S03120100-TRT10010
ORNL/TM-2023/3100

Oak Ridge National Laboratory Optimization of the moderators in the STS preliminary design

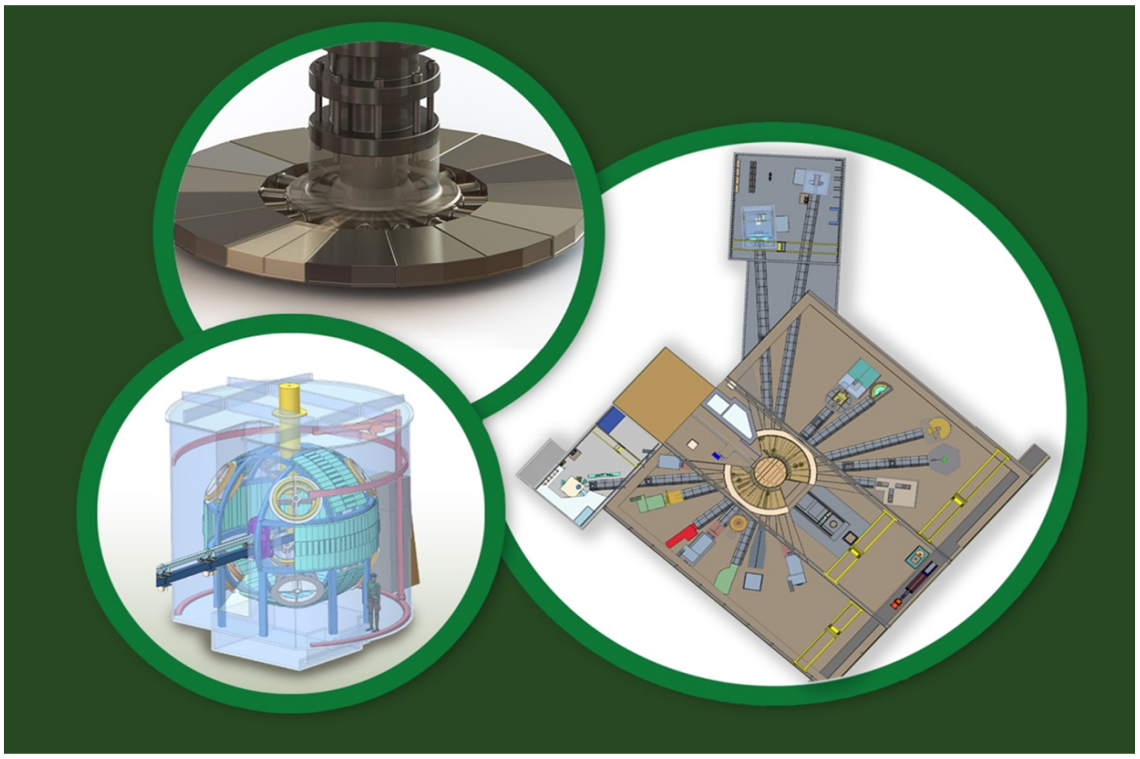

Kristel Ghoos
Lukas Zavorka
Joel Risner
Igor Remec

June 2022

OAK RIDGE National Laboratory

ORNL IS MANAGED BY UT-BATTELLE LLC FOR THE US DEPARTMENT OF ENERGY

S03120100-TRT10010
ORNL/TM-2023/3100

Second Target Station Project

Optimization of the moderators in the STS preliminary design

Kristel Ghoos
Lukas Zavorka
Joel Risner
Igor Remec

June 2022

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-Battelle LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

Moderator Optimization of the STS preliminary design

Optimization of the moderators in the STS preliminary design

LABORATORY ORNL	DIVISION/GROUP Second Target Station (STS) Project	CALC NO. S03120100-TRT10010
Prepared by Kristel Ghoos	Level III Manager Igor Remec	Lead Engineer Jim Janney
Other WBS elements affected S031201		

Signature/Date

	REV 0	REV 1	REV 2	REV 3
Prepared By				
Task Leader				
Level III Manager				
Checked By				
Lead Engineer				

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABBREVIATIONS

CSG	Constructive Solid Geometry
EGO	Efficient Global Optimization
EIF	Expected Improvement Function
ORNL	Oak Ridge National Laboratory
SNS	Spallation Neutron Source
STS	Second Target Station

EXECUTIVE SUMMARY

This report details the results for an optimization of the dimensions of the moderators in the preliminary design of the Spallation Neutron Source Second Target Station (STS). This study uses the optimization algorithms of Dakota[1] and an unstructured mesh model for the moderators in MCNP[2]. More details on the unstructured mesh model and the automated mesh generation can be found in [3]. Parallel to this effort, the same moderator geometries have been optimized using a constructive solid geometry (CSG) MCNP model. More details on this model and its results can be found in [4]. Three optimal designs are selected for each moderator: one that is optimized for maximum peak brightness, one for maximum time-integrated brightness, and one for a combination of peak and time-integrated brightness.

The backbone of the optimization work flow is provided by Dakota. For each set of design parameters requested by Dakota, a new solid geometry is automatically built in Creo[5] and SpaceClaim[6], and subsequently exported to Attila4MC[7] to generate an unstructured mesh geometry for MCNP. After the MCNP calculation is finished, the objective function (e.g., brightness metric) is returned to Dakota. After the new design has been evaluated, a result-file is written, and Dakota proposes the next set of design parameters to be evaluated. The loop continues until a specified convergence criterion has been met.

The design parameters of the cylindrical (upper) moderator include the hydrogen radius, the premoderator thickness (top, bottom, radial), the beryllium radius and the horizontal position of the moderator. The crucial design choice is the hydrogen radius. A radius of 62 mm is shown to provide the maximum time-integrated brightness. The maximum peak brightness occurs with a radius of 40 mm. A combined (middle) design, which balances peak and time-integrated brightnesses, is obtained with a hydrogen radius of 50 mm. The premoderator thicknesses and the beryllium radius are slightly larger in the design optimized for time-integrated brightness than in the design optimized for peak brightness. The sensitivity to these two parameters is relatively small close to the optimal configurations.

The hydrogen vessel and vacuum vessel wall thicknesses are dependent on the radius of the liquid hydrogen due to structural integrity requirements. The increased wall thicknesses for larger vessels significantly penalize the time-integrated brightness, with the maximum obtainable value reduced by more than 10% relative to earlier studies which used fixed vessel wall thicknesses. The impact of the variable wall thicknesses is much less for the peak brigthness and combined brightness designs.

The design parameters of the tube (lower) moderator selected for the optimization are the tube length, the annular premoderator thickness, the beryllium radius and the horizontal position of the moderator. The tube length is the crucial parameter and is chosen large (210 mm) and small (125 mm) in the designs optimized for time-integrated and peak brightness respectively. A combined optimal design has a tube length of 170 mm. The premoderator thickness and the beryllium radius are chosen larger in the design optimized for time-integrated brightness.

1 SCOPE

This study covers the results of the optimization analyses of both moderators for the preliminary design of STS. This report details

- the algorithms that have been used in Dakota,
- the results and interpretation of the optimization and sensitivity analysis using the unstructured mesh models for the moderators (more details on the models itself can be found in [3]),
- a comparison with results obtained with the CSG model (more detailed information can be found in [4]).

2 ACCEPTANCE CRITERIA

Not applicable.

3 METHODOLOGY AND MODELS

3.1 DAKOTA

Dakota is a software package containing many algorithms for optimization, parameter estimation, uncertainty quantification, and sensitivity analysis. It also provides a flexible interface to external simulation codes.

3.1.1 Algorithms of interest

Dakota offers a number of optimization algorithms. The efficient global optimization (EGO) algorithm, which is a gradient-free global optimization algorithm, was used for this study. It typically converges faster than alternative gradient-free algorithms, because it uses an internal surrogate model to predict the most interesting regions in the parameter space. Details can be found in the Dakota Theory Manual [8]. The general procedure of the algorithm is as follows:

- Build surrogate model. This is a Gaussian process (GP) model of the response function(s) as a function of the parameters.
- Calculate the Expected Improvement Function (EIF). The EIF is defined as the expectation that any point in the search space will provide a better solution than the current best solution based on the expected values and variances predicted by the GP model.
- Evaluate the objective function at the point where the EIF is maximized.
- Repeat until convergence.

The EIF can be large either where there is a large model uncertainty (and thus more samples are needed to reduce the uncertainty) or where the model predicts a high objective function. In the first case the algorithm will explore spaces in the parameter domain that have not been sampled sufficiently. In the second case the algorithm chooses parameter combinations that exploit the region where the objective function is predicted to be maximized. Therefore, exploration and exploitation are balanced by using the EIF. The EGO algorithm often converges to a global optimum much faster than algorithms that do not use a surrogate model and 'blindly' sample the full parameter space like a scoping study. The EGO algorithm is expected to work well with up to $\approx$10 parameters according to the typical guidelines.

