Linac power optimisation including beam-velocity effects

Stephen Molloy

January 25, 2012

1 Introduction

Many calculations to determine the required RF power for a linac neglect the fact that velocity changes in the beam will alter the strength of the coupling between it and the accelerating field. This is quite natural in electron machines, where the beam may be assumed to be moving with the speed of light throughout the entire machine.

In proton linacs this is typically not the case, and the changing magnitude of the coupling must be taken into account.

This is demonstrated here by first performing an optimisation of the ESS linac assuming that the coupling parameter, R/Q, is not dependent on the beam velocity, and then a second time where the true dependence is taken into account.

2 Optimisation

Given an entirely free choice as to the amount of detuning to apply to a cavity, and the design of the input coupler (as specified by the resulting loaded quality factor, Q_L), it is possible to achieve a solution that results in complete absorption of the input power by the beam, and zero reflected power.

In practise, this would require the use of a specially designed coupler for every cavity, and would be prohibitively expensive. Thus the goal is not to design a linac with zero reflected power, but rather to try to design one where the coupler for each linac section is designed so as to minimise the total reflected from that section.

Thus a function has been written that accepts as input the details of the desired accelerating voltages in the cavity, as well as the expected beam velocities, and the details of the various cavities, and returns the output the total power reflected from each section.

This is then wrapped into an optimisation routine that is free to vary the Q_L of each coupler in an attempt to minimise the total reflected power.

3 Neglecting the velocity dependence

Figure 1 shows the power distribution in the linac for the case where the nominal values of Q_L are used in each section, and where the velocity dependence of the

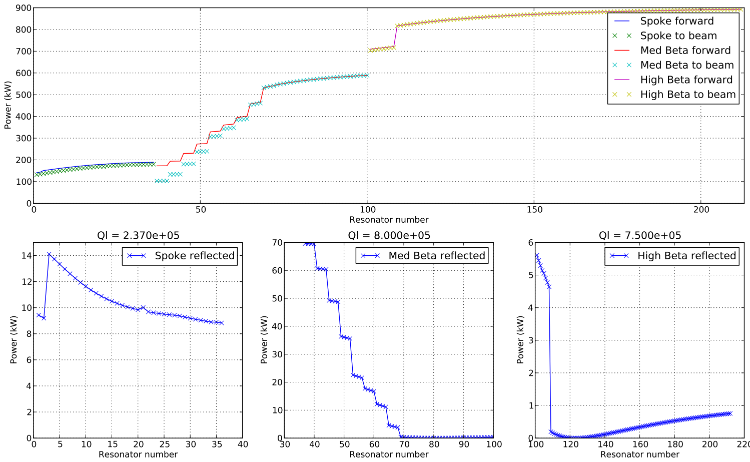

Figure 1: RF power for the unoptimised linac where the velocity dependance of the R/Q is ignored.

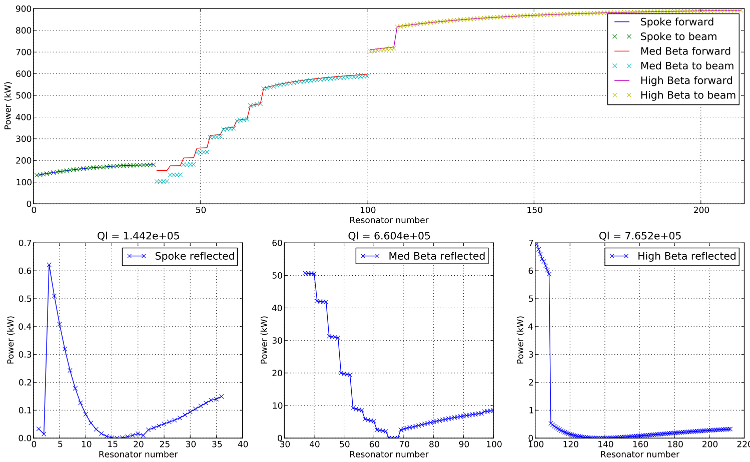

Figure 2: RF power for the optimised linac where the velocity dependance of the R/Q is ignored.

R/Q has been neglected.

It can be seen that the reflected power in the spoke section is consistently high, and that it should be a simple matter to optimise the Q_L here to significantly reduce the total reflected power.

On the other hand, the high beta section appears to be quite close to optimal since any improvement of the low-energy region of this section of the linac will be compensated for by an increase in the total reflected power from the high-energy region.

The medium-beta section appears particularly problematic due to the large spread in the power transferred to the beam. Optimisation of this section is expected to be quite difficult.

Figure 2 shows the optimised version of this linac, again ignoring the velocity dependence of R/Q.

As expected the spoke section shows considerable improvement with a significant change in its Q_L.

The optimised Q_L for the medium & high beta sections do not show such a large difference from the nominal values, and the large spread of accelerating voltages in the medium beta section result in quite a large amount of total reflected power.

