

---

---

## Machine Protection Systems

### Contents

---

<b>16.1</b>	<b>Introduction . . . . .</b>	<b>856</b>
<b>16.2</b>	<b>Single Pulse Induced Failure . . . . .</b>	<b>856</b>
16.2.1	Diagnostic Pulse Protection . . . . .	858
16.2.2	Transition between diagnostic pulses and full beams . . . . .	863
16.2.3	Controlling the Interpulse Difference (MAID) . . . . .	865

---

## 16.1 Introduction

---

One of the most serious operation issues that any future linear collider will face is that of the Machine Protection System (MPS). To produce useful luminosity the beam power and the beam densities must be very high. Unfortunately, these beams will almost certainly damage any material that is intercepted unless extreme care, such as that in the collimation sections, is taken. For example, in the 1 TeV NLC design, the beam power is over 8 MW and a single errant bunch train in the linacs would be sufficient to damage many unprotected accelerator structures. Obviously, this has severe implications on the beam operation during normal running as well as during tuning and commissioning.

The purpose of the MPS is to prevent damage to the collider components in the event of a routine failure or mistake. In addition, it should automatically provide a sequence of beam pulses that can be used as effective diagnostic tools during a startup or a fault period. Furthermore, after a fault, the MPS should be optimized to recover luminosity as quickly as possible in order to minimize the lost time. Of course, the system is *not* intended to provide comprehensive protection against any possible failure; the complexity of the MPS must be balanced against the cost, difficulty, and time for repair of the systems it protects.

The most serious challenge in the NLC is the prevention of 'single pulse induced failure' (SPIF). This is component failure that occurs from an aberrant single beam pulse. Because it is impossible to know the precise trajectory of the upcoming pulse, the MPS must: 1) provide pulses that cannot cause SPIF for tuning and diagnostics and 2) insure that the difference between the upcoming pulse and the one that preceded it is within some limit, known as the 'maximum allowable interpulse difference' (MAID), during normal operation. These two criteria form the basis of the NLC single pulse induced failure machine protection system.

Multi-pulse or 'average power' induced failure is component failure that occurs after a succession of pulses deposit excessive energy on a given component. This type of failure is more familiar from SLC operation and is controlled in a similar fashion, *i.e.*, by using ion chambers, thermocouples, etc., to monitor beam loss.

In the next sections, we will outline the methods that are used to protect against SPIF—as noted the multi-pulse failure mode is protected with a more standard MPS and thus will not be discussed further. At this time, we only have a conceptual description; in the future, we will need to have a detailed solution on an element-by-element basis with greater margins than is outlined in this section. Furthermore, we have only considered protection in the main linac and downstream; we have not considered the MPS issues in the damping rings or bunch compressors although the principles will undoubtedly be similar. Finally, additional detail on the MPS can be found in Sections 7.8, 8.6, 8.7, and 9.2.2.

## 16.2 Single Pulse Induced Failure

---

The strategy to protect against SPIF is to use single bunch, nominal intensity, pulses for most diagnostic purposes and at any time the interpulse difference might be outside the MAID. This strategy results from the extreme energy density of the full intensity, multibunch, NLC beam. It is not practical to build mechanical systems that can withstand the nominal NLC beam except in isolated cases such as the collimation region. It is, however, possible to develop structures that can withstand the impact of single bunches, albeit with somewhat increased emittance. Once such structures are realized, the problem becomes one of ensuring that successive pulses are alike. Thus, the single-pulse protection (SPIF) can be subdivided into protection against a diagnostic bunch, a full current, high emittance single-bunch beam, protection against a multibunch bunch train, and the method of transitioning between the diagnostic and normal operating modes.

The protection against the diagnostic bunches is based upon a passive system consisting of thin spoilers which will increase the beam angular divergence so that, by the time the diagnostic beam strikes another component, it will not cause any damage; to prevent damage to the spoilers themselves, the emittances of the diagnostic beam must be increased by a factor of ten from the nominal beam emittance. Thus, the spoilers will allow the beamline to survive the transport of a single diagnostic bunch without regard to the state of the beamline hardware. All preliminary beam-based alignment, tuning, and diagnostics will be performed using a diagnostic bunch. Of course, if beam is being lost during the transport, these operations would need to be performed at low repetition rates to reduce the average power deposition.