Before starting the EGO algorithm, samples need to be available to build a first surrogate model. If no information is available Dakota starts with evaluating $\frac{(p+1)(p+2)}{2}$ random samples (latin hypercube sampling), in which p is the number of parameters in the problem. It is possible to use previous simulation data as a starting point of the algorithm, using the `import_build_points_file` keyword. This will give the algorithm a 'warm start'. Using more samples is not always better. The algorithm works best when the initial samples are nicely spread out over the domain. Experience indicates that restarting the algorithm without importing previous data often provides better results. If imported data is clustered in variable space, the algorithm will sometimes converge to a bad optimum. Appendix 6 provides an evaluation of the effect of the statistical uncertainty on the result of the optimization.

Figure 1 shows all sample evaluations performed with the EGO algorithm for an example with two parameters. The horizontal and vertical axis show the value of the two parameters and the color of the dot indicates the value of the objective function. Clearly, more samples are chosen in the region close the maximum (red colors) than in the uninteresting region further away from the optimum (blue colors). By

using EGO algorithm more samples are evaluated closer to the maximal region compared to a typical parameter study across the whole parameter space.

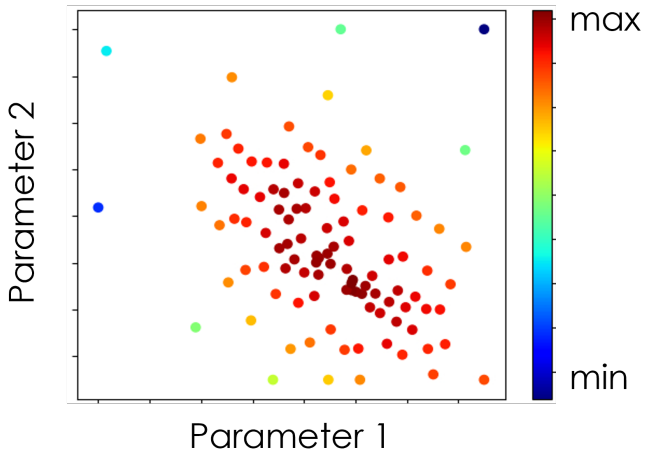

Figure 1. Evaluations with the EGO algorithm for an example with two parameters.

The EGO algorithm in Dakota stops when one of the convergence criteria is reached. The algorithm will stop when

- the EIF is sufficiently small for two iterations in a row (value given in the input file or 10^{-12} by default),
- the distance in the parameter space for the next iteration is sufficiently small (value given in the input file or 10^{-6} by default),
- a maximum number of iterations is reached (value given in the input file or 25*p, with p the number of parameters).

More information on these criteria can be found in the Dakota Reference Manual [9].

When the objective function is the weighted sum of the two or more objectives, and several sets of weights are selected, one optimal design is obtained for each set of weights and the optimal designs constitute the Pareto front. With two objectives, the Pareto front provides a curve of points with optimal designs. Each optimal design is maximized for a different weighted sum of the two objectives. With the `pareto_set` method in Dakota, the optimization is run sequentially for each of the chosen weight combinations. Each optimization run is called a 'set'. Using `pareto_set`, the initial samples that are used to build the first surrogate model for the EGO algorithm are re-used for each optimization run. While the first optimization (with the first combination of weights) requires a full run with no initial information, the next optimizations will require significantly fewer samples because of having a 'warm start'.

It is good practice to perform a parameter study around the optimal configuration. This will check if the obtained design is optimal and also provides information about the sensitivity of the objective function to all parameters. In this study, the `centered_parameter_study` method was used to perform these simulations.

When many samples are available, the user can perform their own fit to the data. Use of a fit smooths out the noise in the data and provides an expression which is easy to evaluate. For this study a cubic polynomial model was used. It is crucial to check if the polynomial fit captures the behavior of the data well. Not all data sets are suitable for a reliable fit. We discourage using this option if there are doubts about the quality of the model.

3.1.2 Interface with neutronics simulation

Dakota communicates with other codes in the optimization process using file exchange. A driver-, parameter- and response-file are specified in the interface block. The driver-script is written in Python. Apart from file exchange, pre- and post-processing of the results is also performed with this Python script. The parameter- and response-files for dakota are typically called 'params.in' and 'results.out'. If a simulation fails, we provide the word 'fail' in the output-file instead of response function values. A simulation can fail because of several reasons, for example when an unfeasable geometry is created and a segmentation fault occurs in MCNP. In this case, no MCNP output is created and the objective function cannot be evaluated.

Dakota offers several possibilities to cope with failed simulations. These following options were evaluated during the course of this study:

- `failure_capture continuation`: Dakota tries to estimate the value for the failed simulation by doing more samples nearby and extrapolating the result. Sometimes many simulations are performed to extrapolate the value for just one failed sample. These extra samples are, however, not taken into account by the optimization algorithm and are only used to estimate the value for the failed simulation. This option is not recommended for this work.
- `failure_capture recover nan`: Dakota omits the sample completely. This causes false convergence in the EGO algorithm. As the result is ignored by the algorithm, the next iteration will ask for the same sample. This causes the algorithm to converge for zero distance and the simulation stops.
- `failure_capture recover 0`: Dakota assumes a response function of 0 (it is also possible to select any other value than 0). This is the option that has been primarily used in this work. The advantage is that the algorithm keeps running and is not aborted. Therefore, samples continue to be accumulated without the need for human intervention. The disadvantage is that a false value (i.e. 0) is assumed. This will cause the surrogate model inside the EGO algorithm to be distorted, and may even cause convergence to values that are not optimal. Consequently, results using this option should be used with caution.
- `failure_capture abort`: Dakota stops the algorithm.
- Apart from these options within Dakota, we have tried to use a simplified model to give a value (in 'results.out') closer to the real value. This, however, is very speculative and is not recommended. Dakota will not separate the real values with those of the simplified model. In that case, the optimum may be governed by the simplified model instead of the real one.

We have not found an ideal way to deal with failed samples. It is always best to avoid failures and to take additional efforts to avoid having to deal with failed samples. When encountering a failure, it is often better to restart the algorithm with a slightly different set of successful samples only.

3.2 NEUTRONICS MODELS

Section 3.2 provides a brief overview of the moderator models and lists the design parameter ranges used in the optimization studies. A detailed description of the models can be found in [3].

3.2.1 Cylindrical moderator

Figure 2 shows a vertical cut of the geometry of the upper moderator. The design parameters are

- the radius of hydrogen vessel R_{mod} (32 to 80 mm)
- the thickness of the premoderator radially and on top R_{pm} (18 to 35 mm)
- the thickness of the premoderator on the bottom H_{pm} (25 to 40 mm)
- the radius of the beryllium vessel R_{Be} (150 to 200 mm)
- the position of the moderator in the direction of the proton beam X_{pos} (-20 to 20 mm)

For some simulations, the thickness of the premoderator R_{pm} is seperated in the radial thickness (still called R_{pm}) and the thickness on top H_{top}. The position X_{pos} is measured relative to the initial position at 87 mm from the outer radius of the target, which was obtained in previous studies. The thicknesses of the aluminum vessel walls are dependent on the radius of hydrogen vessel R_{mod} to ensure structural integrity[10].

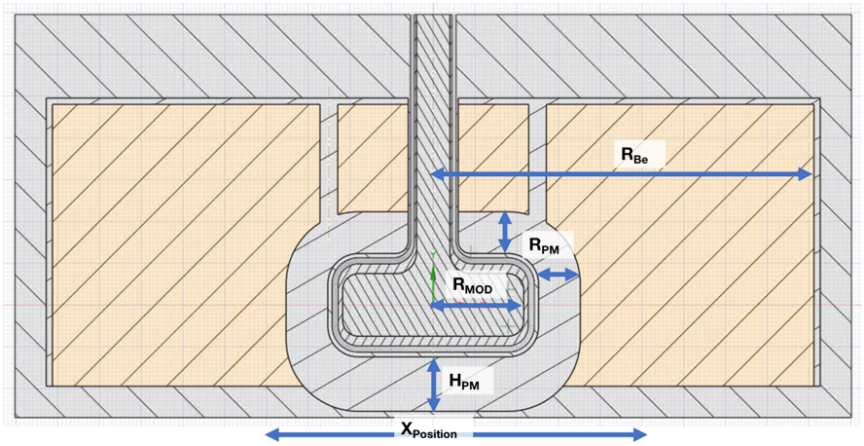

Figure 2. Vertical cut (x-y plane) of the geometry of the upper moderator with indication of the design parameters.

3.2.2 Tube moderator

Figure 3 shows a horizontal cut of the geometry of the lower moderator. The design parameters are

- the length of the long tube (i.e., the tube perpendicular to the proton beam direction) T_{len} (120 to 230 mm)
- the annular thickness of the premoderator R_{pm} (25 to 29.2 mm)

- the radius of the beryllium vessel R_{Be} (180 to 220 mm)
- the position of the moderator in the direction of the proton beam X_{pos} (-20 to 20 mm)

The position X_{pos} is measured relative to the initial position at 87 mm from the outer radius of the target, which was obtained in previous studies. It is noteworty to point out that the triangle is not equilateral but rather isosceles with the angle at the top of 70 degrees. This is necessary to obtain better spacing between the beam lines.

