Splitting the Medium Beta Section

Stephen Molloy

January 4, 2012

1 Introduction

The input coupler to a superconducting (SC) accelerating cavity may be designed to result in zero reflected power from the cavity during acceleration of a beam [1]. For a particular coupler design, this can only be achieved for a particular value of the accelerating field, sychronous phase, beam current, and beam velocity, and so in cases where a significant spread of velocities will be accelerated by a particular cavity design, or where a large spread of accelerating voltages or synchronous phases are required by the lattice designers, a considerable amount of RF power will be reflected from the cavities.

In this case the task is not to optimise the coupler for a specific set of beam & cavity parameters, but rather to optimise the coupler so as to minimise the total reflected energy for the section in which that coupler is installed.

2 Optimising each section

ESS will use three sections of SC cavities – spoke cavities, medium β elliptical, and high β elliptical – and the required accelerating voltages and synchronous phases have been specified by the lattice designers.

For each of these three families, we define the optimum coupler design as that which results in the minimum total reflected power for that section.

Given a spreadsheet containing each of the relevant parameters for all the cavities in a particular section, it is possible to write a function that returns the total reflected power for a specific value of the loaded quality factor, Q_L, resulting from the coupler design. This function may then be supplied to an optimisation routine that determines the value for Q_L that results in the lowest total reflected power.

Figure 1 shows the forward and reflected powers in the situation where the nominal values of Q_L are used for each of the cavity families ($Q_L = 2.37 \times 10^5$, 8×10^5, 7.5×10^5 for the spoke, medium β, and high β cavities respectively). In this case, a total of 1976.5 kW (peak) is reflected from the SC linac.

Figure 2 shows the situation where the Q_L values are optimised as described above. This optimisation results in the total reflected power dropping to 1242.5 kW, which is clearly a significant saving in the amount of power required. Notice that the performance of the spoke cavities in particular is significantly improved by the optimisation, with a 40% change in the value of the coupling.

The peak value of the high β cavity has risen slightly, but the overall integral is obviously significantly improved by this process.

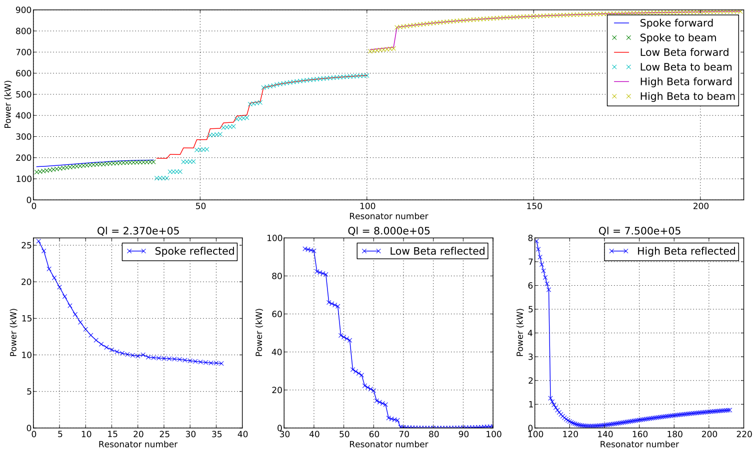

Figure 1: Forwarded and reflected power for the nominal coupler designs.

Also improved is the medium β section, although some significant problems can still be seen here. The optimisation has dropped the peak reflected power from this section from $\sim$90 kW to less than 70 kW, with a small rise in the power reflected from the higher energy resonators.

The peak power at the lower energy cavities (70 kW) is still very high, and it can be seen from the upper plot that this is related to the very low power to the beam required in this section.

The main issue is that the boundary between spoke and low β cavities is also the point at which the frequency of the accelerating RF jumps from 352.21 MHz to 704.42 MHz. From the point of view of the longitudinal beam dynamics, the size of the stable region of phase space would be shrunk by a factor of two if no changes were made to the accelerating voltage and synchronous phase. For this reason, the lattice designers double the synchronous phase to -28^o, and drop the accelerating voltage to a very low level, before gradually moving to the nominal values throughout the next 30 cavities.

Thus, the spread of required voltages and synchronous phases is very large for this family of cavities, resulting in it being very hard to find a suitable optimum.

3 Splitting the coupling

The requirement for the medium β section to include a large number of cavities whose purpose is to match the longitudinal phase space over the frequency transition, suggests that it might be better to view the medium β section as several families of cavities. In other words, it is worth investigating the possible improvements achieved by allowing several families of couplers within this section.

It was decided to investigate the optimisations possible by using three different values of the Q_L throughout the medium β section.