After a diagnostic beam can be transported to the beam dumps without difficulty, *i.e.*, after establishing the initial beam-based alignment, the beam trajectory, feedback setpoints, and energy profile, the repetition rate can be increased to the nominal 120–180 Hz. At this point, the beam emittance can be decreased to nominal and additional bunches can be added to the bunch trains. To verify the beam loading compensation is set properly, the train length will be increased in steps. All subsequent tuning must be performed at the high repetition rate; only the diagnostic beam can be transported at low repetition rate.

To prevent the high rate beams from striking accelerator components, a trajectory window of roughly  $\pm 200 \mu\text{m}$  and an energy window of  $\pm 10\%$  are established about the nominal values; this is the MAID. If the beam deviates beyond these limits in any single pulse, the collider is returned to the high-emittance single-bunch diagnostic mode while the source of the problem is diagnosed from data that was taken during the errant pulse. Extensive logic will be used to prevent erroneous MPS faults due to bad BPMs readings.

This system relies on the fact that there are no transverse deflecting fields that can change sufficiently, within a single interpulse period (roughly 8 ms), to deflect the beams from the operational trajectory window into the accelerator structures or beamline elements. In most cases, this is attained by limiting the strength of all fast correctors and limiting the decay time of the quadrupole and bending magnet fields by using solid core magnets and thick conducting vacuum chambers. The few DC magnets whose fields could change too quickly will have to be interlocked directly to the MPS using a Hall probe or similar diagnostic—separate diagnostics, such as precharge monitor, will be needed for strong pulsed kicker magnets like the damping ring kickers.

In addition, the MPS must verify that the rf systems are operational and correctly phased before the beams are launched into the linacs. However, fairly large energy deviations can be tolerated. For example, a 20% energy deviation in combination with 100- $\mu\text{m}$  random quadrupole misalignments, which are well in excess of what we expect, would only cause 1-mm orbit offsets in the linac.

Thus, the MPS system must only verify that 85% of the rf systems are operational and correctly phased. To this end, all modulators will be polled roughly 100- $\mu\text{s}$  before beam time. At the same time, the klystron timing and phase information will be checked. If there are a sufficient number of failures, the beams will be aborted downstream of the damping rings; this verification procedure is described in greater detail in Chapter 8. Finally, there will be sacrificial spoilers and dumps located in the linac diagnostic station chicanes to prevent energy errors larger than 25% from propagating further down the linac; these are described in Chapter 7.

In the next section, we will describe the mechanical protection system and the diagnostic pulse generation. After this, we will discuss issues associated with the transition from the single diagnostic bunch mode to the full bunch train and then we will describe controlling the MAID.

	<i>Be</i>	<i>C</i>	<i>Al</i>	<i>Ti</i>	<i>Cu</i>	<i>Fe</i>
Radiation Length [cm]	35.7	21.7	9.0	3.7	1.4	1.8
$dE/dx_{min}$ [MeV/cm]	3.1	3.6	4.4	7.2	12.8	11.6
Specific Heat [J/(cm <sup>3</sup> °C)]	3.3	1.9	2.5	2.4	3.5	3.8
Melting Point [°C]	1280	3600	660	1800	1080	1530
Stress Limit [°C]	150	2500	140	770	180	135
$\Delta T$ [°C]	3240	6530	6060	10330	12100	10520
$\Delta T/T_{melt}$	2.5	1.8	9.2	5.7	11.2	6.9

**Table 16-1.** Spoiler material properties and temperature rise due to a single bunch of  $1.3 \times 10^{10}$  having emittances of  $\gamma\epsilon_x = 3 \times 10^{-6}$  m-rad and  $\gamma\epsilon_y = 3 \times 10^{-8}$  m-rad at end of the 500 GeV linacs; these emittances are the smallest values possible and correspond to an rms beam size of  $\sqrt{\sigma_x\sigma_y} = 3.1 \mu\text{m}$ .