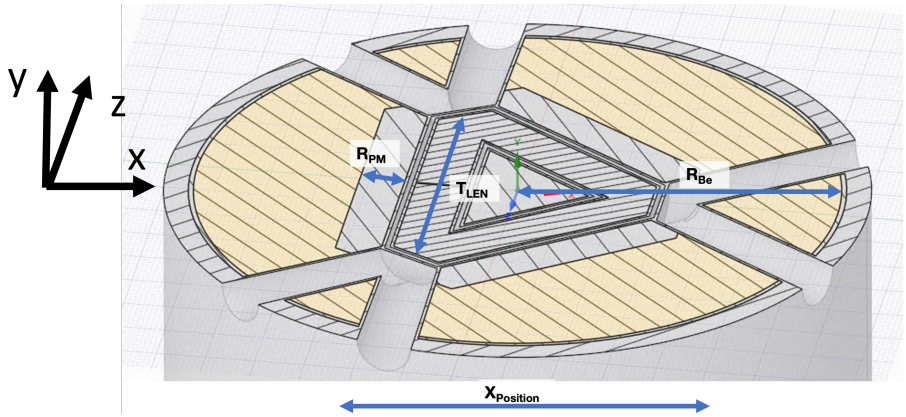

Figure 3. Horizontal cut (x-z plane) of the geometry of the lower moderator with indication of the design parameters. The proton beam is launched towards the target in the x-direction.

3.2.3 Brightness

Both moderators can be optimized either for maximum peak brightness, maximum time-integrated brightness, or any combination in between these extremes. The difference between the peak and time-integrated brightness is illustrated in Figure 4, which represents time-emission spectra for the 5Å neutrons produced by arbitrarily selected configuration of the tube and cylindrical moderators chosen for illustration purposes only. Peak brightness is the curve maximum, whereas time-integrated brightness is the area under the curve. In this configuration, the cylindrical moderator provides 5% higher peak-brightness, 55% narrower pulse (smaller full width at half maximum (FWHM) of the pulse) and a steeper leading edge, while the tube moderator offers 80% higher time-integrated brightness. Some dynamic neutron scattering instruments prefer superior time resolution, steep leading edge, and high peak brightness while other instruments perform best with the maximum time-integrated brightness and lower time resolution. Because of this delicate performance balance, the moderators need to be optimized in a close collaboration with instrument scientists.

The determination of an optimal moderator configuration is based on two performance metrics, which we refer to as peak brightness and time-integrated brightness figures of merit (FOMs), abbreviated as *peak* and *tint*. These FOMS, or weighted combinations of the two, will be used as objective functions for the optimization algorithm. The words FOM and objective function are used as synonyms in this report. Peak brightness FOM is obtained by integrating the flux $\Phi(E,t)$ over a range of neutron energies E at the time t

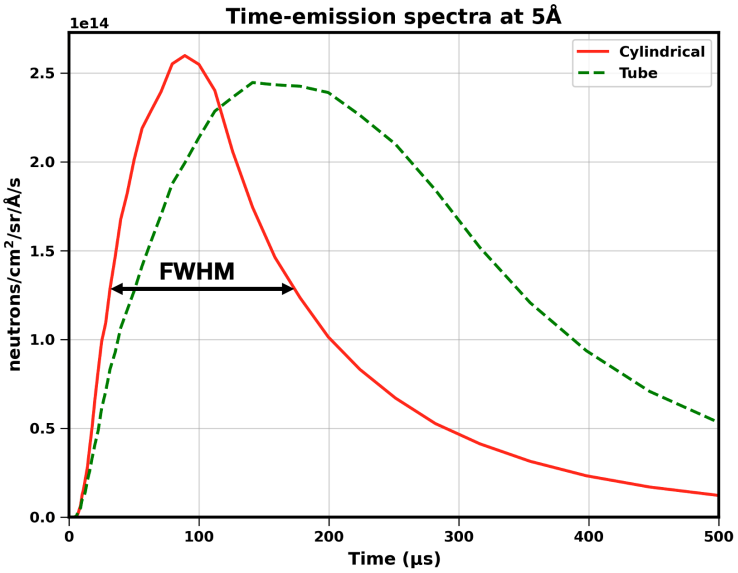

Figure 4. Time-emission spectra for the 5Å neutrons at a random configuration of the cylindrical and tube moderators. Peak brightness is the curve maximum, while time-integrated brightness is area under the curve.

when the flux reaches its maximum value:

$$peak = \int_0^{E_{max}} max_t(\Phi(E,t))dE, \tag{1}$$

where E_{max} is set to 10 meV. The time-integrated brightness FOM is given by

$$tint = \int_0^{E_{max}} \int_{t_{min}}^{t_{max}} (\Phi(E,t))dE. \tag{2}$$

The $\Phi(E,t)$ is the neutron flux calculated with the MCNP simulations from the point detector at 10 meters from the moderator surface. The contributions to the point detector are limited to the neutrons emitted from the viewed surface of the moderator only by a collimator (zero importance zones). The point detector applied in this work tallies the neutron time from the moderator viewed face. Detailed information on the calculation of these metrics can be found in [4].

For the cylindrical moderator, peak and time-integrated FOM are calculated as the sum of two point detectors tallies located at two beam lines. Most simulations were performed with 5 million histories. This gives a relative statistical precision (1σ) of ≈0.3% and ≈0.06% on *peak* and *tint* respectively. One run takes ≈1h on 8 nodes (48 tasks per node) on a Linux cluster. For the tube moderator, peak and time-integrated brightness are calculated as the sum of three point detector tallies. The reported absolute values of the FOMs are the sum of the two or three point detector tally results. Most simulations have been performed with 4 million particles, which corresponds to approximate relative errors of ≈0.2% and ≈0.04% for *peak* and *tint* respectively. One run takes ≈1h 20min on 8 nodes (48 tasks per node) on the Linux cluster. For both moderators, there is a large variability in run duration depending on the moderator size. The error values have been calculated based on the statistical uncertainties reported by MCNP for all energy and time bins. The effect of cross-correlations has not been taken into account. The reported values are thus underestimated, but still indicate the order of magnitude correctly.

4 ANALYSIS AND RESULTS

4.1 CYLINDRICAL MODERATOR

4.1.1 Optimization and sensitivity

For each optimization run (or set), the objective is a weighted sum of the two FOMs:

$$Objective = w_1 \frac{peak}{peak_{max}} + w_2 \frac{tint}{tint_{max}}. \quad (3)$$

The scaling factors $peak_{max} = 1.519 \cdot 10^{-4}$ neutrons/cm^2/proton/s and $tint_{max} = 2.7985 \cdot 10^{-8}$ neutrons/cm^2/proton are estimates of the maximum FOM values. This way both terms in the objective function have similar values and are close to 1.0. The w_1 and w_2 are the weights for each set, with w_1 chosen between 0 and 1, and $w_2 = 1 - w_1$.

A summary of the results is given in Table 1. Sets 1 and 6 are the optimal designs with maximum time-integrated FOM *tint* and maximum peak FOM *peak* respectively. Sets 2 to 5 are optimal designs for a balance of the two FOMs according to the weights listed in the table and equation 3. The most sensitive parameter is the moderator radius R_{mod} which is the largest for the *tint*-optimized design and the smallest for the *peak*-optimized design. Designs with a balanced objective have values inbetween these extremes. The same trend is seen for the radial and top premoderator thickness R_{pm} and the beryllium radius R_{Be}. The bottom premoderator thickness H_{pm} and the horizontal position X_{pos} remain similar for all designs along the Pareto-front.

In this simulation, the Pareto-front is obtained using a total of only 38 design iterations. With five independent design parameters, the EGO algorithm in Dakota needs 21 runs to build its first surrogate model. These 21 initial runs are re-used for each subsequent set, which significantly lowers the number of new runs that are required for each subsequent set.

Table 1. Results of Dakota Pareto set run for the upper moderator

Set	**Nr of runs (new)**	w_1	w_2	R_{mod} **[mm]**	R_{pm} **[mm]**	H_{pm} **[mm]**	R_{Be} **[mm]**	X_{pos} **[mm]**	**Peak [10^{-4}cm^{-2}s^{-1}]**	**Tint [10^{-8}cm^{-2}]**
1	25 (24)	0	1	61.4	31.8	29.2	194	3.2	1.254	2.893
2	25 (2)	0.25	0.75	55.0	29.3	29.0	188	1.8	1.338	2.855
3	25 (2)	0.5	0.5	47.5	29.3	28.7	178	-0.2	1.439	2.750
4	25 (2)	0.7	0.3	40.9	29.5	29.1	181	1.4	1.497	2.577
5	27 (4)	0.85	0.15	38.8	26.0	29.5	159	2.7	1.530	2.446
6	27 (4)	1	0	35.7	24.6	29.0	150	4.0	1.539	2.279

The relative values for *tint* and *peak* FOM for all iterations are shown in Figure 5. The Pareto front is shown with the orange lines. The green and red crosses indicate the design cases that are within 2% of the maximal *tint* and *peak* respectively. The *peak*-optimized design reaches ≈80% of the maximal *tint*-value. Also the *tint*-optimized design reaches ≈80% of the maximal *peak*-value. The optimal designs inbetween have *tint* and *peak* values inbetween 80% and 100%.

