Collimation of Heavy Ion Beams in the LHC
Executive Summary

12 January 2009

John M. Jowett

Summary

The "Phase 2" upgrade of the LHC collimation system should improve the collimation efficiency for both protons and ions. However the physics of heavy ion collimation is quite different from that of protons so requires a separate treatment. This document provides a summary and comparative evaluation of the options available to improve the. collimation efficiency of heavy ion beams in the LHC.

1. Introduction

The "Phase 2" upgrade of the LHC collimation system should improve the collimation efficiency for both protons and ions. However the physics of heavy ion collimation is quite different from that of protons and requires a separate treatment. This note provides a summary and comparative evaluation of the options available to improve the collimation efficiency of heavy ion beams in the LHC. Detailed discussion of the subject is available in various references; here we recall only a few general points. It is assumed that the reader is familiar with the LHC configurations of optics and collimators and the associated terminology.

The collimation of heavy-ion beams in the LHC is different from that of proton beams because heavy ion interactions with matter are intrinsically more complicated than those of protons. While protons essentially retain their identity when scattered by the material of the collimators, heavy nuclei can fragment through nuclear or ultraperipheral electromagnetic interactions (such as electromagnetic dissociation) giving rise to a spectrum of scattered isotopes. Since these have different charge-to-mass ratios from the original nucleus, their magnetic rigidities are different and they follow dispersive trajectories in the optics of the storage ring. The distribution of beam losses is different from the proton case, particularly in the dispersion suppressor sections adjacent to IR7 and IR3 which are used for collimation. Furthermore the basic principle of two-stage collimation, adopted for the LHC long before any study of ion beam losses was made, does not work: the nuclear interaction length in the primary collimator material is shorter than the distance required for the nuclei to be scattered sufficiently by multiple Coulomb scattering to reach the secondary collimators. The principal high mass fragments are therefore lost elsewhere, particularly in the downstream dispersion suppressor.

A further consequence of these differences is that a new technical approach had to be developed (mainly the ICOSIM program, originally written by H. Braun) since the simulation tools used for proton collimation studies cannot be applied to ions.

The collimation inefficiency in the LHC is conventionally evaluated by calculating the level and distribution of beam losses that would occur *if particles were being lost from the beam* at a rate corresponding to a beam lifetime of 12 min. The level of losses thus generated is then compared with the quench limit of the corresponding magnets. But not all beam losses count in this respect: for example, in stable beam conditions the dominant heavy-ion beam losses will be due to reactions occurring in beam-beam collisions which do not contribute to the collimation loss map. The loss levels attributed to the collimation should only occur intermittently, when the lifetime is short for some *other* reason. In reality, such losses will be mostly due to single-beam dynamical phenomena, such as resonances, or scattering processes. This is why the limit in beam intensity due to collimation inefficiency is often described as a “soft limit”: if the appropriate component of the beam lifetime can be kept above 30 min, say, losses a factor 2.5 times higher could be tolerated. In addition there are substantial uncertainties on the level of the quench limit as well as the nuclear cross sections and simulation methods that are used in predicting the loss maps. The final results could easily be a factor of 2-3 out in either optimistic or pessimistic directions. This should be kept in mind in the following.

When the losses reach dangerous levels they will be detected by beam loss monitors and will generally provoke a beam dump, avoiding a quench at the expense of the time needed to refill the machine. It is likely that, most of the time, the Pb-Pb luminosity of the LHC will be limited by other losses due to Bound-Free Pair Production (BFPP) between pairs of colliding nuclei. These also give rises to losses which may exceed the quench limit of certain magnets, this time in the dispersion suppressors adjacent to the experimental straight sections. These losses will be continuous and directly proportional to luminosity, giving a constant (as opposed to intermittent) risk of quenching, particularly while the luminosity is high near the beginning of a store. Insofar as good, stable beam conditions can be obtained in the LHC, this is expected to be the greater risk of beam-dump-induced interruptions of stores.

No previous collider has been limited by effects of these kinds so there is no real operational experience with which to compare our estimates of either collimation inefficiency or BFPP. However we have exploited available opportunities to test our methodology as far as possible at both RHIC and the SPS.

Finally it should be noted that, in the initial year of operation of the LHC with a reduced heavy-ion beam intensity (the “Early” parameters), we do not expect to be limited by either BFPP or losses due to collimation inefficiency. However this period will be important to quantitatively test our present understanding and refine the risk estimates for later runs with higher intensity.

Collimation inefficiency translates rather directly into a limit on total beam intensity. If the LHC is limited in this way for heavy-ion runs, the strategy will be to reduce the number of bunches, keeping the maximum possible single-bunch intensity to provide the highest possible luminosity.

2. Methodology

The SIXTRACK program used to track protons for collimation studies is not easily adapted to ions. A new program, ICOSIM is therefore used for most LHC ion collimation studies. The first version was written by H. Braun in 2003-4 and was further developed recently with the help of a technical student, N. Holden (see the detailed documentation in N. Holden, AB-Note-2008-054).

There are presently some limitations: ICOSIM does not take follow the trajectories of low-mass fragments; imperfections of the orbit and optics are not included in a satisfactory way; and the particle tracking includes some simplifications beyond those made in SIXTRACK. ICOSIM places greater emphasis on including the essential nuclear physics which is much more important for nuclear beams. However it also avoids making the thin-lens approximation for most magnetic elements.

Recent versions include several new models to simulate the nuclear interactions of the beam nuclei with the collimator material. The most frequently used method is a simplified Monte-Carlo simulation based on a large database of cross sections, not only for nuclear fragmentation and electromagnetic dissociation of the beam nuclei ($^{208}Pb^{82}$, …) on all nuclei present in all collimator materials (carbon, tungsten,…) but also all those of the large number of daughter nuclides created by these nuclear reactions (^{207}Pb, ^{205}Tl, ^{204}Hg,…). This database was originally generated by a process involving many of runs of the RELDIS and ABRABLA programs from I. Pschenichnov. Since then we have studied various other methods and have now settled upon one using driver modules from FLUKA under the control of a Mathematica program written by R. Bruce. These cross-sections have been compared with available data (at lower energies) and other methods used in past, e.g., the RELDIS & ABRABLA programs, MARS from N. Mokhov and others. It is also possible to directly use the complete FLUKA program itself within ICOSIM (this was done for the analysis of experimental studies in the SPS). In the case of protons, MARS has also been used in this way.

It has to be stressed that these calculations require extrapolation of many detailed quantities far beyond the energy frontier of present nuclear physics. As our detailed comparisons have shown (see, e.g., the recently accepted paper on ion collimation in the SPS), even within presently accessible energy ranges there are substantial differences among the various available calculations and models.

Special calculations for new collimator types (crystal, magnetic) have also been incorporated recently into the ICOSIM framework (see below).

2.1. Testing of simulations

Loss maps in our experiments with an LHC collimator installed in the SPS have been successfully predicted. A detailed paper has been written and is about to be published in the Physical Review.

The experimental data available so far is essentially limited to single-pass situations. Although a limited comparison of ICOSIM with loss data in RHIC was made by H. Braun a few years ago, it should be stressed that there is little previous experience of the combination of nuclear physics and beam dynamical phenomena relevant in the LHC ion collimation regime. Significant factors of uncertainty must be attached to all predicted loss levels, not least because the uncertainties in the cross sections for the basic nuclear processes themselves are uncertain at the level of a factor of two.

3. Summary of results with present "Phase 1" collimation system

This section presents some typical results for Pb-Pb collisions illustrating the loss maps with the present installation (so-called Phase 1) of collimators in the LHC.

The ICOSIM runs shown or referred-to in this section have been carried out by G. Bellodi, based on appropriate optical configurations and a fairly detailed model of the LHC aperture.

3.1. Collision energy

Figure 1 shows results for the LHC V6.503 optics in the Ion Collision configuration for Beam 1 at collision energy, the Phase 1 collimation setup with all the standard settings of the collimators (TCPs at 6σ, TCS at 7σ, TCT at 8.3σ and TCLA at 10σ). The caption explains what each plot represents and the same description applies to several other sets of results that follow.