## 16.2.1 Diagnostic Pulse Protection

### Spoiler Materials

The purpose of the spoiler system is to allow a nominal intensity, single bunch, beam to be transported throughout all systems at low repetition rate without concern for damage. As stated, this will be performed using thin spoilers. If the beam is steered sufficiently far off-axis to intercept an accelerator element, it must first pass through one or more spoilers. These will increase the beam angular divergence so that, by the time the single bunch beam strikes another component, it will not cause any damage.


The length of the spoilers is a trade-off between the increase in the multiple scattering and the heating due to the electromagnetic shower in the spoiler. The projected angular distribution of the beam after a spoiler can be described by a gaussian distribution [Particle Data Book]

$$f(\theta_{x,y})d\theta_{x,y} = \frac{1}{\sqrt{2\pi}\theta_0} \exp^{-\theta_{x,y}^2/2\theta_0^2} \quad \theta_0 = \frac{13.6 \text{ MeV}}{E} \sqrt{t} (1 + 0.038 \ln(t)) \quad (16.1)$$

where  $t$  is the length of material in units of the radiation length. Thus, we expect a 0.2 radiation length spoiler to increase the beam angular divergence of a 500 GeV beam by  $11.4 \mu\text{r}$ ; this is hundreds of times larger than the incoming angular divergence.

Table 16-1 lists properties for a number of different materials including the  $dE/dx_{min}$ , the specific heat at room temperature, the melting point, and the stress limit. The stress limit is based on the tensile strength, the modulus of elasticity, and the coefficient of thermal expansion for the material. When a beam strikes the material, there is a sudden local temperature rise that may create local thermal stresses. If the temperature rise exceeds the stress limit, micro-fractures can develop in the material. In addition, it has been observed in experiments, that if the local temperature rise exceeds four times the stress limit, the shock wave due to the thermal rise will cause the material at the surface to fail completely or “delaminate” [Walz 1973, Walz 1996].

In the spoilers, we are not actually concerned by micro-fractures or deformations that might develop when the temperature exceeds the stress limit. These will not degrade the performance of the spoilers and would likely be partly re-annealed with further heating. Thus, the allowed temperature rise is limited by either the melting point of the material or four times the stress limit at which point the material will fail catastrophically.



**Figure 16-1.** EGS simulations of temperature rise in 0.2 radiation length long  $Ti$  spoilers with a single bunch of  $1.3 \times 10^{10}$ . The dotted and solid lines show the temperature rise in the front and back halves of the spoilers for six different incoming beam sizes:  $\sigma_r = 0, 3.1, 10, 20, 30,$  and  $50 \mu\text{m}$ .

Table 16-1 also lists an estimate of the temperature rise  $\Delta T$  assuming a single bunch of  $1.31 \times 10^{10}$  particles with a beam size of  $\sqrt{\sigma_x \sigma_y} = 3.1 \mu\text{m}$ . This beam size corresponds to emittances of  $\gamma \epsilon_x = 3 \times 10^{-6}$  m-rad and  $\gamma \epsilon_y = 3 \times 10^{-8}$  m-rad at the end of the 500 GeV linac which are the smallest emittances that could be delivered from the damping rings and assumes that there are no emittance dilutions through the end of the linacs. The temperature rise is calculated from a simple analytic model:

$$\Delta T = \frac{N}{2\pi\sigma_x\sigma_y} \frac{dE/dx_{min}}{\text{Spec. Heat}} \quad (16.2)$$

which ignores the effect of the shower buildup and the variation of the specific heat with temperature.

In both the linacs and the beam delivery, the spoilers are constructed from titanium. Although other materials such as beryllium or graphite would be able to handle higher beam densities, titanium is the most practical choice, balancing the spoiler survival against the length of spoiler required and the vacuum and handling properties. Unfortunately, the surface temperature rise due to this low emittance beam is over five times the melting temperature of  $Ti$ . Thus, to prevent damage to the spoilers themselves, the emittance of the diagnostic beam must be increased significantly.