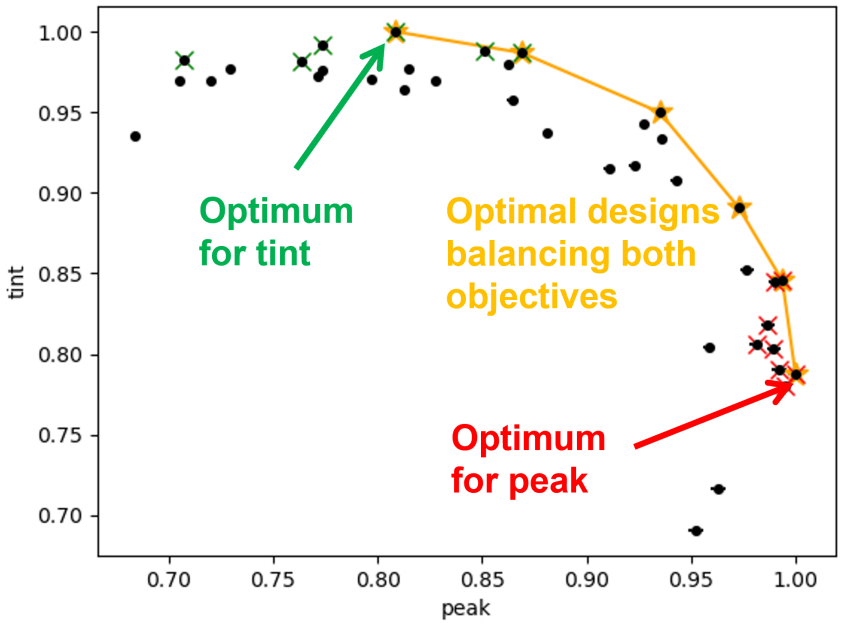

Figure 5. Relative *tint* and *peak* FOM for all iterations in one Dakota Pareto set run for the cylindrical moderator. The orange line indicates the Pareto front. The green and red crosses indicate the designs that are within 2% of the maximal *tint* and *peak* FOM respectively.

To smooth out the noise and have a simple functional description to work with, a cubic polynomial is fitted through the data (see section 3.1.1). Here we use 196 MCNP runs in total (coming from exploration runs and other optimization runs). In Figure 6, the *tint* and *peak* FOM are plotted from all MCNP runs. The orange line is the Pareto front using a cubic polynomial model. This polynomial fit has also been used to evaluate the effect of the statistical uncertainty of the solution in Appendix F. More plots of this data can be found in Appendix D.

It is interesting to inspect the values for the objective functions and the choice of parameters along the Pareto front. In Figure 7, the dots show the values obtained with the Pareto set run with MCNP (corresponding to Table 1). The solid lines are obtained by performing Pareto optimization with FOM obtained from the cubic polynomial fit and using much finer grid for the weights. From left to right in the figure, we walk along the Pareto-front from the *tint*-optimized design (w_1=0) to the *peak*-optimized design (w_1=1), with w_1 defined in equation 3. The range over which the optimal designs vary is also indicated. The results from the polynomial fit confirm and strengthen the previous findings: the *tint*-optimized moderator design has a larger radial dimension and larger premoderator thickness than the *peak*-optimized moderator design. For the other parameters, the results of the polynomial fit do not match the MCNP results well and, therefore, no strong conclusion can be taken.

We select three optimal designs: the *tint*-optimized design, the *peak*-optimized design, and a middle design that balances both objectives. The middle design is chosen as the design that has the minimal distance to utopia. Utopia is the (unfeasible) design which has both the maximal *peak* and *tint* FOM, i.e. point (1,1) in

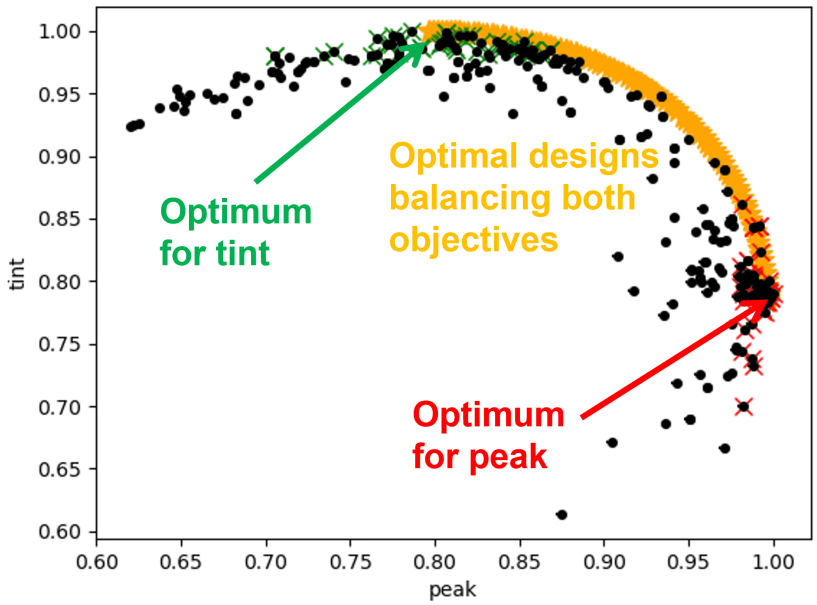

Figure 6. Relative *tint* and *peak* for all MCNP runs for the cylindrical moderator. The orange line indicates the Pareto front based on a polynomial fit. The green and red crosses indicate the designs that are within 2% of the maximal *tint* and *peak* respectively.

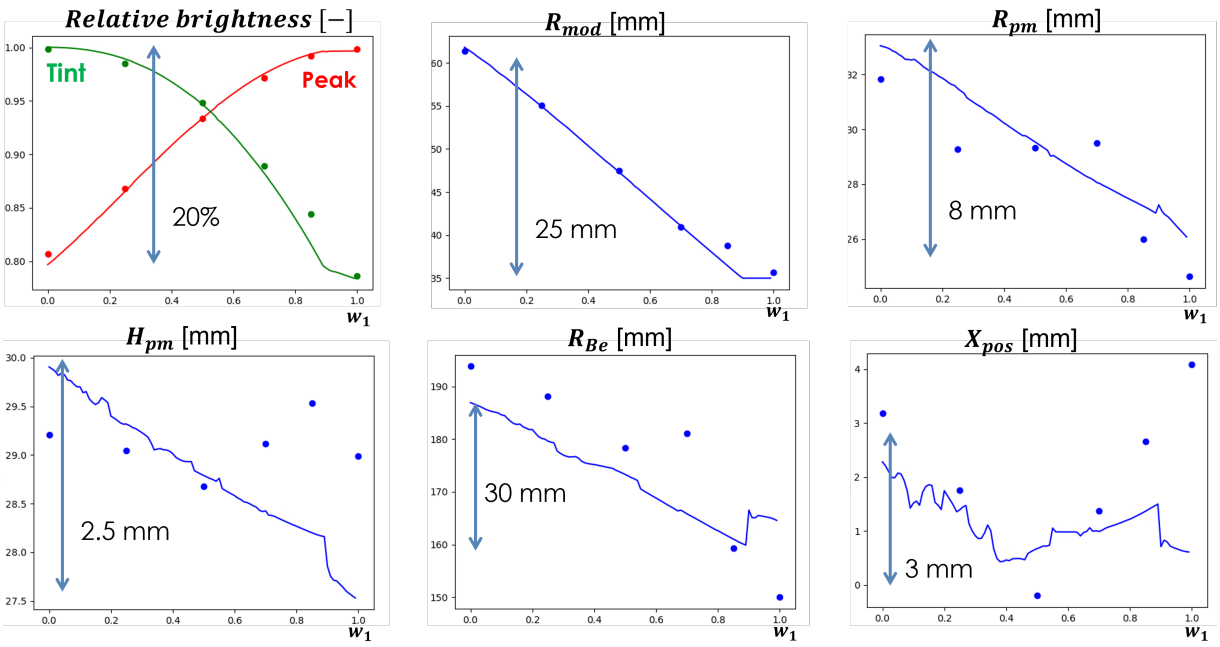

Figure 7. Objective functions and the parameters along the Pareto front obtained with MCNP runs (dots) and along the Pareto front obtained with a polynomial fit (lines) for the cylindrical moderator.

Figure 6. The distance can be calculated as $\sqrt{(1-peak)^2+(1-tint)^2}$, with *peak* and *tint* taken as their relative value to the maximum. From the simulations, the middle design with $w_1 \approx 0.5$ results in the minimal distance, with a peak and time-integrated FOMs that are each 94% of their maximum values. Table 2 lists the design parameters and brightness FOMs for the three chosen designs. Geometry plots of these models can be found in Appendix C.

Table 2. Optimal designs based on the surrogate model.

	R_{mod} **[mm]**	R_{pm} **[mm]**	H_{pm} **[mm]**	R_{Be} **[mm]**	X_{pos} **[mm]**	**Peak** **[10^{-4}cm^{-2}s^{-1}]**	**Tint** **[10^{-8}cm^{-2}]**
Time int	62	33	30	187	2	1.218 (80%)	2.899 (100%)
Middle	47	29	29	173	1	1.435 (94%)	2.734 (94%)
Peak	35	26	28	165	1	1.524 (100%)	2.276 (79%)

The sensitivity of the FOMs to the design parameters is investigated by using a centered parameter study around the optimal designs. In these studies, one parameter is varied in 5 mm increments while keeping all others constant. The optimum is sensitive to the varied parameter if a large change in *peak* and/or *tint* is observed. The results are shown in Figures 8, 9 and 10 for respectively the *peak*-optimized design, a design close to the *tint*-optimized design, and the optimal middle design. The dots in these figures are results with MCNP.