In performing this optimisation, the total reflected power for the entire linac was reduced from 1536 kW to 894 kW, representing a saving of 643 kW (peak power).

4 Including the velocity dependence

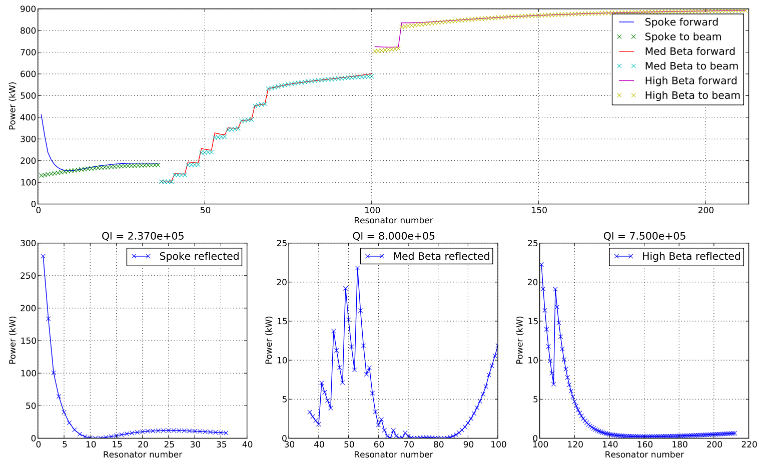

Figure 3: RF power for the unoptimised linac where the velocity dependance of the R/Q is included.

Figures 3 and 4 show the situation when the velocity dependence of the R/Q is included, and it can be seen that this relationship leads to a very large difference in the profile of the reflected power.

While the high beta section can be seen to experience a larger amount of reflected power, it is still relatively low, and of no major importance to the design of the machine.

Some differences can also be seen in the medium beta section, with a distinct improvement in the maximum reflected power in this part of the machine.

The major difference between these plots and those in figures 1 and 2 can be seen in the spoke cavity section. Here the low energy cavities reflect up to

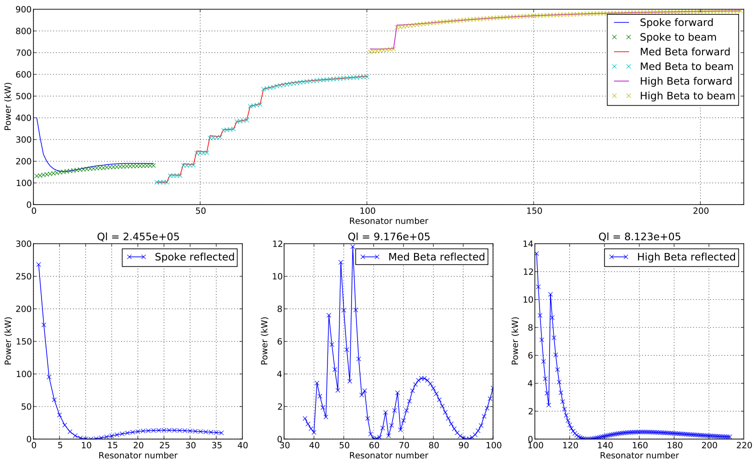

Figure 4: RF power for the optimised linac where the velocity dependance of the R/Q is included.

300 kW, which is almost three orders of magnitude higher than calculated when ignoring velocity effects.

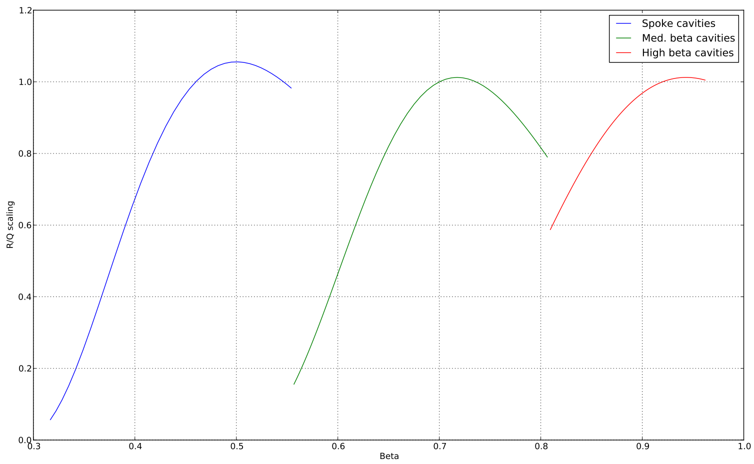

Figure 5: Scaling of the R/Q for each cavity in the linac.

Figure 5 shows the R/Q scaling for each cavity in the linac, and some hint as to the reason for the large reflected power from the beginning of the linac can be seen in the fact that the R/Q falls to a very low level in this region.