These analytic calculations have been supplemented with EGS simulations. In Figure 16-1, the maximum temperature rise is plotted in each of eight linac spoilers separated by 1.8-m for round gaussian beams having sizes of  $\sigma_r = 0, 3.1, 10, 20, 30,$  and  $50 \mu\text{m}$ ; it is thought that the round beam case will model a flat beam with similar density although it may slightly overestimate the temperature rise. The dotted and solid lines show the temperature in the front and back halves of the spoilers and the dashed horizontal line shows the melting temperature of  $1800^\circ\text{C}$ . Notice that the temperature rise in the back half of the spoilers is roughly twice that in the front half due to the buildup of the electromagnetic shower.

Clearly, all spoilers downstream of the first will survive a single pulse of any incoming beam size. But, the first spoiler will be damaged unless the incident beam has an rms size greater than  $\sqrt{\sigma_x \sigma_y} > 10 \mu\text{m}$ . Thus, to prevent damage to the spoilers, either both the horizontal and vertical emittances of the diagnostic beam must be increased by a factor of ten or the vertical emittance could be increased by a factor of 100 to yield the required beam size.

## Spoiler Placement

In the linacs, the spoilers are primarily needed to protect the accelerator structures since these have the smallest aperture while, in the beam delivery, the spoilers are needed to protect the vacuum chamber and magnets. Unlike the spoilers, it is important to limit the temperature rise of the other accelerator elements to a value below the stress limit listed in Table 16-1. Deformations or micro-fractures in the accelerator structure irises would probably increase the multipactoring and the dark current from the structures and could distort the acceleration field patterns. Similarly, deformations of the magnet poles or coils could lead to large multipole fields or shorted coils.

We can get a first estimate of the requirements by looking at Table 16-1. Here, the temperature rise was calculated assuming a diagnostic pulse with an rms beam size of  $3.1 \mu\text{m}$ . In this case, the temperature rise in  $Cu$  is estimated to be over  $12000^\circ\text{C}$  while the stress limit is  $180^\circ\text{C}$ . Thus, to decrease the temperature to a more reasonable value, the incident beam density must be decreased by a factor of 70. Actually, the requirements are much greater because of the electromagnetic shower. In practice, we need to decrease the beam density of a 500 GeV beam by over a factor of 3000.

Now, to determine the spoiler placement, we need to determine the failure scenarios. Ultimately, we will have to consider the failure modes on an element-by-element basis but at this time we will only consider three global scenarios:

1. First, as a “worst case scenario,” we consider a FODO channel where a large deflection arises at the focusing quadrupole; a FODO array describes most of the linac beamline and much of the beam delivery beamline. Such a deflection could arise if one or more poles of the quadrupole magnet becomes shorted or if the magnet mover runs away to an extreme value, typically limited to  $\pm 1 \text{ mm}$ . This deflection will offset the beam in the next defocusing magnet which then deflects it further and possibly into an accelerator element.
2. Most other scenarios such as large energy errors or smaller amplitude deflections will cause the beam to oscillate at large amplitudes before being lost into an element.
3. Finally, there are a number of special situations which are not covered by either of these two cases. In particular, failures of the bending magnets or some specialized strong quadrupoles could directly drive the beam into an element with a very large deflection angle.


Unfortunately, there is no way to provide passive protection against this third case. Thus, these components will have to be directly monitored and interlocked to the MPS. In contrast, the second case is straightforward to protect against by placing a number of spoilers along the beamline with apertures smaller than that of all other elements. This leaves the first case which is more difficult to passively protect than the second case because there is relatively little distance in which the beam size can be increased before it intercepts an accelerator element. Of course, we could adopt the solution used for the third case, *i.e.*, direct monitoring, but this is less desirable than a passive system because of the large number of elements that would have to be monitored and the potential reliability problems.

To evaluate the requirements to protect the first case, we can calculate the distance from the defocusing quadrupole to the position at which the beam strikes the accelerator structure or the vacuum chamber  $\Delta L$ :

$$\Delta L = \frac{L_c}{2} \frac{(r_a/r_{QD} - 1)}{1 + 2 \sin \psi/2}. \quad (16.3)$$

Here,  $\psi_c$  is the phase advance per FODO cell,  $L_c$  is the FODO cell length, and  $r_a$  and  $r_{QD}$  are the aperture of the chamber/structures and the trajectory offset in the defocusing magnet.