The largest sensitivity is seen by varying the moderator radius R_{mod} and the premoderator thicknesses R_{pm} and H_{pm}. A variation of ≈5 mm in these parameters will cause a 1% decrease in performance (*peak* and *tint*) around the optimum. Further away from the optimum, the change is larger (see figures for more details). The radius of the beryllium vessel R_{Be} and the horizontal position of the moderator X_{pos} have a much smaller sensitivity. To obtain a 1% decrease in performance around the optimum, these parameters have to change by ≈20 mm.

It is noteworthy to point out that the design used for the sensitivity study in Figure 9 is not the optimum as listed in Table 2. It can be observed in Figure 9 that *tint* can be improved by choosing a larger R_{pm} and H_{pm}. Indeed, the values of these parameters in the *tint*-optimized design from Table 2 are larger than the design around which the parameter study is performed (indicated with stars in the Figure 9).

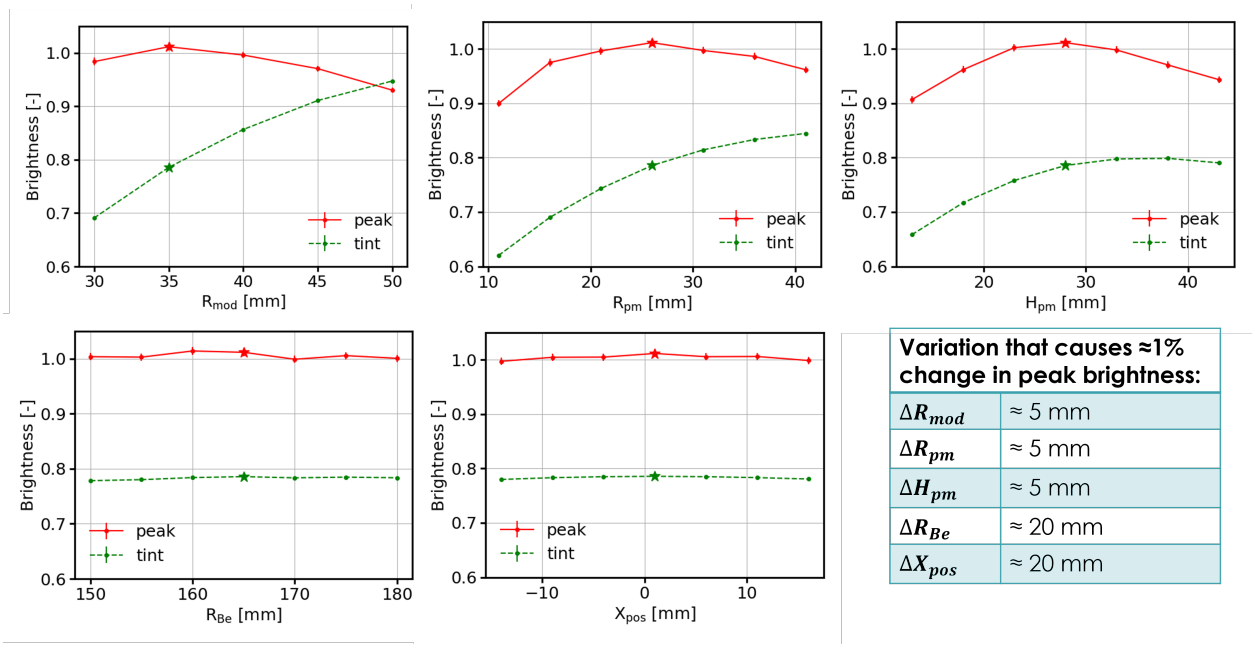

Variation that causes ≈1% change in peak brightness:	
ΔR_{mod}	≈ 5 mm
ΔR_{pm}	≈ 5 mm
ΔH_{pm}	≈ 5 mm
ΔR_{Be}	≈ 20 mm
ΔX_{pos}	≈ 20 mm

Figure 8. Sensitivity study around the *peak*-optimized design.

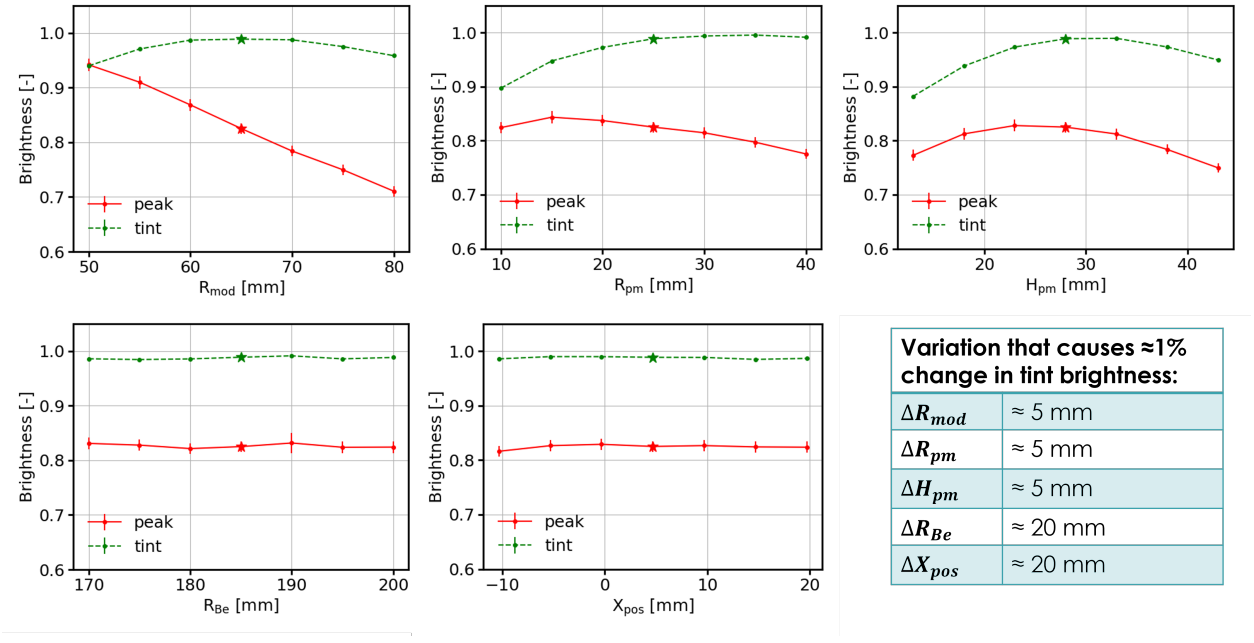

Variation that causes ≈1% change in tint brightness:	
ΔR_{mod}	≈ 5 mm
ΔR_{pm}	≈ 5 mm
ΔH_{pm}	≈ 5 mm
ΔR_{Be}	≈ 20 mm
ΔX_{pos}	≈ 20 mm

Figure 9. Sensitivity study around a design close to the *tint*-optimized design.

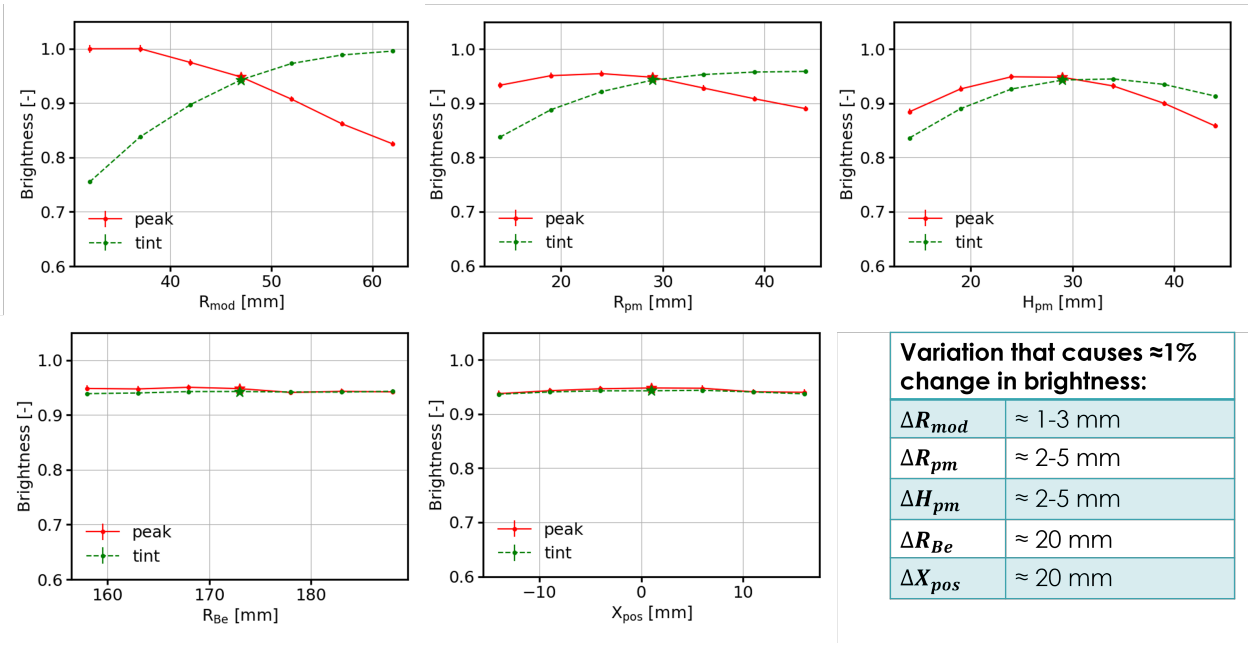

Variation that causes ≈1% change in brightness:	
ΔR_{mod}	≈ 1-3 mm
ΔR_{pm}	≈ 2-5 mm
ΔH_{pm}	≈ 2-5 mm
ΔR_{Be}	≈ 20 mm
ΔX_{pos}	≈ 20 mm

Figure 10. Sensitivity study around the optimal middle design.

