A PHYSICS BASED APPROACH FOR MAGNET CONTROL IN A BOOSTER AND STORAGE RING *

S. M. Hartman†, S. F. Mikhailov, Y. K. Wu,
Free Electron Laser Laboratory, Duke University, Durham, NC 27708, USA

Abstract

At Duke University, a booster synchrotron was recently commissioned as part of the High Intensity Gamma-ray Source (HIGS) upgrade. For the ramping magnet power supply controls, we followed an approach previously implemented for the Duke storage ring controls. The high-level operator interface is presented in terms of the physics quantities of the accelerator, i. e. the effective focusing strengths of the magnets. This approach allows for a tighter integration of the control system with physics modeling programs and facilitates machine studies and operation. The approach also simplifies operations of the accelerators by presenting an operator interface nearly independent of machine energy. For the booster, nonlinearities of the magnets, the result of its extremely compact footprint, are incorporated into the low-level software while providing a high-level of machine tunability. For the storage ring, feed-forward compensations built on the effective strength of the magnets simplify tuning of the magnetic lattice over a wide range of electron beam energies or wiggler settings. This approach provides for a good match to the diverse operational modes supported by the Duke storage ring.

INTRODUCTION

Accelerator control systems are typically built in terms of engineering units, controlling magnet power supplies in terms of their currents, for example. Physics applications and modeling programs, on the other hand, work in terms of the physics units of the accelerator, such as effective bending or focusing strength of magnets. To bridge these two systems, the current trend has been to develop a layer of middleware. The middleware serves as a common interface for various high-level physics applications to interact with the control system. At the Duke Free Electron Laser Laboratory, we have chosen a different approach. Instead of developing high-level, host-based applications to bridge the divide between engineering and physics, we have instead moved the mapping to the low-level of the control system, into the front-end computers. By moving this layer to the front-end computers, we can take advantage of their real-time operating system characteristics, providing predictable and repeatable responses to complex machine tuning. In addition, we can reduce the sensitivity to network load or responsiveness by incorporating related systems into a few front-end computers.

We have moved towards developing the control system to present the high-level operator interface in terms of the physics units of the machine. To this end, all magnet controls for the storage ring and its full-energy injection booster synchrotron are presented in terms of their effective focusing strengths. Rather than presenting magnets in terms of power supply current, they are presented in the same way an accelerator physics model would present them. Orbit trims are in mrad, quadrupoles in m^{-2}, sextupoles in m^{-3}, etc. The low-level software encapsulates any non-linearities of these elements and presents a nearly energy-independent accelerator interface. This machine interface is meant to match the virtual accelerator of the model, allowing easy transition from simulation to running accelerator.

STORAGE RING

The Duke storage ring is a 0.24 to 1.2 GeV electron storage ring used as a driver for UV-VUV Free Electron Lasers (FEL). The FELs are also the driver for a Compton gamma-ray source. The Duke storage ring operates over a wide range of parameters (see Table 1) and supports a user program requiring frequent changes of electron beam energy and photon or gamma energy. It is not unusual to operate the machine at several energies during a single day.

Table 1: Basic operational parameters of the Duke storage ring.

RF Cavity Frequency:	178.5472 MHz
Harmonic Number:	64
Bunch Pattern:	any arbitrary pattern
Operation energy:	0.24–1.2 GeV
Circumference:	107.46 m
Single Bunch Current (max.):	95 mA
Multi-bunch Current (max.):	280 mA
Wiggler K (OK-4 FEL, lin. pol.)	0–5.0
Wiggler K (OK-5 FEL, circ. pol.)	0–4.5

For the Duke storage ring, the initial motivation to present the accelerator in terms of the physics quantities was to provide simpler and more consistent energy ramping. Prior to the commissioning of a full-energy synchrotron booster in 2006, injection into the ring was from a 270 MeV linac. The electron beam energy was