Assuming a phase advance of  $100^\circ$  per cell, we find that, to shadow all the downstream elements, spoilers located at the quadrupoles would have to have a radius roughly  $\frac{1}{4}$  the minimum radius along the beamline. Although this may



**Figure 16-2.** Possible trajectories of a single bunch through a beam line as a result of completely or partially shorted quadrupole legs. Every trajectory encounters a spoiler at least 10m before reaching the wall or another element. The quadrupoles are assumed to be spaced at 30m and have a focal length of 20m in this example.

be the simplest solution, it implies very small aperture spoilers with correspondingly large wakefields; the transverse wakefield of a single spoiler would be comparable to that of an entire 1.8-m accelerator structure.


This becomes simpler in the beam delivery beamline where the vacuum chamber aperture can be increased to relatively large values between the magnets. In this case, the spoilers can also have relatively large radii which reduces the wakefields. Then, the placement of the spoilers must be chosen so they are located sufficiently far from the magnets, where the vacuum chamber constricts again, so that the spoiled beam will not damage the elements. This is illustrated schematically in Figure 16-2 where 0.25 radiation length  $Ti$  spoilers have been placed 10m from the magnets to increase the beam size to roughly  $\sqrt{\sigma_x \sigma_y} \sim 140 \mu\text{m}$ . The magnets and vacuum chamber will be further protected with sacrificial absorbers located immediately upstream of the chamber constriction. Both the spoiler placement and the absorbers are discussed further in Chapter 9.

The linac is more difficult to protect and the solution is different for the beginning and end of the linac. At the beginning, the low energy end, the quadrupoles are only separated by a few accelerator structures. In this case, a spoiler with an aperture radius of 1 mm,  $\frac{1}{4}$  of the iris radii, would increase the total transverse wakefield seen by the beam an unacceptable amount; as stated, each spoiler would have a transverse wakefield comparable to that of an entire accelerator structure.


Fortunately, at low beam energy the angular divergence due to the spoilers is relatively large,  $\theta_0$  scales inversely with the beam energy, and the required spot size to prevent damage to the structures is relatively small. Thus, the spoilers can be placed close to the location that the beam first strikes a structure. This is illustrated in Figure 16-3 which shows the temperature rise in the irises of an accelerator structure due a 30 GeV single bunch with  $1.3 \times 10^{10}$ . The four different curves correspond to initial beam sizes of  $\sigma_r = 10, 20, 30,$  and  $50 \mu\text{m}$ . Notice that, with an incoming  $30 \mu\text{m}$  spot size, the temperature remains less than the stress limit of  $180^\circ\text{C}$ . A  $30 \mu\text{m}$  spot size can be attained at 30 GeV by passing through a 0.2 radiation length spoiler 15 cm before striking the structure.

The picture is completed by looking at the failure modes. The magnet movers have a maximum range of  $\pm 1$  mm while in typical quadrupole designs, if a pole is shorted, the field on axis is roughly  $B_{pole}/5$  and causes a  $45^\circ$  angle deflection of the beam.

Assuming a  $100^\circ$  cells, which is stronger focusing than in most of the linacs, a 1 mm offset of the first quadrupole would cause a 1.5 mm offset in the second quadrupole. In the linac, this beam would then strike the accelerator structure irises, having a 3.9 mm radius, after 60% of the distance to the next quadrupole. Similarly, assuming a single shorted pole of a linac quadrupole with a 7 mm bore radius, the beam is also offset in the next quadrupole by 1.5 mm in both the horizontal and vertical planes and again the beam will strike the structures after 60% of the distance to the next quadrupole.



**Figure 16-3.** EGS simulations of the temperature rise in the accelerator structure irises due to a 30 GeV single bunch of  $1.3 \times 10^{10}$ ; the different curves correspond to initial beam sizes of  $\sigma_r = 10, 20, 30,$  and  $50 \mu\text{m}$ .