4.1.2 Separate radial and top premoderator thickness

A parametric study was conducted to determine whether using separate values of the radial and top premoderator thickness, R_{pm} and H_{top}, rather than a single value for both, has a significant effect on the peak and time-integrated brightness. In this study, each parameter (R_{pm} and H_{top}) was varied in 2.5-mm intervals around the optimal design. All other parameters were kept constant. The results are shown in Figure 11. A variation of 5 mm around the optimal configuration remains within 1% of *peak* and *tint* around the optima, which is consistent with the results in the previous section.This study confirms that varying the radial and top premoderator thickness separately does not significantly influence the optimal configuration.

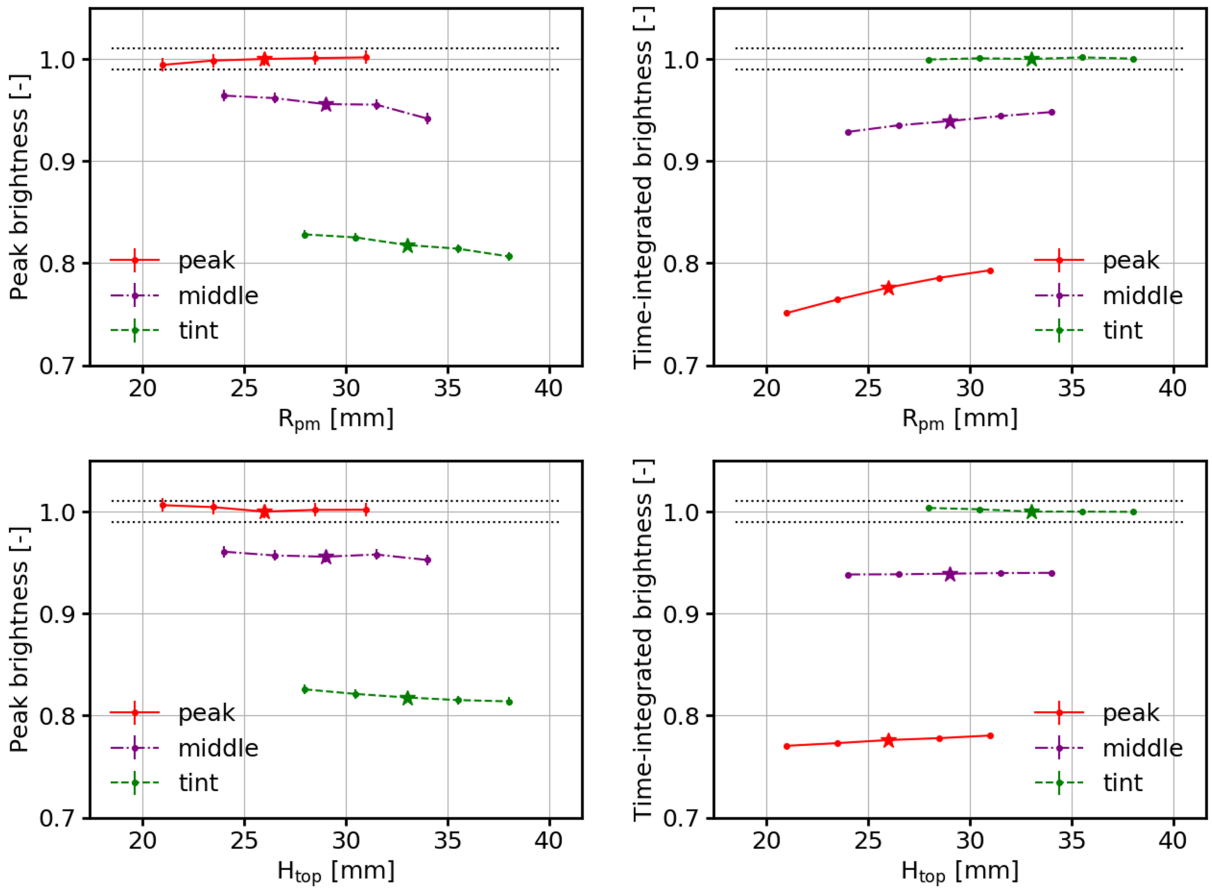

Figure 11. Sensitivity study around the optimal designs (*peak*, *tint* and middle) for separate radial and top premoderator thicknesses, R_{pm} and H_{top}. The dotted black lines indicate ±1% around the optimum.

4.1.3 Effect of aluminum vessel thickness

In previous STS moderator optimization studies, the thicknesses of the hydrogen and vacuum vessels were fixed at constant values. For the current study, vessel thickness values from [10] were used. These values are functions of the moderator radius and are dictated by structural requirements. Table 3 gives an overview of the wall thicknesses for three moderator radii that are close to three optimal designs. It also lists the values that were used in the previous optimization study. Larger R_{mod} result in increases in the thickness of all vessel walls. The previous fixed values are in between the *peak*-optimized design and the middle design.

A Pareto set optimization was run with the fixed wall thickness values in Table 3. Table 4 compares the

Table 3. Vessel wall thicknesses for several designs

	H-vessel radius R_{mod} [mm]	H-vessel wall flat [mm]	H-vessel wall cylindrical [mm]	Vacuum vessel wall flat [mm]	Premoderator vessel wall flat [mm]
peak-optimized	40	4.24	1.77	2.06	2.8
middle	50	5.35	2.8	2.61	3.38
tint-optimized	62	6.67	4.34	3.26	4.06
previous fixed values		4.6	2.1	2.25	3.0

Table 4. Optimal designs with a fixed and variable thickness

	R_{mod} [mm]	R_{pm} [mm]	H_{pm} [mm]	R_{Be} [mm]	X_{pos} [mm]	Peak [10^{-4}cm^{-2}s^{-1}]	Tint [10^{-8}cm^{-2}]
Peak optimized design							
Fixed	40	27	29	200	4	1.494	2.484
Variable	35	26	28	165	1	1.524 (+2%)	2.276 (-9%)
Time int optimized design							
Fixed	80	26	28	200	1	1.260	3.312
Variable	62	33	30	187	2	1.218 (-4%)	2.899 (-13%)

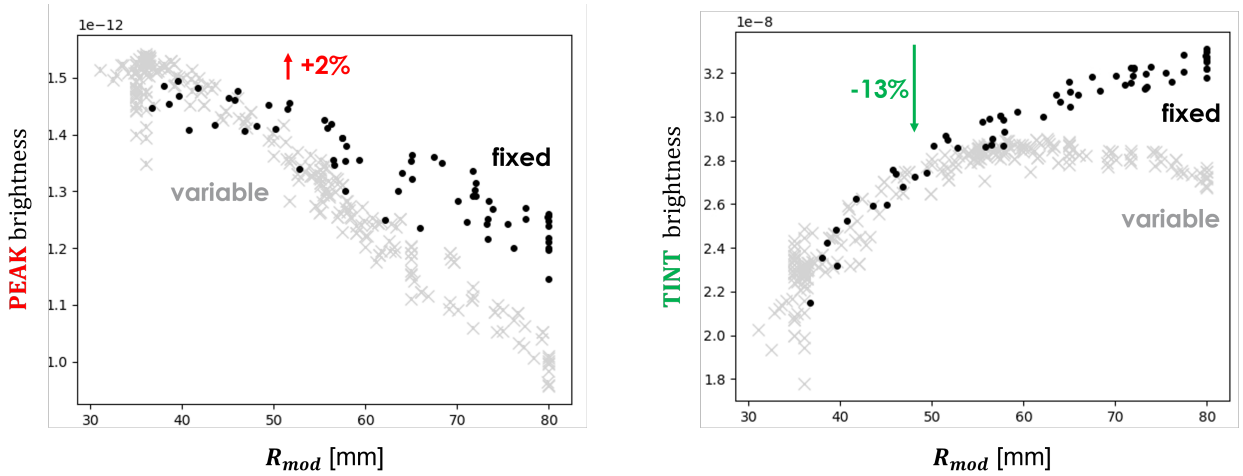

Figure 12. Brightness FOMs *peak* and *tint* as a function of the moderator radius R_{mod} with variable (grey crosses) and fixed (black dots) wall thicknesses for the cylindrical moderator.

optimal designs for maximal *peak* and *tint* with fixed and variable thicknesses. Geometry plots of both models are provided in Appendix C. In Figure 13, the Pareto front for the results with fixed (black dots and purple line) and variable thicknesses (grey crosses and orange line) are compared. Figure 12 shows *peak* and *tint* FOMs as a function of the moderator radius.