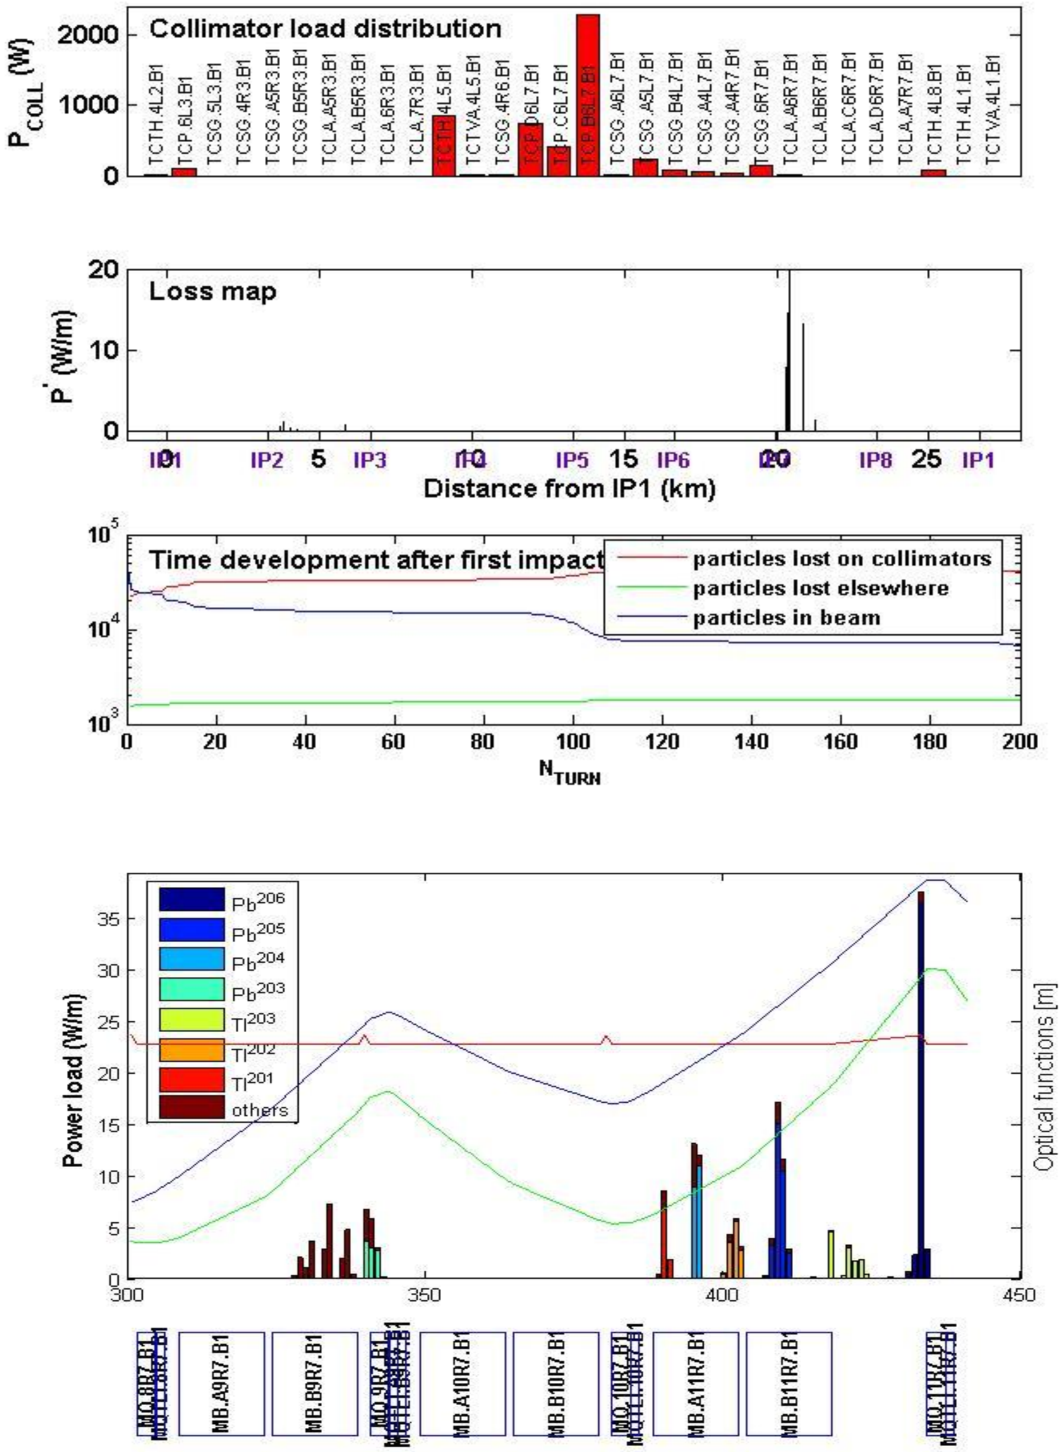

Figure 1: ICOSIM results for the LHC V6.503 optics in the Ion Collision configuration for Beam 1 at collision energy, the Phase 1 collimation setup. The top plot shows the load on individual collimators; the next plot shows the loss map in the whole ring; the third plot shows the time development of the intensity and losses during the 200 turns of the nonlinear tracking (in the sense of ICOSIM, not including magnetic imperfections) simulation and the bottom plot shows the details of the loss map in the dispersion suppressor right of IP7, where quenches are most likely. This description holds for several similar plots in later figures.

The differences with results previously shown for V6.500 are minor. The fact that the TCTVs in IR2 are not included in V6.503 might explain the few losses near IR2.

3.2. Injection and Intermediate energies in the ramp

Besides collision energy shown in Figure 1, it is also important to verify the loss maps at injection energy and at intermediate energies in the ramp. It turns out the loss maps change significantly. Figure 2 shows results for Pb beams at injection energy $E = 177A\,\mathrm{GeV}$. The distribution of losses in the IR7 dispersion suppressor is very different from collision energy, with a greater contribution of lower-mass nuclides. However the maximum power density is less and, of course, the quench limit will be substantially higher when the magnets are operated at lower field levels.

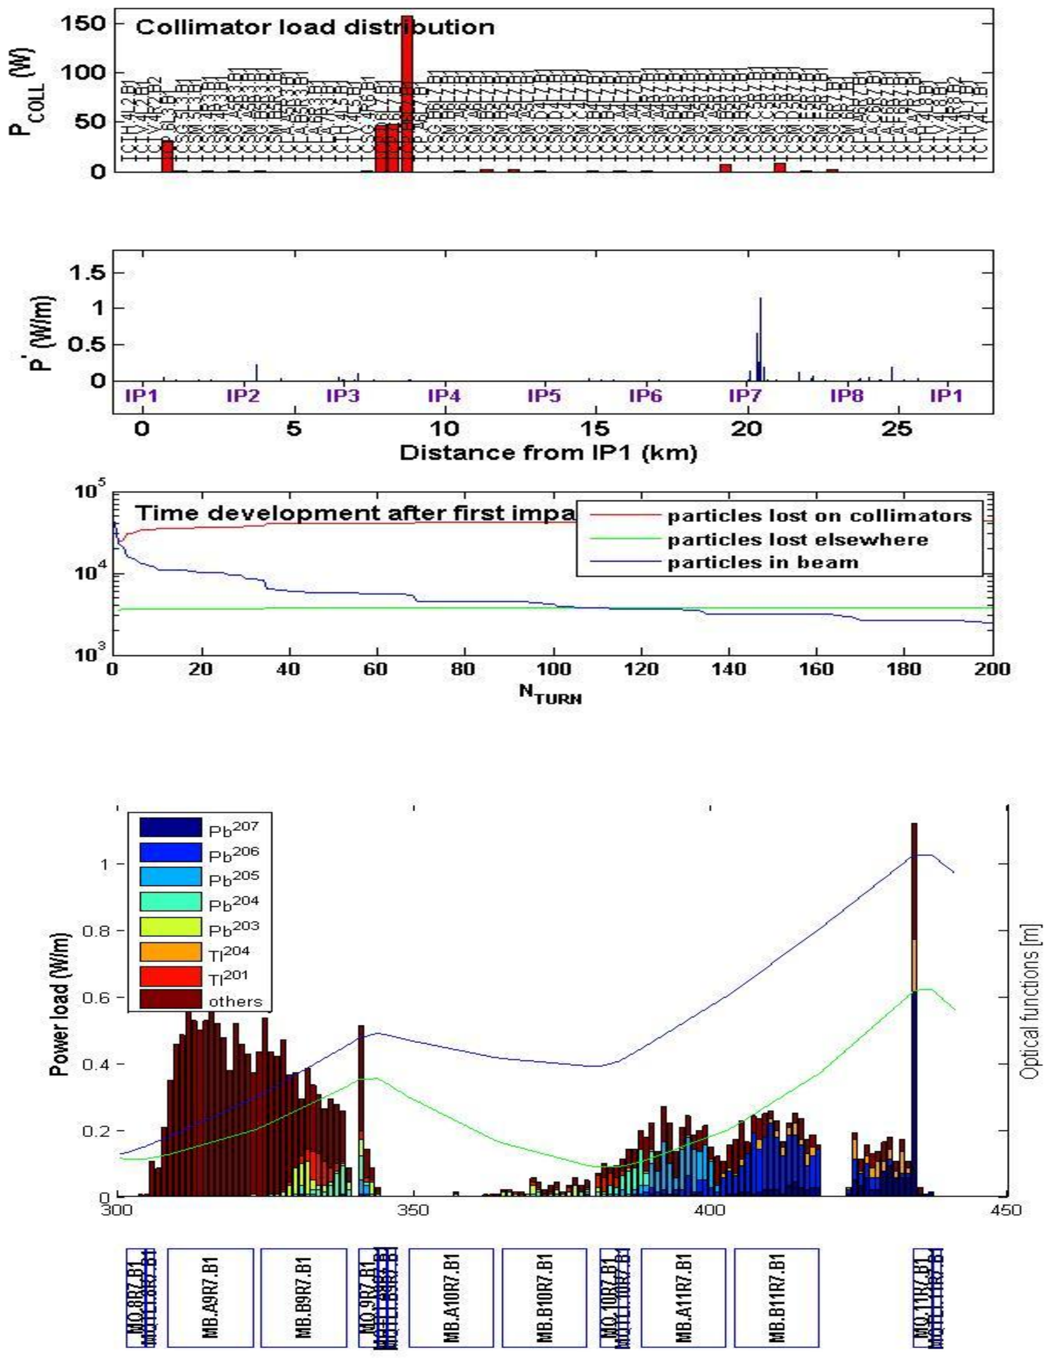

Figure 2 Standard Collimation scheme for Pb Beam 1 at injection energy.