**Figure 16-4.** EGS simulations of the temperature rise in the accelerator structure irises at the end of the 500 GeV linac due to a single bunch of  $1.3 \times 10^{10}$  that has passed through eight spoilers separated by 1.8 m; the four curves correspond to initial beam sizes of  $\sigma_r = 10, 20, 30,$  and  $50 \mu\text{m}$ .

Both of these cases, are protected against by placing 0.2 radiation length spoilers with 2 mm radii at the end of every accelerator structure. As discussed in Chapter 7, the transverse wakefield deflection of the spoiler is roughly 10% of that due to an accelerator structure. Provided that the spoilers could accurately mounted to the structure ends, they would be aligned along with the structures as discussed in Chapter 7.

This solution works well until the beam energy is above roughly 400 GeV. At this point, the scattering due to the spoilers is small and a large amount of energy is contained in the electromagnetic shower. Thus, to be effective, the beam must intercept the spoilers a long distance before striking the irises. An example is illustrated in Figure 16-4 which shows the temperature rise in the irises due to a 500 GeV bunch of  $1.3 \times 10^{10}$  that has passed through eight 0.2 radiation length spoilers separated by 1.8 m. The four curves correspond to initial beam sizes of  $\sigma_r = 10, 20, 30,$  and  $50 \mu\text{m}$ .

Unfortunately, in both cases that were considered for the beginning of the linac, a single shorted pole or a 1-mm offset, the beam would only intercept six spoilers before striking an accelerator structure; this would lead to a temperature that is roughly a factor of two above the stress limit.

There are two solutions: first, we could adopt solution (3), direct monitoring of the magnets, or, second, we could use a single spoiler with a 1 mm radius located at the quadrupoles. Although the transverse wakefield of a spoiler with a 1-mm radius is roughly five times greater than that with a 2-mm radius, the single small radius spoiler would replace ten of the larger radius spoilers since there are ten accelerator structures between the quadrupoles at the end of the linac. Thus, the integrated wakefield is the same. Additional benefit is gained by placing these elements before the quadrupole magnets since a spoiler will induce a large energy spread in the beam which leads to a large decrease in the beam density after the beam is deflected by the subsequent quadrupole. Finally, the spoiler could be directly mounted on and pre-aligned to the quadrupoles; thus, they would be aligned to the beam when the quadrupole beam-based alignment is performed. At this time, we have not determined the best option and will reserve that decision for the future.

### Diagnostic Pulse Emittance Enlargement

As noted, to prevent destruction of the protection spoilers, the diagnostic beam emittance product  $\epsilon_x \epsilon_y$  must be increased by a factor of 100. This enlargement must be turned-on and -off at the full repetition rate of 180 Hz without significantly steering the beam. The logical place to perform the enlargement is either in or just after the damping ring. At this time, we have considered two solutions: first, we could induce a large horizontal and/or vertical dispersion oscillation in the damping ring wigglers using pulsed quadrupoles. These would be located in the dispersion suppressor sections at the end of the arcs. As an example, to increase the vertical emittance by a factor of 100, skew quadrupoles can be used with an integrated strength of 7 kGauss—these change the tunes by less than 0.01 while increasing  $\gamma \epsilon_y$  to  $2.5 \times 10^{-6}$  m-rad. Second, we could use a pulsed chicane, at the exit of the damping ring, to direct the beams through a gas filled chamber. Passing the beam through 1-m of Ar at a pressure of 1 Torr will increase the emittance product by 100. Alternately, one could use a thin foil (50  $\mu$ m) of Be although in this case the power density may be a problem.