Because the aluminum wall thicknesses for the *peak*-optimized design are less than the original fixed wall thicknesses, the peak FOM increases slightly ($\approx$ 2%) for this configuration. The time-integrated FOM is relatively insensitive to the reductions in the wall thicknesses. The curves of the wall thicknesses as a function of R_{mod} in [10] do not extend below 40 mm. If those curves are extrapolated to smaller values of

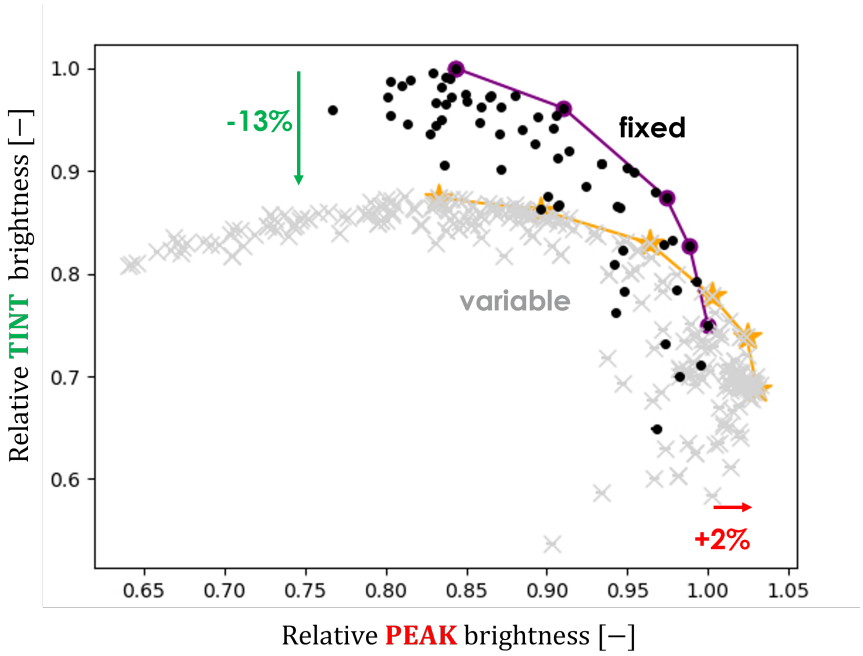

Figure 13. Relative *peak* and *tint* for designs with variable (grey crosses) and fixed (black dots) wall thicknesses. The Pareto fronts are indicated with the orange and purple lines for variable and fixed thicknesses respectively.

R_{mod}, the thinner aluminum walls result in a slight increase in peak FOM. However, structural integrity constraints place a lower bound on all of the wall thicknesses. While the minumum wall thicknesses were not explicitly calculated for R_{mod} values less than 40 mm, a minumum wall thickness of 2 mm was considered to be a reasonable lower bound. As some of the MCNP simulations have wall thicknesses lower than 2 mm, this additional constraint is taken into account in the next simulations.

The aluminum wall thicknesses for the *tint*-optimized design are significantly larger than the original fixed wall thicknesses. Because of this, the time-integrated brightness reaches a maximum at an R_{mod} value of ≈60 mm and then decreases with increasing values of R_{mod}. Consequently, the maximum time-integrated brightness is reduced by ≈13% relative to the original study. Figure 12 suggests that the maximum *tint* value may occur for an R_{mod} value greater than 80 mm with the fixed values of the wall thicknesses. This is confirmed by the optimization study in [4], which shows a maximum value of *tint* at an R_{mod} value of ≈90 mm.

The parameter study with R_{mod} around the previous optimum R_{mod}= 35 mm was repeated with applying the 2 mm minimal thickness of the aluminum walls. The blue and red symbols in Figure 14 show the results with and without this lower limit. The error bars indicate the 1σ precision. From these results, we see that there is no significant difference for the results with and without 2 mm limit. With 30 mm ≤ R_{mod} ≤ 42.5 mm, the peak FOM remains within 2% of the *peak*-optimized design of section 4.1.1. Because of this broad optimum and the increasing engineering difficulties with smaller R_{mod}, a design with R_{mod}= 40 mm was selected as *peak*-optimized design.

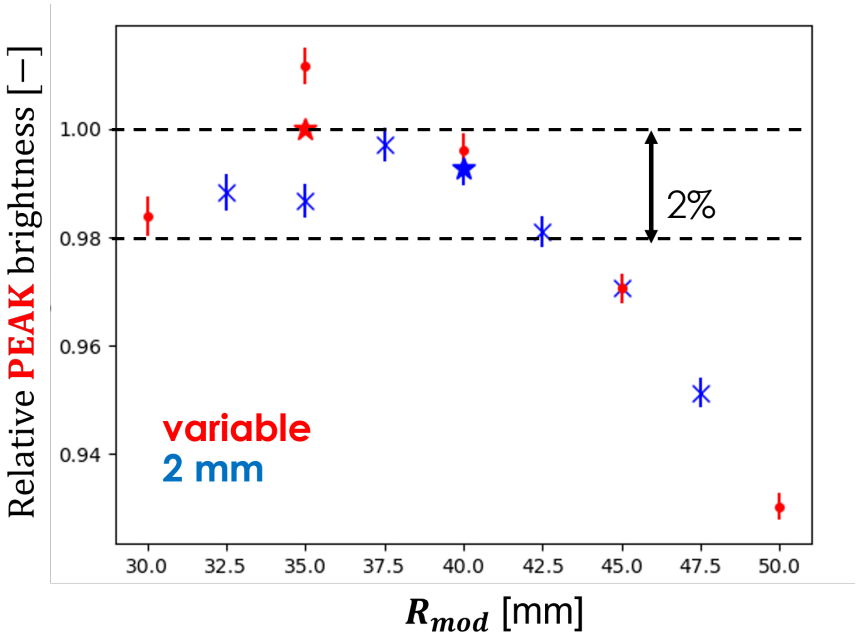

Figure 14. Parameter study around the *peak*-optimized design. A lower limit of 2 mm wall thickness is taken into account in the blue symbols. The blue star with R_{mod}=40 mm is suggested as the *peak*-optimized design instead of the previously selected red star with R_{mod}=35 mm.

4.1.4 Effect of location of point-detector tallies

In previous simulations the brightness FOMs for the cylindrical moderator were calculated as the sum of two point detector tallies (for more details, see [3]). Here, we take four point detector tallies located in four beamlines. The tally locations are chosen in an asymetric way, as can be seen from Figure 15 which gives a graphical view of the geometry around the moderator and the selection of the beamlines in which the tallies are placed.

The results of the Pareto set optimization with *peak* and *tint* as the sum of the FOMs of these four tallies are shown in Table 5. The resulting designs are very similar to those obtained before. In Figure 16 the average *peak* and *tint* as well as the individual tally results are shown for the Pareto front with four and two tallies in respectively red and blue colors. The average FOM with four tallies is ≈1% higher than the average with two tallies. The difference between FOMs of the four tallies separately is ≈4-7%. Tally 4, which gives the largest brightness metrics is forward directed with regard to the proton beam.

We conclude that the two point detector tallies give a good representation for the FOMs in all beamlines.

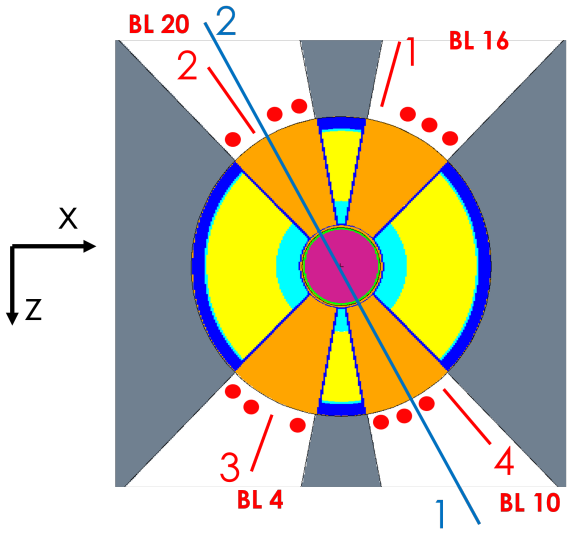

Figure 15. Schematic representation of the beamlines around the upper moderator and the choice of tally locations for the two-tally (blue) and four-tally simulations (red line = tally, red dot = no tally). Also the beamline numbers (BL) for the chosen four-tally positions are indicated.

Table 5. Results of Pareto set Dakota run with four tallies.