Other energies in the ramp have been treated, showing how the broad injection loss map gradually transforms into the one characteristic of collision energy with localized discrete peaks.

These simulations are carried out with a specific scaling of the collimator gap settings. Moreover, in all cases, the initial beam distribution parameters and diffusion rate are scaled to obtain a constant $1\,\mu$m average of the impact parameter distribution on the primary collimators on the first turn. This is in fact the value which gives the worst collimation inefficiency so the simulations are maximally pessimistic in this sense; see Figure 3.

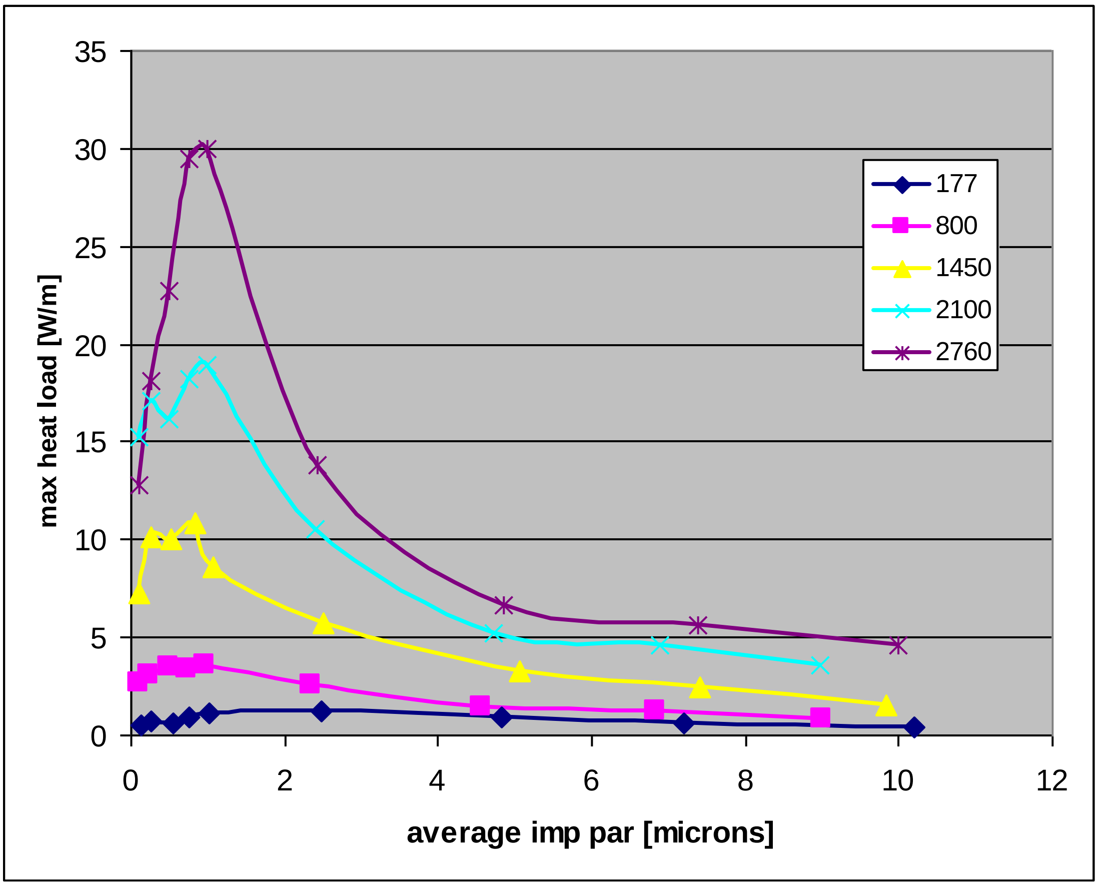

Figure 3 Dependence of the beam distribution on the TCPs (first turn) on the average impact parameter , shown for Beam 1, Phase 1 collimation scheme; the coloured curves correspond to different collision energies, the caption indicating the value of the energy in (*A* GeV).

Comparing the maximum power density in the loss map with the expected variation of the quench limit with energy, it is possible to identify the energy in the ramp where the maximum of the loss map first exceeds the quench limit. Present data suggests that this occurs around $E = 1.7A\ \mathrm{TeV} = 4.3\ Z\ \mathrm{TeV}$.

3.3. Hybrid collimation

In view of the potential difficulties (related to the irradiation of electronics) in using IR7 for betatron collimation in the initial period of operation of the LHC, the option of using IR3 for combined betatron and momentum collimation of ion beams has been evaluated (by G. Bellodi, presented at the Collimation Working Group of 22/4/2008). As expected, the scheme could be usable for the Early ion beams but does not produce any improvement with respect to the nominal configuration using both collimation insertions. While the global collimation efficiency is similar to the value for the usual collimation scheme,fact there are substantial new losses, exceeding the quench limit for nominal intensity, in other parts of the ring.

3.4. Operational considerations

In heavy-ion operation, losses occur in different places from proton operation. It is necessary to instrument the LHC to detect these in order to anticipate possible quenches and dump the beams.

Accordingly, a substantial number of additional beam loss monitors have been specified for heavy ion operation and have been installed in the ring; this is described in detail in LHC Project Note 402, Beam

Loss Monitors for Heavy Ion Operation . Note, in particular, that BLM thresholds for beam dump have been established to be the *same as for protons.*

These BLMs form an integral part of our strategy for operating the LHC with heavy ion beams.

4. Measures to improve heavy-ion collimation efficiency

The simplest and most effective way to reduce the collimation losses in cold magnets would be to increase the amplitude of the “dog-leg” transverse displacement in the collimation sections. This was rejected many years ago because it would require civil engineering work to widen the tunnel. We mention it here as it is the first item on a list of possible solutions which progress from the straightforward and highly effective, yet requiring substantial modification of the present LHC layout and installation, to compact installations of devices relying on untested principles that remain in the R&D phase.

4.1. Cryogenic Collimators

Until recently, modifications of the cold dispersion suppressor sections, to install collimator or any other substantial equipment, has been considered too disruptive. However the development of cryogenic collimators for the FAIR project at GSI and the needs for Phase 2 proton collimation led to the proposal (by R. Assmann and T. Weiler) to install such collimators in the dispersion suppressors downstream of the primary collimators in IR7 and IR3. From the point of view of ion collimation, this is also a direct method to intercept ion fragments before they are lost in the cold sections. These collimators would be installed where the dispersion function is already large and the dispersive trajectories of the ion fragments created at the primary collimators are well separated from the main beam. Similar devices could have even more important role in heavy-ion operation, intercepting secondary beams of BFPP ions emerging from IPs.

Two such collimators are required in each downstream dispersion suppressor, i.e., 4 around IP7. Their installation requires modification of cold parts of LHC. Because some dipole magnets have to be moved to make space for them at appropriate locations, the geometry of the machine is changed locally, with IP7 being moved about 0.9 mm towards the centre of the ring. A new LHC layout has been created (by JMJ) for both rings and used to rematch IR7 to compensate these changes. The new optics, shown in Figure 4, has exactly the same transfer matrix as the standard one for both rings so can be directly “plugged-in” to any existing LHC optical configuration.

The new layout for cryogenic collimators causes IP7 to move by nearly 1 mm towards the centre of the ring and reduces the circumference by 1.8 mm (it may be possible to modify these values in some further iterations of the design).