### 16.2.2 Transition between diagnostic pulses and full beams

As discussed, a diagnostic beam, a high-emittance full-current single bunch beam, will be used to recover from a fault or to start up after an off period. To then establish full current operation requires a sequence of steps that include ramping the repetition rate, the emittance, and the number of bunches. Specifically, after the diagnostic pulse can be cleanly transported to the beam dump, the repetition rate must be increased to the nominal full rate. This is necessary since the MAID is only guaranteed for short periods. Next, the beam emittance can be decreased to the nominal and finally the bunch train can be lengthened. It is this last step that is probably the most difficult because of the changes in beam current and beam loading which can have significant dynamical and thermal effects in the damping rings and linacs. A single bunch diagnostic pulse was chosen since its transport dynamics are, in many ways, similar to that of the full bunch train. Nevertheless, the transition from the diagnostic pulse and full intensity operation must be automatically checked at several steps along the way.

## Linac Operations

When transitioning from the single bunch beam to the full bunch train care will be needed to verify the beam loading compensation. For the full bunch trains, the beam loading is roughly 25% in the main linacs. Thus, if the loading compensation is not set, the trailing bunch will have an energy roughly 25% too low.

To verify the loading, we will increase the number of bunches per train in steps. At this time, it seems reasonable to increase the number of bunches from a single bunch to 10 bunches per train. In this case the maximum energy deviation should be less than 3%, even if the loading compensation is not set. At this point, the loading profile can be accurately determined and thus it seems reasonable to increase to a full train in the next step; details of the beam loading compensation and control can be found in Chapter 8.

At each stage during the ramp the collider should operate for enough time to allow the beam-based feedback systems to stabilize. The time required should be less than a second (100 pulses).

In addition, active feedforward will be required to control the rf power delivered to the accelerator structures. Without changing the temporal profile of the rf power to the structures, they will cool significantly during full bunch train operation and will heat up after an MPS fault when operating with a single bunch. This arises because the full train absorbs roughly 25% of the rf power. As described in Section 8.7, the solution to this problem is to adjust the klystron phase profile so that the additional power from the rf pulse compressors is sent to the loads and is not delivered to the structures.

## Damping Ring Operation

The damping ring rf system has been designed to accept full current trains at the maximum repetition frequency but it is heavily beam loaded. Thus, changes in the total current must be accommodated with the appropriate beam scheduling information which includes the repetition frequency and a measure of the incoming charge which is required by the rf feedforward algorithm.

There are two sequences that can be used to ramp the rings from zero to full current. The first is used after an extended off period or for initial checkout at reduced current. The second is used for ramping to full current quickly as required, for example, following a spurious MPS ion chamber trip. In either case, the number of bunches per train will be increased in steps from 1 bunch per train (bpt) to 10 bpt and then to 90 bpt.

In addition to the feedforward algorithm to adjust the rf systems for changes in the beam current, feedforward will be needed for thermal regulation. The damping rings emit an enormous amount of power as synchrotron radiation and higher-order modes which is removed by the cooling systems. Thus, the feedforward system, likely consisting of both flow control and heating elements, will be required to prevent significant thermal fluctuations during changes in the stored beam current.

**Recovery procedure after extended off period** For single bunch operation, beam loading of the rf cavities in the damping ring may be neglected. However, the cavity loading angle must be adjusted to at least partially minimize the reflected power during extended periods of low current operation. Because the reflected power depends only weakly on loading angle with 1 bpt, the tuners positions are moved to minimize reflected power at 10 bpt. This allows for efficient ramping of the current from 1 to 10 bpt. The injection procedure is given below:

1. With the beam off and the extraction kicker deactivated, the tuners are adjusted to  $\phi_{z,10}^*$ , which is the tuning angle for which the reflected power is minimized at 10 bpt.

2. Since feedforward for reduced beam loading is not required, the tuner loops are left closed and the ring is sequentially filled with 4 trains of 1 bpt.
3. The extraction kicker is activated and the ring is operated for as long as desired with single bunches using single-turn extraction and injection.
4. The repetition frequency is then increased, if desired.
5. The number of bunches per train is then increased to 10 bpt. The principle is the same as that for full-current (90 bpt) injection, with the exception that the voltage and phase offsets for feed-forward are adjusted for 10 bpt using Eqs. 4.64 and 4.65. (If the current of the incoming beam is known, then the offsets are automatically adjusted by the feedforward controller.)
6. To resume operation at 90 bpt, the beam is turned off and the tuner setpoints are adjusted to  $\phi_{z,90}^*$  and the rapid recovery procedure is followed. (Here  $\phi_{z,90}^*$  is the tuning angle for which the reflected power is minimized at full current—Section 4.5.3.)