* Supported by US DoE grant DE-FG02-01ER41175 and by US AFOSR MFEL grant FA9550-04-01-0086.
† hartman@fel.duke.edu

ramped in the storage ring to operation energy, from 270 to 1200 MeV. By moving the physics to engineering mapping to the low-level systems, the ramping process was greatly simplified. It became possible to step a single variable, beam energy, and fan it out to all relevant front-end computers. The front-end computers would calculate the new magnet current setting to maintain the same effective focusing strength at the new energy. Non-linearities in the magnets, due to saturation at higher energies, were dealt with by using lookup tables for the mapping of field strength to current. Ramping beam energy in a linear manner resulted automatically in ramping power supply current in the appropriate non-linear manner. Similarly, ramping from one lattice to another can follow a linear K-value curve. This minimizes betatron tune shifts as compared with ramping quadrupole current values linearly [1].

The approach also became a useful way to deal with combined function magnets. The Duke storage ring is unique in using arc quadrupole magnets, with inner and outer coils driven asymmetrically, to produce the needed sextupole field. The difference between inner and outer coils can be as high as 40% in these magnets. The high-level operator interface tunes these magnets simply in terms of K_1 and K_2. The low-level software transparently handles the appropriate mapping to the multiple power supplies involved to produced the desired field.

Another benefit came in the simplification of the building of feed-forward compensation schemes. Compensation tables are built into the low-level software to provide appropriate compensation using trim magnets for any field differences in magnets driven in series. For the dipole magnets, for example, with all the main coils driven in series, the trim coils are set to provide appropriate compensation for the individual magnets. Similar feed-forward compensation is used for quadrupole magnets. During the tuning of beam energy or the lattice, these compensations are automatically propagated to relevant trim magnets.

Insertion devices can also be compensated in this manner. The OK-4 FEL, an optical klystron configuration with two long wigglers separated by a buncher magnet, can produce an unwanted quadrupole effect resulting in betatron beating and tune shifts at different wiggler settings. To maintain lattice matching and tunes, a set of ΔK_1 values were calculated for nine families of quadrupoles in the wiggler straight section. Simply adding this ΔK_1 value to the lattice's K_1 value in the front-end computer results in an effective feed-forward scheme. This makes it possible to tune the wiggler over a wide range ($K_{wiggler}$ from 0 to 5), with minimal impact on the electron beam.

With the OK-5 FEL, which in its final configuration will be an optical klystron with four wigglers and three bunchers, a similar scheme will be developed. An orbit correction feed-forward system has already been deployed. An additional feed-forward for non-linear effects using octupole magnets is under development.

Care has been taken to cluster related systems into the same front-end computer as much as practical. This reduces the required network load and sensitivity to network variations. For example, tuning the OK-4 FEL field adjusts the wiggler setting, but also results in nine quadrupole families, main coils plus trims, tracking the wiggler setting to maintain lattice parameters. Control for all of these quadrupole families is clustered in the same front-end computer. At the level of the operator interface, only one knob, the normalized wiggler strength, is changing. But at the machine level, eighteen power supplies are tracking the change to produce only the desired field change.

BOOSTER SYNCHROTRON

With the Duke booster synchrotron, elaborate feed-forward schemes were not required. However, the booster magnets do suffer from very significant field saturation as a result of its very compact footprint (see Table 2). This saturation, uncompensated, leads to large tune variations and changes in chromaticity during the energy ramp [2].

Table 2: Basic operational parameters of the Duke booster.

RF Cavity Frequency:	178.5472 MHz
Harmonic Number:	19
Injection energy:	0.24–0.27 GeV
Extraction energy:	0.24–1.2 GeV
Extraction energy resolution:	1 MeV
Minimum operation cycle:	1.3 seconds
Circumference:	31.902 m
Bending radius:	2.273 m
Peak dipole field (1.2 GeV):	1.76 Tesla

The control system interface to the booster offers a great deal of tunability. In addition to the ramping mode, the booster can be controlled in a steady-state mode, with a level of control similar to that expected in a storage ring. Tuning changes made to the lattice in the steady-state mode are automatically propagated to the ramping waveforms and are available for the next ramp cycle [3]. By implementing the operator interface in terms of the effective focusing strength of the magnets, the tuning changes can be propagated with appropriate energy dependent compensations already taken care of. For example, adjusting the tunes or chromaticity in the steady-state mode automatically propagates the effective focusing change to the higher energy portion of the ramping waveforms. The energy can then be ramped to a higher value and an additional correction can be made. In this manner, a tune shift or chromaticity correction can be built into the ramping waveform dynamically. The more significant energy related changes are built into the waveform, providing a smoothed curve for the newly developed waveforms.