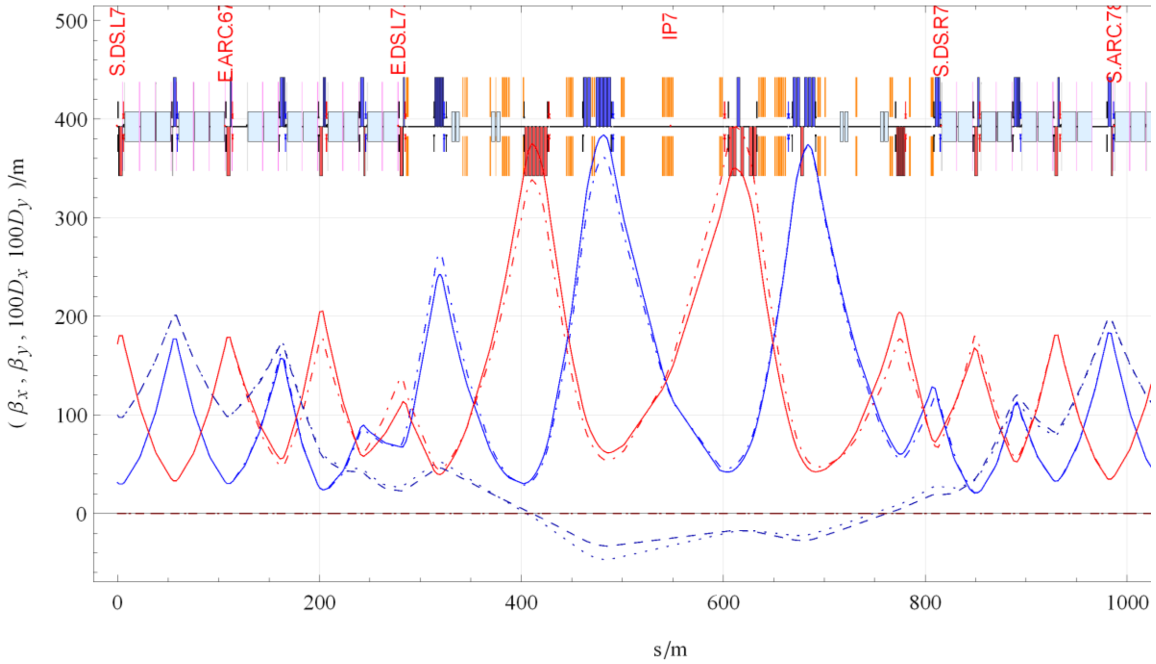

Figure 4: The new cryogenic collimator optics for Ring 1 (solid curves), compared with the old (dot-dashed curves) in the following plot, where β_x is shown in blue, β_y in red and D_x in black. Collimators are indicated symbolically in orange.

A similar optics has been matched for Ring 2. All quadrupole strengths remain within comfortable ranges and the circumferences of both LHC rings are reduced by 1.87 mm.

Next, in Figure 5, we show ICOSIM results for the new optics with all Phase 2 metallic collimators but with the cryogenic collimators retracted.

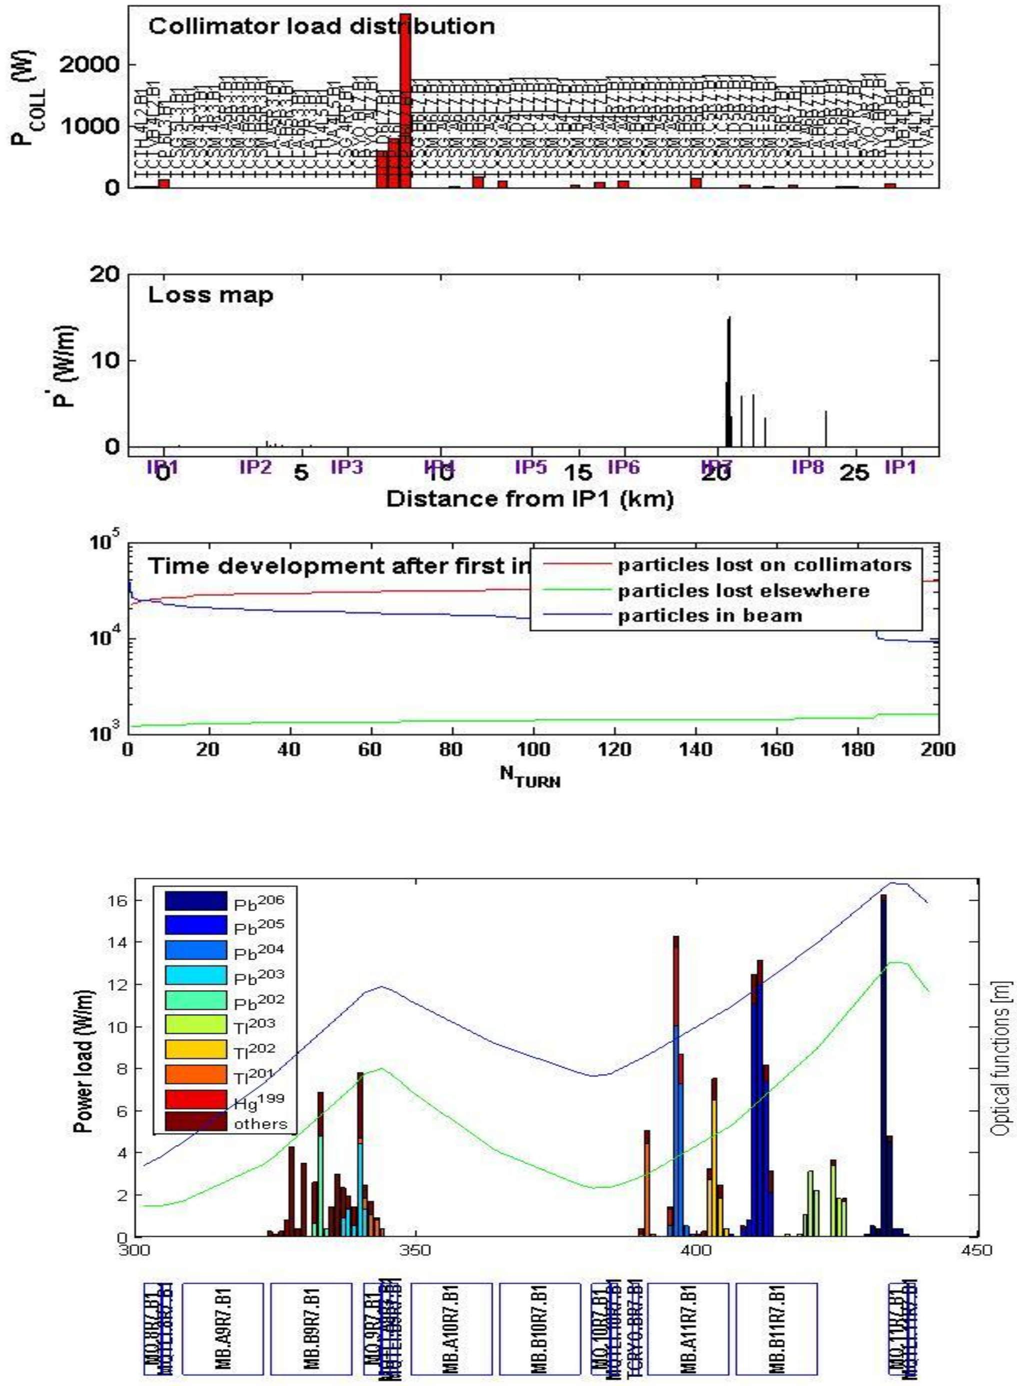

Figure 5: Simulation results for the Phase 2 optics with all cryogenic collimators retracted.

The results are quite similar to the Phase 1 collimator loss maps shown in Figure 1. From this we can conclude that installing only the Phase 2 metallic collimators does little to reduce the losses of ion fragments in the cold magnets around IR7.

In Figure 6, we show results for the identical optics with the difference that the cryogenic collimators are inserted at 45σ:

