF should be established for all data belonging to PPM stations. Such a separation in functionality might make the transition less painful.

3.2 Device Controllers

The most fundamental change in device controller function is the addition of the facility for user resets via the

machine time lines. Interrupts will be generated in a time line receiver-decoder card, upon which the controller will set up indexing for its next cycle transactions with the equipment and the storage of data acquisitions. As a rule of thumb, the receipt of the code tells the controller what to do next; while the scheduling of the changes of pointers, etc., must be made consistent with the real-time schedule for the device. Once the next user has been set up, non-user-specific pulses from the time lines, with or without local delays added, will provide start, stop, and other synchronization signals.

The principal change in the device data structures for PPM will be the use of the presently designated MODE byte in the packet header for the PPM user information. A particular SLD transaction will then be addressed via the logical device (LD) number and the mode byte which will contain a bit mask. In reports, this mask will always have a single bit set indicating the particular user to which the report refers. The usual command will also contain only one bit set; but, especially for maintenance or initialization operations, it will be possible to send a command to any combination of valid users in one transaction. This transaction protocol will be equally applicable to CLDs. In command packets and in default reports (except for LD#0 as discussed below), the value 0 in the user field means non-PPM. For PPM SLDs in the device controller each SLD structure in the default report will have a depth corresponding to the maximum number of users. The full default report will therefore contain all possible fields. In order for the station to avoid scanning and reporting data which is invalid, the device controller will pass information on valid user cycles since the last scan through the LD#0 ("device controller" device) report. In this data structure, the MODE byte will contain an 8-bit mask which will contain a non-zero bit for each user for which a cycle was executed.

CLD reports will be made on an individual user basis, with the mode byte again being used to indicate the user code. In special reports from certain controllers at the Linac, where high data rates make extensive buffering necessary, the MODE byte = 0 means either "device" is non-PPM, or data from all cycles is included. For the Booster, groups of cycles will occur for proton injection at a frequency of 7.5 Hz. The packet switching character of the network data transactions necessitates the buffering of multiple pulses for CLD reporting purposes. The occurrence of a reset event will signal the termination of current buffer filling and the initiation of a report transaction (Special Report Ready).

An additional change which may be desirable for the device controllers is the implementation of a standard time stamping method for device reports. The issue of whether this facility needs to be implemented only for CLD reports or

also for select SLD reports needs further discussion. In any case, it would be appropriate for a timing event from the supercycle generator (Supercycle Start) to reset all controller cycle number registers.

3.3 Stations

As discussed above, many of the changes to the stations which are required for the support of PPM concern database and data structures. Tables of pointers and other data will have to be divided into those parts which use only the device name as a key, and those which use the structure corresponding to the name plus user code as the key. The proposed handling of simultaneous commands to multiple users for a given SLD will mean additional looping will have to be built into the update of the live database.

The major change for the station code, in addition to the database aspects, comes in the area of scan task management. It is undesirable for the station to report invalid data in response to a report request; but there is no guarantee that a particular user cycle has occurred since the last scan. In fact there will typically be no cycles at all of some users during the execution of a particular programmed supercycle. For this reason, the device controller, as described above, will report on all users active since the last default report was made. The station watch task and the deferred report task should use that mask in LD#0 to only act on the valid user data. In contrast, an immediate report request of the station should elicit a normal response since, by definition, that data should be what the station has currently in its database. It might remain for further discussion, whether or not a total absence of data or very stale data should be indicated somehow even in an immediate report. The question of time stamp usage (i.e., current cycle number or cycle on acquisition) might need review.

Because the Booster is capable of multiple uses during one AGS cycle and may also run independent of the AGS, it will be important to appropriately schedule Booster controller scans. Since the station scan cycle is timed from AGS T_0 , it will be necessary to make available a transaction by which the schedule may be changed to accompany a new supercycle setup, without re-initializing the whole controller database in the station. These same considerations apply to the Linac, but with somewhat less importance.

3.4 Host Software

The principal aspect of the application of PPM to the host or console environment is the assignment of a **user context** for that environment. Either by individual program, or by console as a whole, the operating individual will be

queried by the system to choose a user context. This context will correspond to a specific user code referenced by a user name given at supercycle setup. For instance, user code #1 might be assigned to "AGS Physics (HEBT)" and user code #2 might be "Booster Commissioning". There will be no pre-assigned designations within the four codes available to the Booster/AGS. The Linac codes of 5-7 will be equally general but reserved for local use. Of course, it may become common usage to assign user code #1 to the physics program, if it is running or not. However, that is not dictated by the controls structures. Having established a context for the task, it is envisioned that the basic logical device and parameter tools referenced by the program(s) will use this information along with device names to produce the requisite device keys. Within the process environment, the operator may then think in terms of names alone when referencing the relevant PPM and non-PPM devices.