For each ramping power supply, a ramping waveform describes the ramp in terms of the effective focusing strength of the magnets. As in the storage ring control scheme, orbit trims are in mrad, quadrupole trims in m^{-2} and sextupoles in m^{-3}. The index to the waveform array defines the energy and provides the synchronization between the elements. When an array of these physics units is tuned, the lower-level software running on the front-end computer calculates a new array corresponding to the power supply current needed to produce this field at each energy. This engineering units array is then used to reprogram the ramping DAC, a VME based arbitrary waveform generator. The arrays provide a linear energy ramp from 0 MeV to 1.2 GeV.

From this approach, it then becomes a simple matter to ramp to or from any arbitrary energy. In the high-level operator interface, variables for injection energy (to match the linac energy) and extraction energy (to match the storage ring energy) can be easily adjusted. These variables are then used to program the start and stop points of the ramping waveforms to produce the energy ramp. In this manner, the booster can be quickly set to extract at any energy from injection energy to the 1.2 GeV maximum. We have also taken advantage of this flexibility to expand the energy range for injection. For the nominal 280 MeV linac, we expected booster injection in the 270 to 280 MeV range. Due to an issue with a linac accelerating section, we have injected in to the booster at energies as low as 230 MeV. Adjusting the booster injection energy is much simpler and quicker than optimizing linac energy, so the booster can simply be adjusted to track any changes in injection energy. This has greatly simplified operation of the injector system.

CONCLUSION

With the non-linear mappings built into the front-end computers, the physics quantities can be added in a simple linear manner. For example, a quadrupole's setting may be the sum of the lattice K_1 value plus multiple corrections (ΔK_1 for tunes correction, wiggler compensation, etc.). Another example is an orbit corrector, whose value may be the sum of an operator tuning value, a compensation value, and a correction from a high-level orbit feedback program, all expressed in mrad. The nonlinear mapping of this sum to the power supply current only needs to happen once on the front-end computer rather then for each correction value. This ensures synchronization for the various inputs and provides a system capable of dealing with inputs from multiple high-level applications.

By moving the mapping between the physics model and the engineering machine to the low-level front-end computers of the control system, we have developed very flexible and robust accelerators. The storage ring and booster can easily be tuned over a wide range of operation parameters, simplifying the task of accelerator setup. Compensations built into the low-level software make operation of the machine much simpler. Using beam-based techniques, the mapping can be refined and the control system updated over time as necessary.

The overall approach is to create a "virtual accelerator" using standard control system tools. This allows the operator interface layer to present the accelerator in the terms used in the physics model. High-level physics applications can be developed without concern for engineering implementation. The virtual accelerator provides an intuitive interface for accelerator physics studies and for accelerator operations and tuning.

REFERENCES

[1] Y. K. Wu, S. Hartman, S. F. Mikhailov, "A Physics Based Control System for the Duke Storage Ring" Proceedings of IEEE Particle Accelerator Conference (PAC2003), Portland, Oregon, 12–16 May 2003, 2482–2484 (2003).

[2] S. F. Mikhailov, S. M. Hartman, J. Li, V. G. Popov, Y. K. Wu, "Challenges for the Energy Ramping in a Compact Booster Synchrotron," Proceedings of IEEE Particle Accelerator Conference (PAC2007), Albuquerque, New Mexico, 25–29 June 2007, 1212–1214 (2007).

[3] S. M. Hartman, S. F. Mikhailov, V. G. Popov, Y. K. Wu, "A Physics Based Approach for Ramping Magnet Control in a Compact Booster," Proceedings of IEEE Particle Accelerator Conference (PAC2007), Albuquerque, New Mexico, 25–29 June 2007, 515–517 (2007).