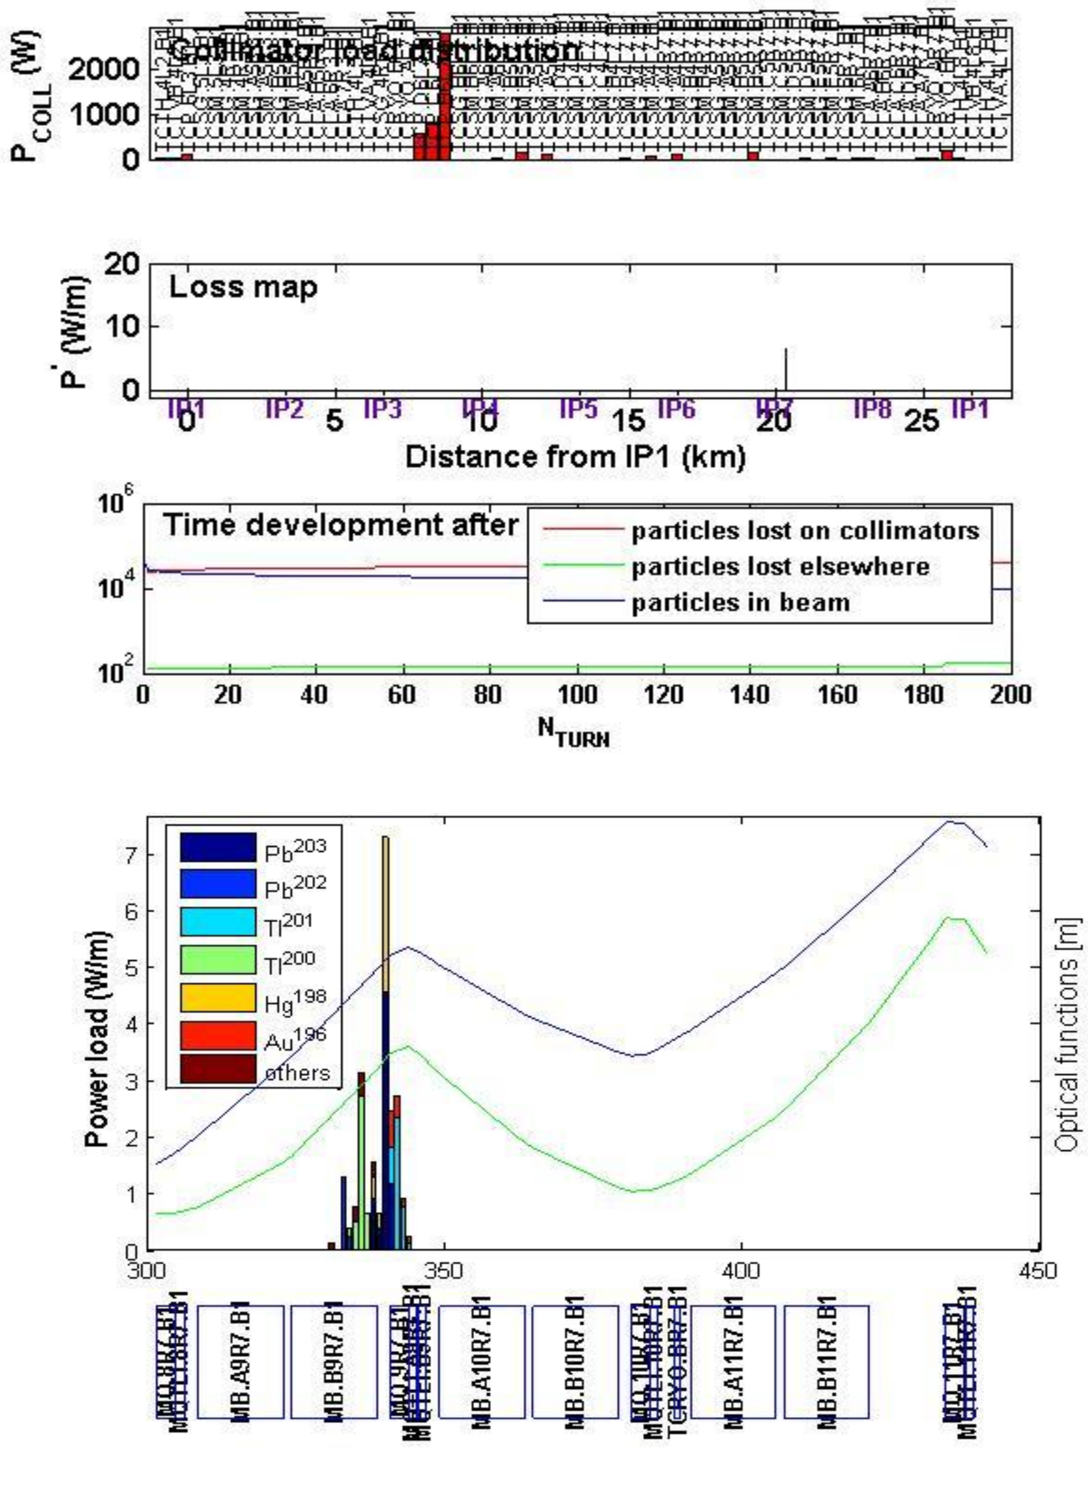

Figure 6: Simulation results for the Phase 2 optics with both cryogenic collimators set at 45σ.

Even with this very wide setting of the cryogenic collimators This is a significant improvement. Going further, it turns out that the losses in the cold section around IR7 vanish completely at around 30σ. Figure 7 shows results for the ideal optics with the cryogenic collimators inserted at 15σ:

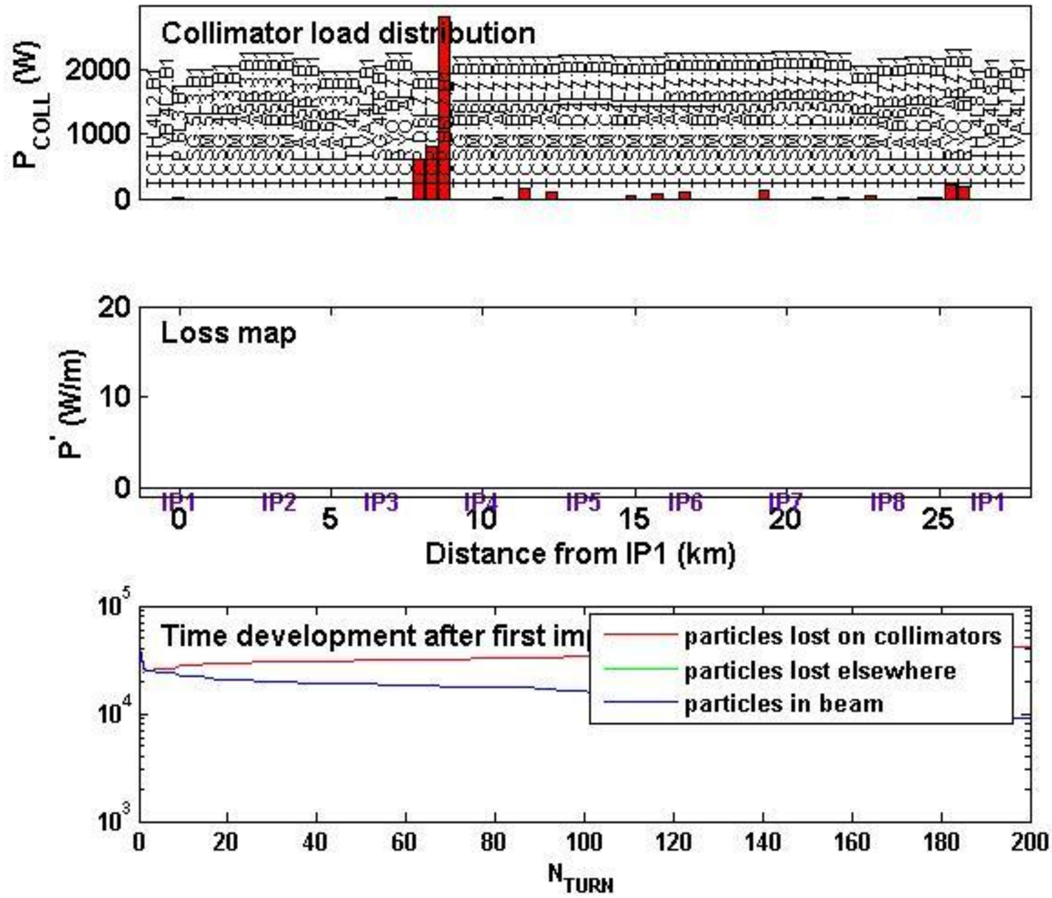

Figure 7: Simulation results for the Phase 2 optics with both cryogenic collimators set at 15σ.

At this level of simulation statistics (150 000 ions in the initial beam), we have the striking result that all losses due to ion fragments scattered out of the primary collimators are eliminated.

Similar results were found in an early mismatched version of the optics, where the beta-beating was of the order of 20%, giving an indication of the effectiveness of this solution in a real, somewhat imperfect, machine. It is clear that the result is robust and essentially independent of the various uncertainties involved in computing the loss levels mentioned earlier.

The cryogenic collimators do alter the load distribution on certain other collimators (e.g., near IP5); more detailed studies are under way to rebalance these but this does not seem to be a significant problem.

While further detailed studies are desirable, we expect these results to be rather robust against collimator tilts and other imperfections of the machine.

4.2. Magnetised Collimators

A magnetised collimator has material jaws outside which is a layer of very strong non-linear magnetic field. The magnetic field should decay rapidly away from the jaw surface so as not to affect the dynamics of the main beam. Errant ions should be deflected by this magnetic field toward secondary collimators rather than interacting with matter. This would avoid their fragmentation and possibly restore the two-stage collimation principle for ions.

Magnetised collimators were discussed for the LHC in 1993 (P.J. Bryant et al, CERN SL/93-15). The risk of unacceptable vibration from cooling their magnetic excitation coils, questions related to their alignment and the difficulty of ensuring that their highly non-linear fields did not have detrimental effects on the main beam seem to have been among the reasons why the idea was not adopted at that time (B. Jeanneret, private communication). They were suggested again more recently by H. Braun as a possible solution for heavy ion beams.

So far no resources have been available to study the possibility of a hardware implementation and it is far from clear whether all the previous concerns about such devices could be overcome.

With the help of a technical student, N. Holden, a model of a magnetic collimator has recently been implemented in ICOSIM and some preliminary simulation results have been obtained. These are described in some more detail in AB-Note-2008-054.

In simulations with the normal set of 3 (horizontal, vertical and skew) primary collimators in the LHC, each positioned at 6σ of the beam distribution, the method of preparation of the initial distribution of ions "about to be lost on a collimator", treats the horizontal and vertical planes equally. This leads to a distribution in which the skew collimator has the highest load and is responsible for creating most of the nuclide fragments that are later lost in the dispersion suppressor. To some extent, this is a consequence of the way ICOSIM works. In reality, loss mechanisms transporting particles to large amplitudes (e.g., a drift of a tune onto a resonance) will usually be dominantly either horizontal or vertical. From this it follows that any of the three primary collimators may, at different times, be the principal source of off-momentum ion fragments so each might have to be substituted by a magnetised collimator. Thus, a suitable installation of magnetised collimators will probably involve at least two, possibly three, devices of different types. These facts also complicate the setup and interpretation of the simulations and a number of runs with specially prepared initial distributions and collimator jaw settings have been necessary to fully understand and interpret the results. A full discussion would be quite lengthy. Here we only show a reference case for the Phase 1 collimation system with normal collimators at their normal settings followed by a case in which the horizontal primary collimator is replaced by the new magnetic collimator model.