One immediate consequence of this point of view is that the entire machine tree and archive structure is independent of PPM. The tree reflects, as it has previously, the process areas of the accelerator complex. The application of the user context to a particular process consists merely of associating the user code with the device names found in the tree. One result of this distinction is that there are not separate "user" areas for Protons, Heavy Ions, and Polarized Protons; but, there are, as appropriate, separate tree areas. Remembering that user code assignments are entirely dependent on immediate usage in a supercycle, one might imagine (some time hence) a supercycle with several users, all utilizing proton beams, each using some areas of the tree differently.

In the same spirit, archives are really identified with a user only in the separate contexts in which they are saved and restored. An archive would be created within a given context and would contain the current user name as assigned at supercycle definition. The restoration of an archive would be from the full spectrum of choices, across all users. The user context applied to the restoration would be that of the operator restoring the archive, not that in which it was saved. A simple example comprises the case of the restoration, in the context of the physics program, of the results of a successful study seeking some improvement in the running conditions. It will be desirable for there to be an option available to the operator on whether or not to restore non-PPM devices along with PPM devices. Non-PPM devices would always be archived, however.

Acknowledgments

The author would like to thank those who have had a hand in helping to define PPM as it applies to the AGS complex: Leif Ahrens, Tom Clifford, Bryan Culwick, Elise DeCarlo, Brian Oerter, Ken Reece, Joe Skelly, and Alan Stevens.

References

1. Clifford, T., Frankel, R., and Stevens, A.J., Overview of the AGS Distributed Control System, October 1986, ACS Tech Note ACS/TN/86/001 (1986).
2. Stevens, A.J., Adapting the Apollo Control System to a PPM Environment, ACS Internal Note ACS/IN/87/002 (1987).
3. Griffiths, C. and Katz, R.A., Database Descriptions and Applications, ACS Internal Note ACS/IN/89/001 (1989).

Appendix

Some examples of Supercycle Time Line Use

Fig. A1: This diagram shows a representative segment of a time line program which may be in use during Booster commissioning. The AGS is running the SEB program with beam injected along the old HEBT line. This setup is running as User #1. In turn, the Booster is accelerating four successive pulses of protons, but without transfer to the AGS. This program is assigned to User #2. Users #3 and #4 are unused in this example, and codes related to these options do not appear on the time line. As a background operation, the Linac operates its local program according to local control options applied to the User #5 cycles scheduled during the supercycle. The various peaker (PKR) events indicated might come from the relevant machine's Gauss Time Line (GTL).

Fig. A2: This diagram shows a possible time line segment during operation of the AGS SEB program with full Booster injection (User #1), accompanied by a Booster study utilizing a single Booster proton cycle (User #2). The AGS magnet ramp has been shown to include not only the injection "front porch", but an additional flattop for operation of a second cycle of the dilution rf system (VHF). The AGS extraction flat-top is artificially short. The Booster study cycle was placed such that it fit on a single-page diagram, whereas it might actually be scheduled well onto AGS extraction flat-top. Adjustment of dwell periods and ramp parameters may be needed to satisfy scheduling constraints imposed by the synchronization to the line clock.

Fig. A3: This example represents a full supercycle consisting of a cycle of the AGS running the SEB program with Booster injection (User #1), accompanied by commissioning work utilizing heavy ions (User #3). An additional Booster cycle (User #2) has been specified to provide a separately tunable means of removing possible effects of different magnet history on the proton cycles. The labels BEX and AIJ are simply indicators of the process of Booster-AGS transfer. Reset events have been placed so as to optimize buffered data reporting. In addition, it has been necessary to accommodate the change in injection field for the Booster between the two modes of running. In this case, the proton injection field is the "baseline", with the heavy ion cycle matched to it.

Fig. A4: A final example shows the AGS FEB ramp (User #1) scheduled along with a test of Booster polarized proton accumulation (User #2). In this example, there is very limited use of the Linac for local clients. The label BIJ indicates the injection of separate pulses of polarized protons from the Linac.

FIG. A1: AGS (HEBT INJ) & BOOSTER COMMISSIONING

FIG. A2: AGS (BOOSTER INJ) & BOOSTER STUDY

FIG. A3: AGS (BOOSTER INJ w/ SETUP CYCLE) & BOOSTER H.I. COMMISSIONING

FIG. A4: AGS (BOOSTER INJ) & BOOSTER POLARIZED PROTON COMMISSIONING