The R/Q scaling cannot be the whole story however, since the scaling for the medium beta section falls to a similarly low level, but without the very large

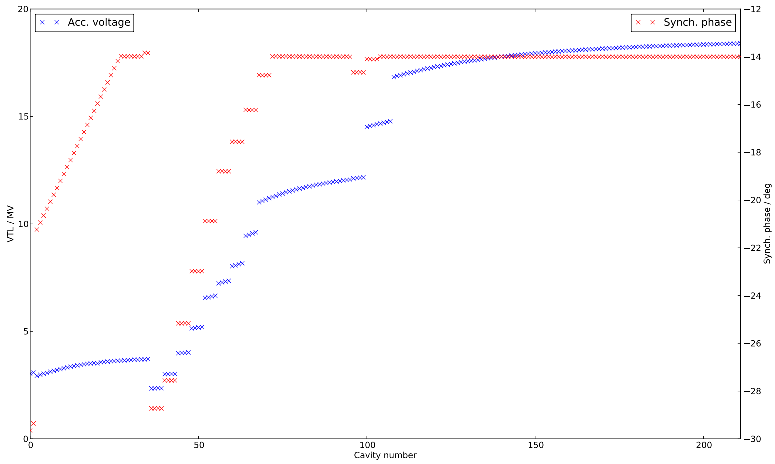

Figure 6: Voltage & phase profile for the linac.

increase in reflected power. By examing the voltage and phase profiles required by the linac designers in figure 6, an explanation can be found.

In the case of the cavities in the medium beta section, the low R/Q is compensated for by the fact that the required cavity voltages have been set by the linac designers in a way that is approximately proportional to the R/Q scaling of the cavity. This means that the optimal detuning angle of the cavity is approximately equal to the desired synchronous angle, and it can be shown that this solution reduces the reflected power to zero.

The voltage profile for the spoke cavities does not approximate the R/Q scaling curve in this way, resulting in the very large reflected powers observed in figure 4. In effect, the most optimal solution for these cavities requires a Q_L that results in a beam injection time significantly after the point at which the transient reflected power has passed through zero, just as in the case of accelerating a bunch current much lower than the nominal value.

A similar effect can be seen at the low energy end of the high beta cavities, although with a magnitude much less than that observed in the spoke section. Looking at the reflected power for the high beta section in figure 4, a steep incline can be seen for the cavities numbered 110–120. For the high beta cavities lower than number 110, a similar incline can be seen, but with a discontinuity that corresponds to the jump in required cavity voltage moving from cavity 108 to cavity 109. The fact that this behaviour is so similar to that at the low energy end of the spoke section, as well as the fact that reducing the cavity voltage for the first few high beta cavities results in a drop in the reflected power, confirms the hypothesis that the large reflected powers observed in the spoke cavities are related to a mismatch between the desired voltage and the R/Q scaling of the low energy cavities.

5 Conclusions

Given the differences between figures 1/2 and figures 3/4, as well as the Q_L output of the optimisations, it is very clear that the velocity dependent scaling of the R/Q is vital to the result obtained.

The optimisation shown in figure 4 showed that suitable Q_L values can be obtained for the medium and high beta sections, however the spoke cavity presents a significant problem due to the resulting reflected power.

As shown in figure 5, the R/Q scaling of the spoke cavities falls very low in this region, however the fact that similar low levels are also observed in the medium beta section implies that this cannot be the principle cause of the problem.

It was concluded that the low R/Q scaling in the medium beta section was masked by a corresponding drop in the required cavity voltage, and that the significant differences between the voltage profile and the R/Q profile in the spoke section cause a significant amount of reflected power, even for a well optimised cavity.

The large reflected powers observed in the spoke cavity section are certainly undesirable, and there are several alternatives that may be considered to improve the situation.

Eliminate low energy end of the spoke section Much of this problem could be avoided by increasing the beam energy at the start of the spoke section by extending the Drift Tube Linac (DTL) to accelerate the beam to 65 MeV rather than 50 MeV. This would eliminate the need for the first eight spoke cavities, and allow the Q_L optimisation to converge on a solution that would result in reflected powers with a magnitude similar to that observed in the medium & high beta sections.

Decrease the design velocity of the spoke cavities The problem is strongly related to the fact that the velocity of the first few cavities is so much less than that for which the cavities were designed, and so the size of the reflected power would be reduced by reducing the design velocity of the cavities to a value that is closer to the average value of the beam in this section.

Figure 5 shows that the current value is strongly skewed towards the high energy end, and that a reduction in this value could be expected to have a significant benefit for the reflected power.

Additional coupler designs in the spoke section Having a Q_L better matched to the expected beam conditions could also help to reduce the reflected power, by allowing filling of the cavity at a rate that matches the rate at which the beam is taking away energy.

Given the steepness of the slope in figure 5, the low energy region might require two additional coupler designs in order to reduce the reflected power to an acceptable level.

Each of these three solutions could be used to reduce the reflected power levels to a reasonable level, however additional studies would be needed to determine the other effects these would have on the performance of the machine.