Figure 8 shows a reference case with the usual graphite primary collimators set at their standard settings of 6σ of the beam distribution. The initial distribution of particles "about to hit the collimators" prepared in the initial part of the ICOSIM run was saved and re-used in a number of comparison runs with magnetic collimators and varying collimator gap settings. This distribution shows the usual high losses in the dispersion suppressor (peaking at ~ 15 W/m on the right of IP7, reflected in the time evolution of the beam intensity over the 100 turns of the simulation.

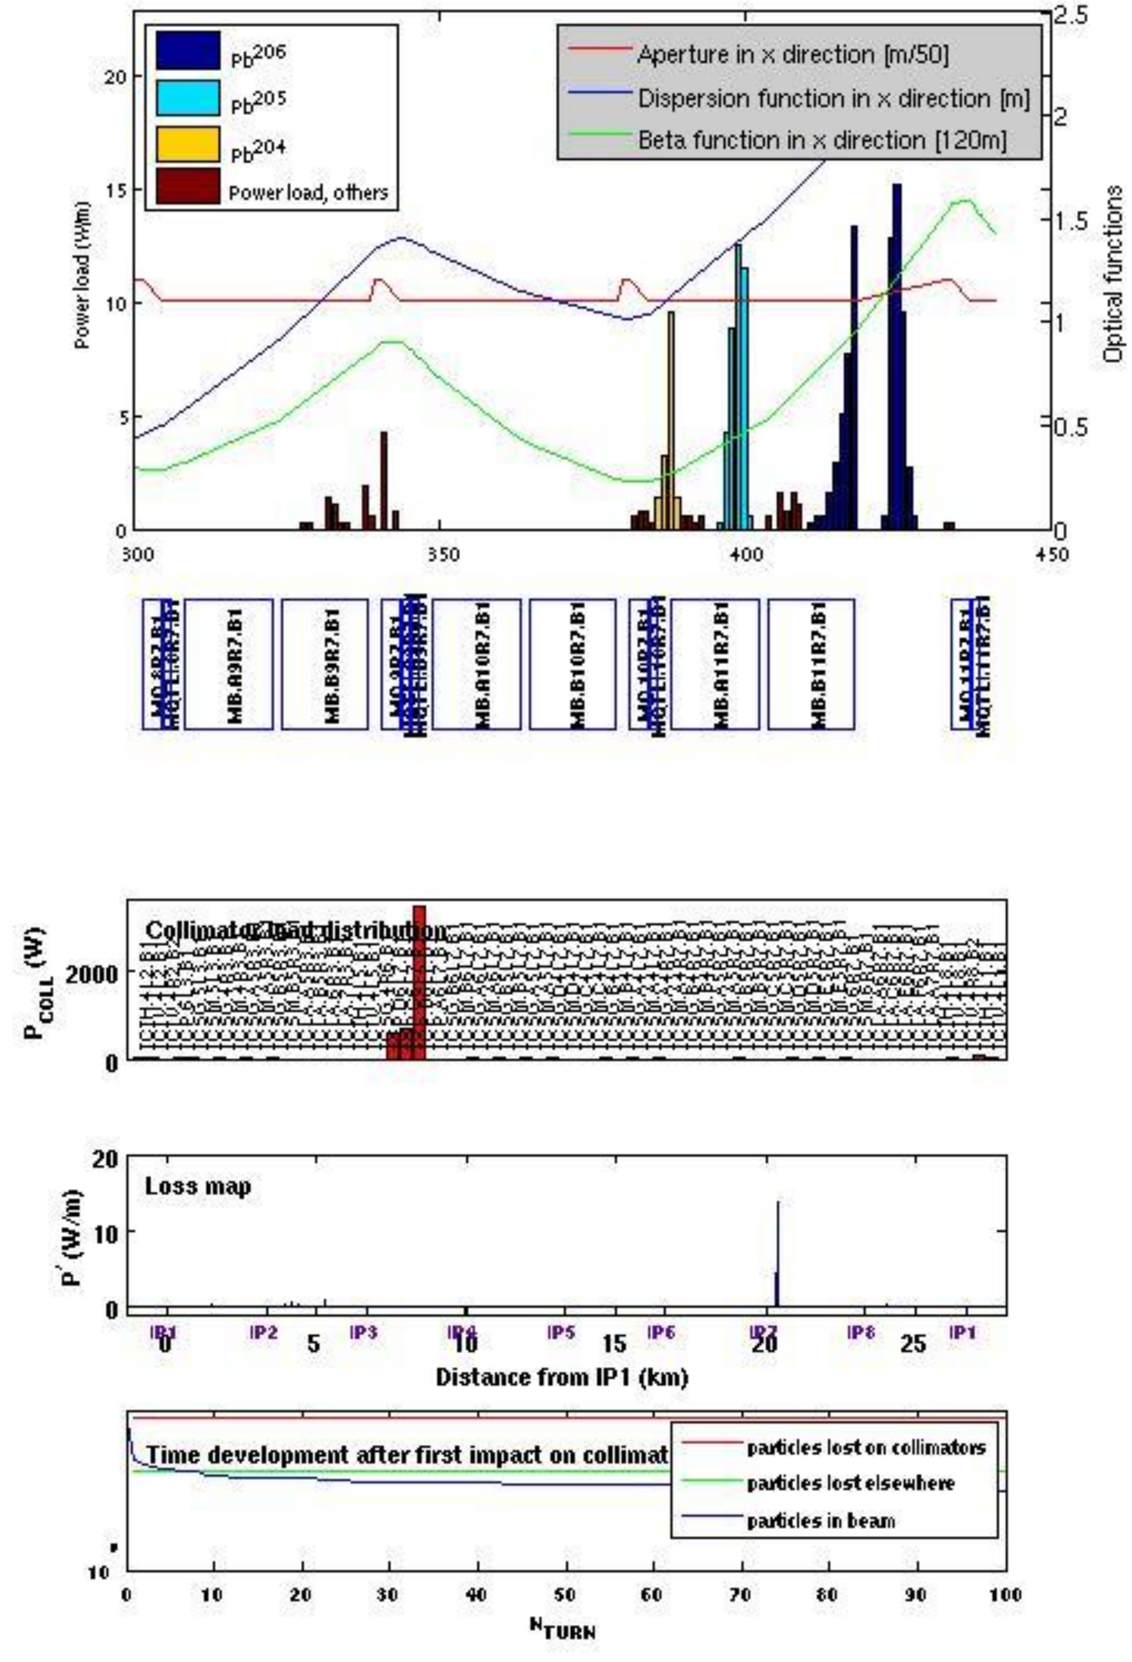

Figure 8: Reference case with the usual graphite primary collimators set at their standard settings of 6σ of the beam distribution.

The initial distribution of ion impact parameters on the horizontal primary collimator is shown in the following plot, where $x = 0$ corresponds to the edge of the collimator material.

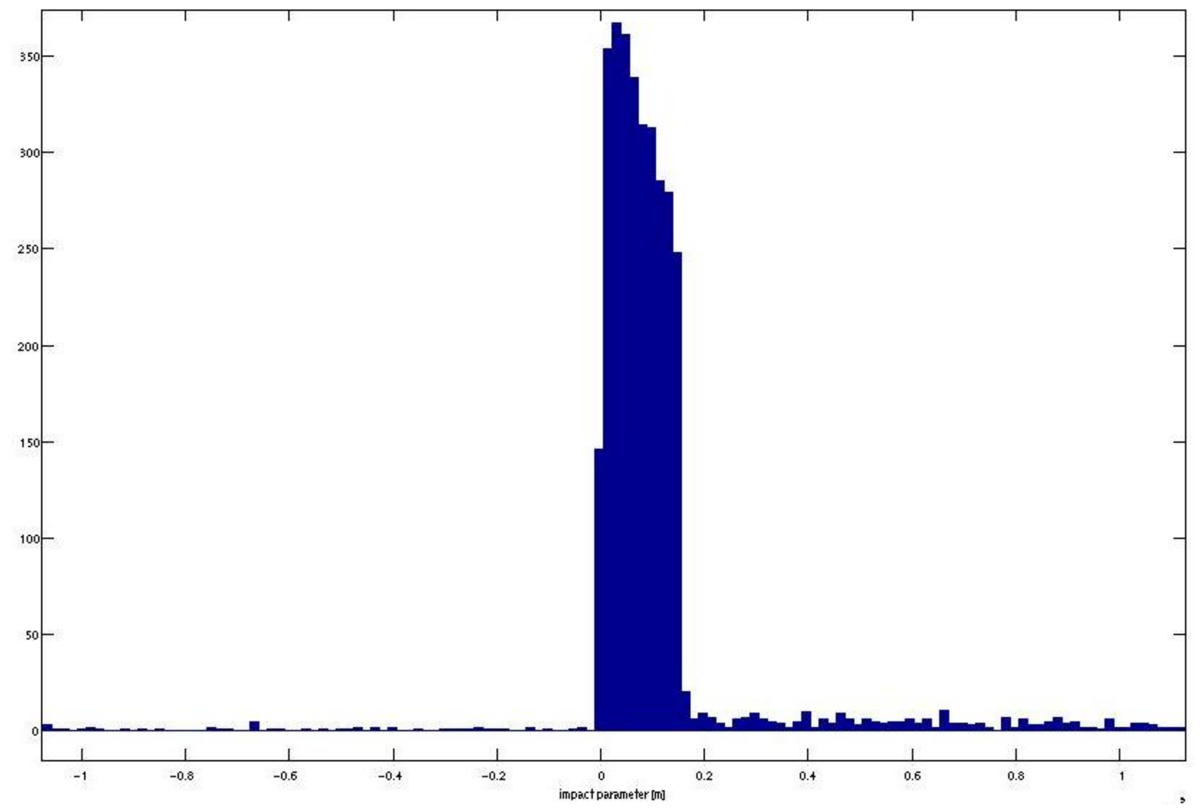

Figure 9: Initial distribution of ion impact parameters on the horizontal primary collimator