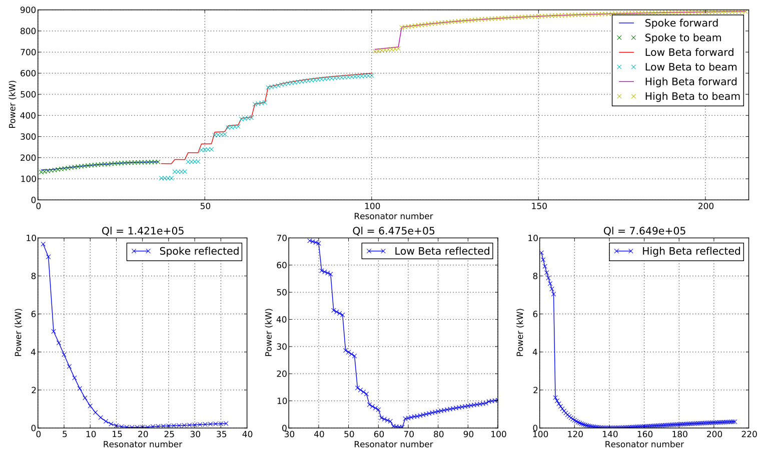

Figure 2: Forwarded and reflected power for optimised couplers.

This calculation was done by wrapping the previous optimisation (i.e. that of setting the Q_L of a section of cavities to minimise the total reflected power from that section) within another based on the break-points between the different coupler families.

In other words, the number of cavities in each of the three families was left as a free parameter, and the optimiser would vary these in order to find the combination that resulted in the lowest reflected power. At each stage of this optimisation, another optimising routine would be called in order to calculate the best Q_L for each of the sections.

The results can be seen in figure 3, where the red lines mark the breaks between the new Q_L families, and the results for the spoke and high β sections are identical to those in figure 2.

It can be seen that the optimisation has resulted in the break-points between coupler families both lying in the region where the longitudinal matching is being performed, and where, therefore, the voltages and phases are furthest from the primary operating point of the cavity.

It can be seen that this optimisation offers a vast improvement for the performace of the medium β section, with the peak reflected power dropping from 70 kW to 11 kW, and with the reflections from each of the three sections being considerably lower than previously. For example, each of the cavities in the final section reflect only a few hundred Watts (except for the first four), showing that this optimisation allows this range of cavities to be tuned to almost perfectly match the design beam parameters.

After this optimisation step, the total reflected power is 334.2 kW, representing a saving of over 1.6 MW from the nominal case, and ~900 kW from the initial optimisation. Note that these numbers correspond to the peak powers, and the average power saving from the nominal parameters is 64.4 kW (due to the duty cycle of 3.9%.

It is also worth noting the optimal values calculated for the Q_L of each of the

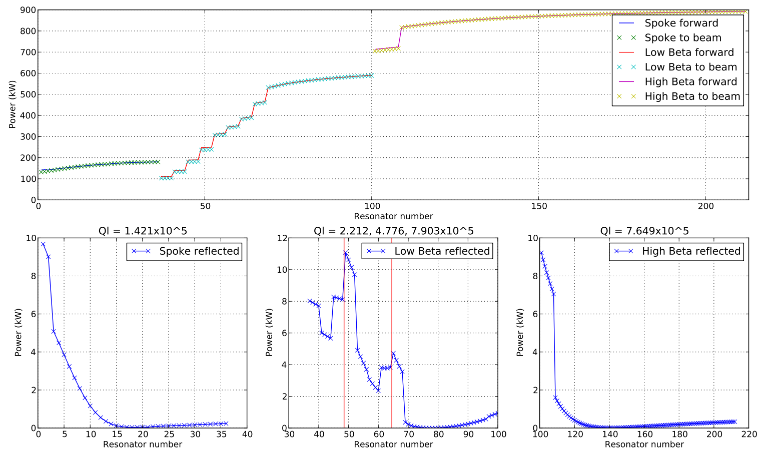

Figure 3: Forwarded and reflected power for the situation where three different coupler designs are used in the medium beta section.

three new medium β families – 2.212×10^5, 4.776×10^5, and 7.903×10^5. The fact that each of these values are significantly different from each other supports the decision to split this section into three coupler families, rather than two. If it was found that one of the optimal Q_L values lay quite close to another (within, say, 10%), it could be argued that those values were sufficiently close to allow merging of the families. The significant differences observed here, therefore, do not support an argument to move to two Q_L families.

4 Conclusions

It was found that the nominal values of the load quality factors due to the input couplers resulted in a total reflected power of 1976.5 kW (peak) over the ESS SC linac, and that by optimising each section this could be reduced to 1242.5 kW (peak).

The main feature hampering the performance of this optimisation was the need to match the longitudinal optics across the frequency jump between the spoke and medium β sections, since this resulted in a large number of cavities with a very low accelerating voltage and large synchronous phase angle.

It was shown that it is possible to improve the optimisation if the medium β section may be allowed to contain three different families of couplers. An optimisation based on this resulted in a total reflected power of 334.2 kW (peak) – a 1.6 MW saving on the nominal case.

While it should be expected that the need to design two additional couplers will increase the expense of the machine, it is a simple matter to tune a coupler to a particular Q_L, and so the total cost increase should be significantly smaller than the savings due to the reduced power usage over the lifetime of the machine.

References

[1] S. Molloy, "Required power for detuned, loaded, SC cavity", Tech Note ESS/AD-158, 2011