**Rapid recovery procedure** The fastest possible recovery from no beam to full current operation requires a minimum of 12 cycles (or 66.7 ms at 180 Hz); the actual process will be slower to allow the feedback systems sufficient time to stabilize between changes. The procedure is as follows:

1. With the tuners fixed at  $\phi_{z,90}^*$ , the extraction kicker deactivated, and feedforward deactivated, inject 4 trains of 1 bpt.
2. Activate the extraction kicker and feed-forward. Begin single-turn extraction and injection. Inject 4 trains of 10 bpt.
3. Inject 4 trains of 90 bpt. (The presence of the 10 bpt trains in the ring is taken into account in the feedforward algorithm.)
4. Close tuner loops and proceed with nominal operation.

### 16.2.3 Controlling the Interpulse Difference (MAID)

When operating with more than one bunch or when operating with the nominal beam emittances, the MPS must guarantee that the trajectory of the upcoming pulse is within the MAID. In order to do this, every device that can change the beam energy or trajectory by more than the MAID in the interval between pulses must be checked before allowing the permit to the scheduler. For transverse deflecting magnetic fields, the best way to do this is to actually prevent those fields from making large changes during the interpulse period. This is practical for all but the strongest magnets since the typical magnet  $L/R$  decay times can be made  $\sim 200$  ms. High bandwidth magnets, such as linac fast feedback correctors, must not have enough strength to exceed the MAID. One consequence of this is that low repetition rate, full train intensity operation, is not possible.

Fast, powerful, pulsed dumper magnet and extraction kicker systems must provide an electronic warning of their behavior prior to each pulse. This warning must be timely enough to stop beam extraction from the damping ring (in the case of the linac) or to fire protection dumper magnets.

The linac klystrons must also provide a summary status signal that can be used for MPS. Since the NLC linac MAID requires the energy difference to be less than 10 completion of these tests for each klystron is not required. The

klystron pre-pulse checking system must be most effective for warnings about common mode failures that affect many klystrons. On a pulse-to-pulse basis, a klystron's amplitude, phase or trigger timing can vary outside nominal limits and both must be checked. The modulator switch high voltage will have a charging step in its cycle that takes place about  $100\ \mu\text{s}$  before beam time. A comparison of that amplitude against the expected value will be used as an input to the MPS permit. The RF drive will also be tested for amplitude at about that time. The phase is more difficult to measure and will require that a phase measurement be made using an independent reference system.

To verify the beam loading compensation when going to the longer bunch trains, the train length will be increased in steps. Presently, we believe that the loading could be checked using only one intermediate train length, namely, going from one bunch per train to 10 bunches per train and then to a full 90 bunches per train; this needs further verification and increasing the steps has minimal implications for the rest of the systems. In addition, to prevent thermal changes of the accelerator structures when changing from low repetition rate to high rate and from the single- to the multibunch modes, the klystron phases are varied so that the additional power which normally would accelerate the missing bunches is dumped into loads rather than into the accelerator structures; this is discussed further in Chapter 8.

---

## References

---

- [Particle Data Book] Particle Data Group, "Particle Properties Data Book", *Phys. Rev. D*, **50**, 1173 (1994).
- [Walz 1973] D. Walz, D. Busick, T. Constant, K. Crook, D. Fryberger, G. Gilbert, J. Jasberg, L. Keller, J. Murray, E. Seppi, and R. Vetterlein, "Tests and Description of Beam Containment Devices and Instrumentation—A New Dimension in Safety Problems", SLAC-PUB-1223 (1973).
- [Walz 1996] D. Walz, "Justification for Temperature Rise and Thermal Stress Limits", NLC-Note-22 (1996).

## Contributors

---

- Vinod Bharadwaj
- John Irwin
- Michiko Minty
- Ralph Nelson
- Tor Raubenheimer
- Sayed Rokni
- Marc Ross
- Dieter Walz