In the case of a magnetic collimator, the material part of the collimator jaw is set back by $1.5\ \mu$m from the 6σ position with the field falling very rapidly over a length scale of 1 mm as shown in Figure 10. This means that a substantial fraction of the incoming particles are now affected by the magnetic field rather than hitting the collimator material. The following results, in Figure 11, show a closely related run, with the same initial distribution as Figure 8, except that the horizontal primary collimator is replaced by a magnetic collimator with a peak field of 2 T just outside the material part of the collimator jaw.

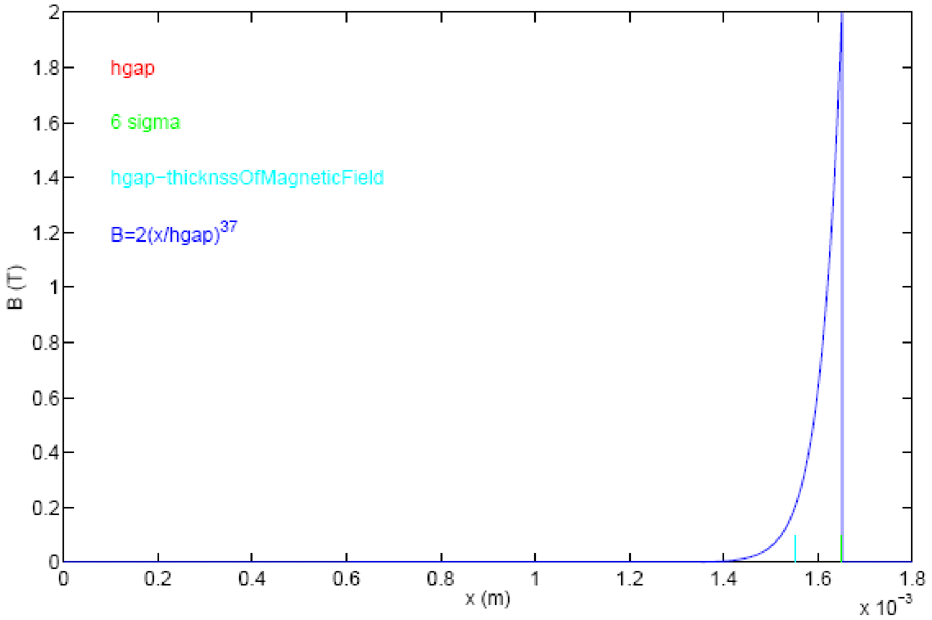

Figure 10: Vertical magnetic field for a model of a horizontal magnetic collimator.

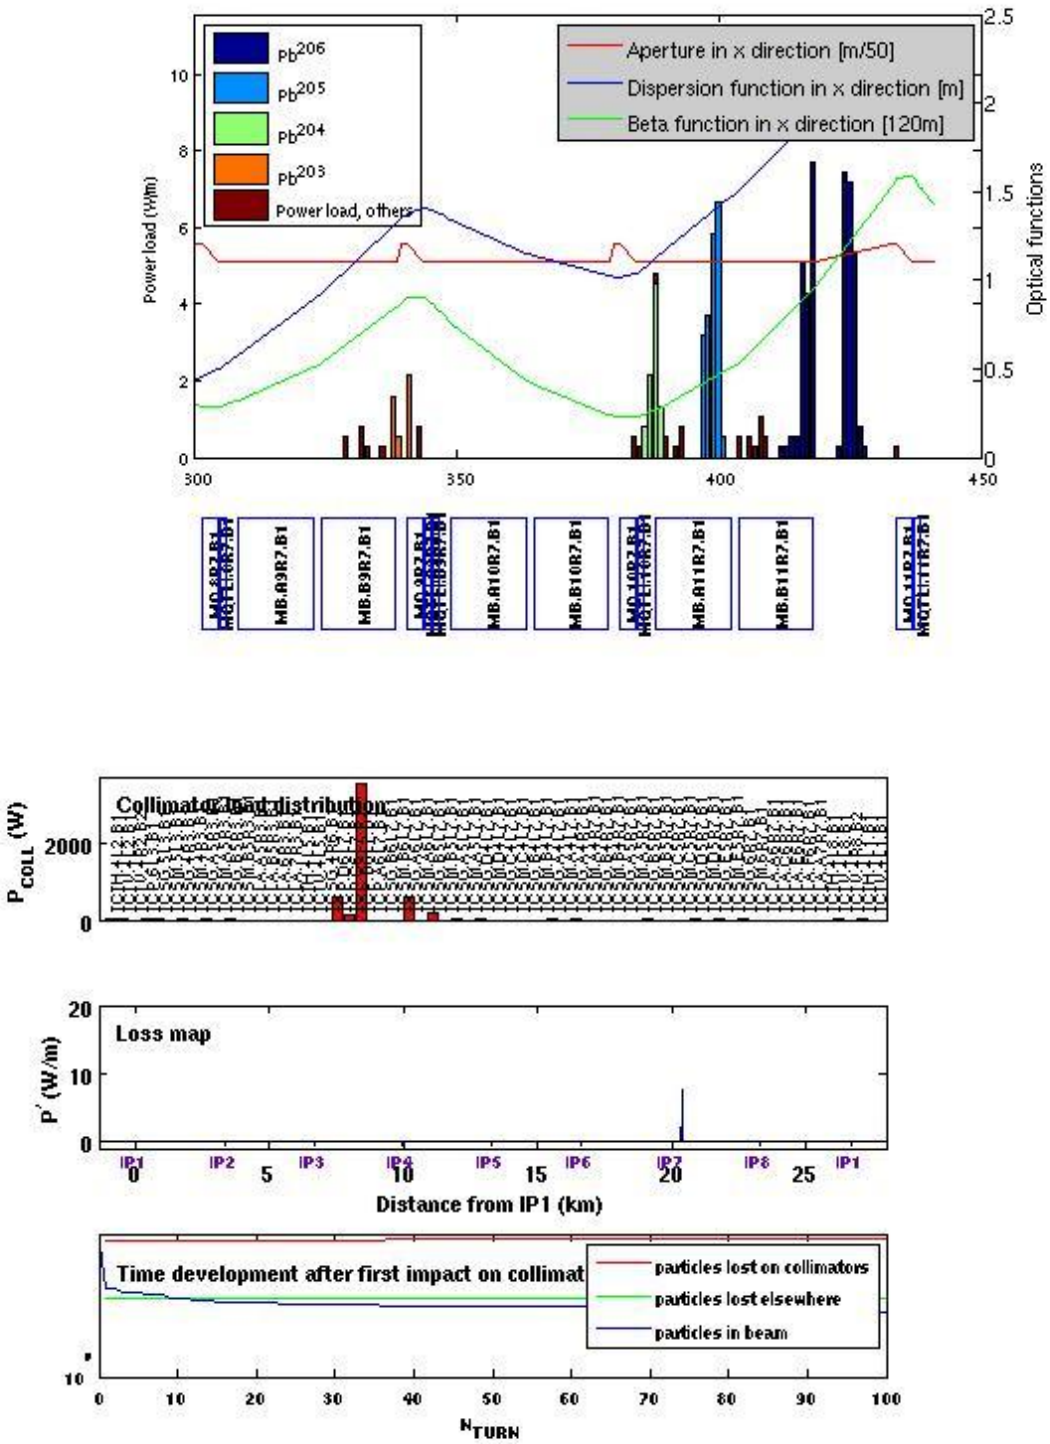

Figure 11: Simulation results for the model magnetic collimator replacing the horizontal primary collimator in the case shown in Figure 8

With these settings, the total losses from the beam are very similar to the case of the normal collimators above as can be seen in the following plots. It is important to arrange this in order to have a valid comparison of the loss maps.

In both cases, the maximum load occurs on the third (skew) primary collimator. However there is a significant reduction in the load on the second (horizontal) collimator when it is replaced by a magnetic collimator. This load reappears further downstream on secondary collimators. At the same time there is a significant reduction in the peak loss in the dispersion suppressor, from ~ 15 W/m to about ~ 8 W/m.

By making fine adjustments to the collimator settings or by withdrawing some collimators completely in the simulation it is possible to re-distribute the main load from the skew to the horizontal or vertical primary collimator. We have made a number of other simulation runs (see AB-Note-2008-054) to study such configurations and achieve a better detailed understanding of the relation of ICOSIM results to what one might expect in reality. Based on these studies, we consider the results selected for display above to be a reasonable indication of what one might expect to see in reality. Our interim conclusions are that one horizontal magnetic collimator installed in each ring with the high peak field of our model could reduce the unwanted losses in cold magnets by factors of order 2. It may be possible to do somewhat better by replacing further primary collimators by magnetic collimators at varying tilt angles around the beam direction.

Extensive further studies would be needed to fully understand the effects of longitudinal collimator tilts, other machine imperfections, etc., most of which are likely to degrade the expected performance of the system.

It has to be emphasised that the magnetic collimator model implemented here is highly idealised and achieves a very strong magnetic field, decaying very rapidly to zero outside a thin layer on the collimator jaw. It is not at all clear whether such a device can be built. Further detailed simulations will be required to reach precise engineering specifications. If design studies then suggest that the hardware can be built, and that they can be installed close enough to, or in place of, the existing primary collimators, prototyping work will have to be carried out with high priority in order to fabricate the collimators in time for installation as part of Phase 2 of the LHC collimation system.

4.3. Optics changes

It has been suggested that the collimation efficiency for ions could be improved by making adjustments to the optics of IR7 in such a way as to cause the fragments with altered magnetic rigidity (effective momentum deviation) generated in the primary collimators to be deviated more strongly towards the secondary collimators. This is equivalent to increasing the R16 element of the linear transfer matrix between the primaries and the secondaries (and not equivalent to changing the periodic dispersion function). In fact there is very little scope to do this.

Analysis of the possibilities shows that, in the normal optics the largest value of this matrix element (a few cm of "locally generated dispersion") occurs at the collimator TCSM.D4L7.B1" which is not far downstream of the primaries. This quantity would have to be increased substantially in order to collimate the nuclide fragments.

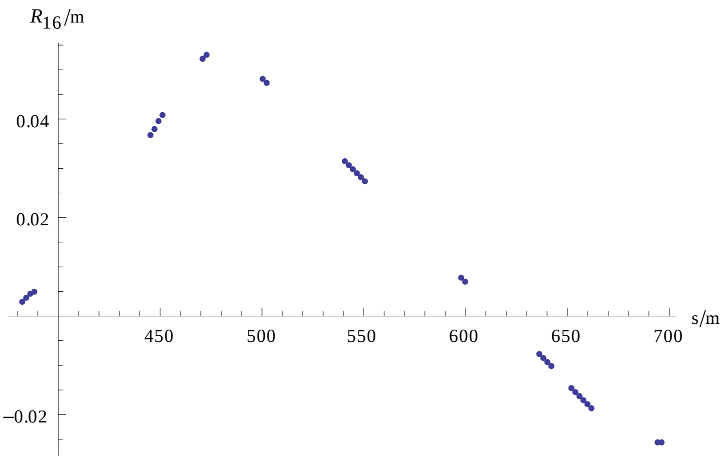

Figure 12 R_{16} matrix element downstream of primary collimators in Beam 1.

The only magnets that could modify this quantity are the set of defocusing quadrupoles which are already close to their maximum strengths. Any increase of the matrix element would require them to go beyond it. Even if this was done, by installing stronger magnets, the optics of the insertion would have to be rematched to the rest of the ring and it is doubtful that this could be done without exceeding aperture constraints.

Table 1 Strengths of quadrupoles between primary and secondary collimators and their limits.

	KEYWORD	s	K1	K1MIN	K1MAX
MQWA.E5L7.B1	QUADRUPOLE	19853.	-0.0013355	0.000080216	0.001499
MQWA.D5L7.B1	QUADRUPOLE	19857.	-0.0013355	0.000080216	0.001499
MQWA.C5L7.B1	QUADRUPOLE	19861.	-0.0013355	0.000080216	0.001499
MQWB.5L7.B1	QUADRUPOLE	19865.	-0.000032658	-0.0012848	0.0012848
MQWA.B5L7.B1	QUADRUPOLE	19869.	-0.0013355	0.000080216	0.001499
MQWA.A5L7.B1	QUADRUPOLE	19872.	-0.0013355	0.000080216	0.001499

4.4. Crystal collimation

Use of the channelling phenomenon in a bent crystal to deflect protons or ion beams has been discussed and studied experimentally for many years. Interest in the use of channelling to deflect ions was stimulated by the theoretical expectation that nuclear fragmentation and electromagnetic dissociation (EMD) of the ions would be suppressed. More recently the volume reflection phenomenon observed for protons (by W. Scandale et al) promised much greater efficiency than channelling. However consideration of the physical processes involved shows that the suppression of EMD expected in channelling is not expected in volume reflection and this mode of deflection by a bent crystal will not be useful for ions.

Earlier studies of crystal collimation with ion beams in RHIC were inconclusive.

A test with ion beams in the H8 beam line at the SPS was planned for 2007. Unfortunately this was cancelled due to loss of beam time for ion studies in 2007 and delays connected with the RF. There have been no ion beams in the CERN complex in 2008. It may finally become possible to perform this experiment in 2009 once ion beams have been fully commissioned in the SPS. An even more interesting experiment may be possible with circulating beams and a crystal collimator installed in the SPS ring itself. However, it is not yet clear how much time will be available for these experiments in the 2009 schedule. They will have to be well prepared and ready to take data when the opportunity arises. It is difficult to predict how conclusive the results will be and whether one will then be in a position to install reliably operating devices as part of Phase 2 of the LHC collimation system.

Despite the success of some recent experiments with protons, crystal collimation, remains very much an R&D topic.

4.5. Electron-lens as collimator

For completeness, it should also be mentioned that the use of electron lenses (as implemented in the Tevatron) has been proposed by V. Shiltsev as another possible solution for ion collimation. The idea is conceptually similar to the magnetic collimator: a strongly nonlinear field is used as primary collimator to deflect errant beam ions towards the secondary collimators without giving them the chance to fragment by interaction with a material collimator. In this scheme the nonlinear field would be provided by a hollow electron beam.

This option has not been studied at CERN and would require a substantial hardware development and installation.

5. Future Phases of the LHC Ion Programme

Future operational modes of the LHC are expected to include hybrid p-Pb collisions and collisions of lighter ions. For p-Pb mode, the collimation requirements in the proton ring will be essentially the same as in p-p operation (except that the intensity will likely be less) and those in the Pb ring will be similar to those in Pb-Pb operation. Therefore no new collimation requirements are anticipated for this mode.

Up till now, no resources have been devoted to the study of lighter ion beams (e.g., Ar, Ca) in the LHC. However these beams should be of higher intensity than the Pb beams and the magnetic rigidity spectrum of the ion fragments will be broader (although cross-sections will be smaller). Without some detailed study, it is not easy to judge how the loss maps for collimation inefficiency will differ from those for p-p or Pb-Pb operation. Since the Phase 2 collimation system will be expected to cover the requirements of all future operating modes of the LHC, it is important to ensure that it can also cover this case.

The tools we have prepared can be applied to generate fragmentation and electromagnetic dissociation cross-sections for these cases and then to simulate loss maps. A first study has been started for the case of $^{40}Ar^{15+}$ nuclei ($E = 3.15 A\,\mathrm{TeV}$), and preliminary results indicate that, for the same number of nucleons per beam (the intensities are expected to scale roughly like this), the collimation inefficiency problem is just as important as for Pb nuclei. Indeed the peak power in the loss maps scales approximately as the total energy in the beam, i.e., for the same bending field and nucleon-nucleon luminosity, implying the same number of nucleons in the beam, the peak power in the loss map is increased roughly by the ratio of the energy per nucleon. The loss map patterns in IR7 are modified in ways that reflect the broader spectrum of magnetic rigidity changes created by the fragmentation processes. Initial simulations also confirm the efficacity of the cryogenic collimator scheme for these beams.

6. Conclusions

The solution of implementing cryogenic collimators in the dispersion suppressors shows the most promise as a relatively straightforward, direct and robust solution to the problem of collimation inefficiency for heavy ion beams in the LHC. It is also the least dependent on results of future beam tests whose timing and results are difficult to predict. Its interest will be all the greater if, as appears likely, it is also the preferred solution for proton collimation.

Similar installations (with probably only one cryogenic collimator per dispersion suppressor) around the experiments taking ion collisions could also intercept the secondary beam of Pb81+ ions created by bound-free pair production in Pb-Pb collisions, potentially also removing the most important performance limitation for heavy ion operation of the LHC. Studies to verify the effectiveness of such a scheme and clarify the hardware and optical modifications to the experimental IRs will be started (by myself and Roderik Bruce in Lund) as soon as possible. We will also have to check whether the experiments can accept the transverse displacements of their interaction points.

It is therefore important to advance the development of the necessary hardware and embark on studies of the integration of cryogenic collimators into the dispersion suppressors of the LHC. These will have to include FLUKA studies of the cryogenic collimator design to verify its capacity to absorb the loss load and any secondary losses to the cold magnets.

Recent progress with the study of magnetized collimators suggests that, in carefully established conditions, they may also provide a useful improvement in collimation efficiency. However it remains to be seen if the hardware can be developed and built in time.