Set	**Nr of runs (new)**	w_1	w_2	R_{mod} **[mm]**	R_{pm} **[mm]**	H_{pm} **[mm]**	R_{Be} **[mm]**	X_{pos} **[mm]**	**Peak** **[10^{-4}cm^{-2}s^{-1}]**	**Tint** **[10^{-8}cm^{-2}]**
1	25 (24)	0	1	59.5	34.0	30.1	187	3.4	2.543	5.882
2	36 (15)	0.25	0.75	54.5	33.1	31.1	200	2.4	2.679	5.821
3	25 (2)	0.5	0.5	46.9	31.2	30.0	177	8	2.921	5.606
4	35 (12)	0.7	0.3	40.8	29.0	29.1	172	-4	3.078	5.241
5	27 (2)	0.85	0.15	38.9	28.6	27.4	178	14	3.083	5.042
6	28 (4)	1	0	35.5	28.8	26.7	178	19	3.094	4.723

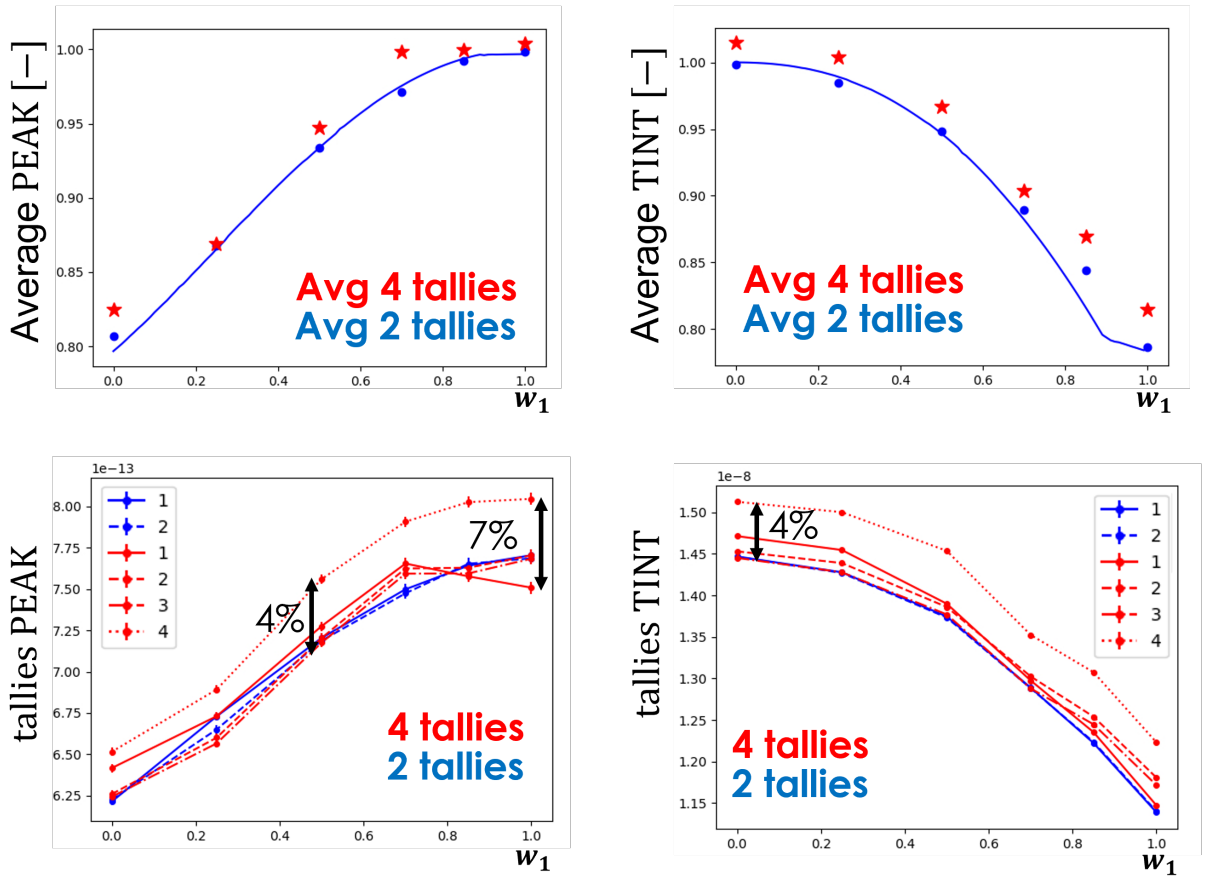

Figure 16. Brightness FOMs obtained from optimization runs with two (blue) and four (red) tallies. The top figures show the average *peak* and *tint* FOM values. The bottom figures show all tally results separately.

4.1.5 Final optimization run with model updates

To conclude the optimization of the upper moderator, the optimization algorithm was run again with all updates combined. The lower limit for the wall thickness is set at 2 mm, tallies are located at four beamline locations and the radial and top premoderator thickness can be varied seperately. In addition to these changes, a newer MCNP executable that corrects a bug in point detector tallies was used. With the new executable a slightly lower (≈3-4%) value is expected for the brightness metrics.

The EGO algorithm was run for the six weighted objectives as before (see equation 3). with six parameters, the initial number of samples needed to build the first surrogate model was 28. The maximum number of iterations for each set was set to 10. The results are shown in Table 6. Only one of the sets (set 5) did not converge within 10 iterations. One iteration in this set failed (and returned a zero solution), which interfered with a fast and reliable convergence of the algorithm. It is clear from looking at the results that the unconverged 5th set does not give the expected design as the solution of the 4th set has both a higher *peak* and *tint*, while a higher *peak* is expected. The relative *peak*- and *tint*- values of all samples are shown in Figure 17, in which the optimal points are connected with the orange line. The green and red crosses indicate solutions that are within 2% of the maximal *tint* and *peak* respectively.

Overall, very similar designs result from this final optimization run compared to the first optimization run. All model updates have a minor effect on the brightness metrics and the parameters of the optimal designs.

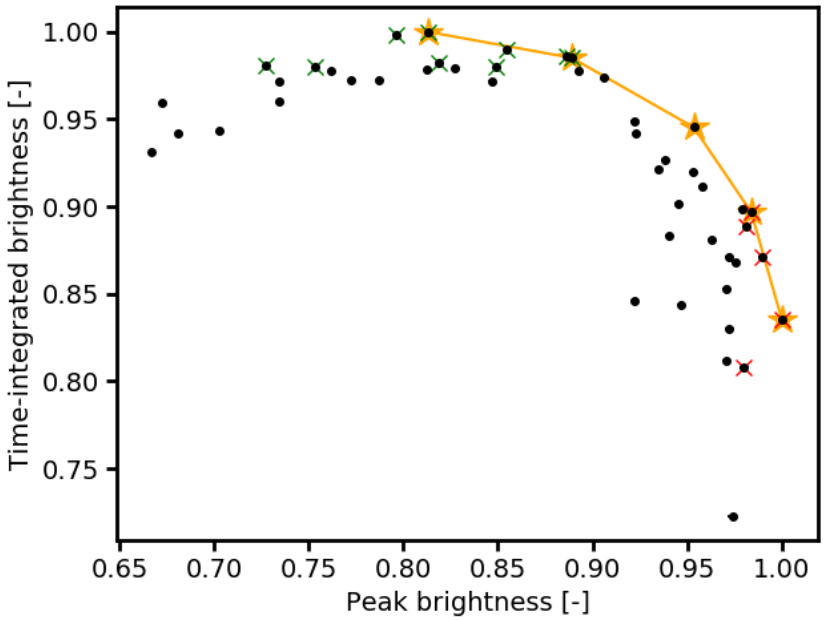

Figure 17. Relative *tint* and *peak* for all MCNP runs for the final upper moderator optimization. The orange line indicates the Pareto front. The green and red crosses indicate the designs that are within 2% of the maximal *tint* and *peak* respectively.

Table 6. Results of final Pareto set run for upper moderator. The results of set 5 are typed cursively to indicate that these results did not properly converge.

Set	Nr of runs (new)	w_1	w_2	R_{mod} [mm]	R_{pm} [mm]	H_{pm} [mm]	H_{top} [mm]	R_{Be} [mm]	X_{pos} [mm]	Peak [10^{-4}cm^{-2}s^{-1}]	Tint [10^{-8}cm^{-2}]
1	33(31)	0	1	61.5	35.0	31.3	25	183	2	2.416	5.609
2	32 (2)	0.25	0.75	53.2	32.5	31.5	28.4	178	-1	2.642	5.525
3	31 (1)	0.5	0.5	46.7	30.1	30.4	25.0	173	-1	2.833	5.302
4	32 (2)	0.7	0.3	42.9	25.7	29.2	25.0	172	4	2.924	5.030
5	*38 (10)*	*0.85*	*0.15*	*41.6*	*25.6*	*31.5*	*35*	*167*	*5*	*2.917*	*4.984*
6	32 (2)	1	0	38.9	24.3	28.9	25.0	178	-4	2.972	4.686

4.1.6 Final optimal configurations

The final optimal designs are listed in Table 8. This selection of designs has been made based on a comparison between the CSG [4] and UMG results based on the sum of two point detector results, see Table 7. The brightness FOMs for the final configurations in Table 8 are based on beamline 15. The simulations have been performed with 50M particles to ensure sufficient accuracy to construct neutron sources which are used in the design of the instruments. A detailed comparison between the pulse shapes obtained for these designs with the CSG and UMG model can be found in [3].

Table 7. Comparison between optimal solutions obtained with UMG and CSG.

	R_{mod} [mm]	R_{pm} [mm]	H_{pm} [mm]	R_{Be} [mm]	X_{pos} [mm]	Peak [10^{-4}cm^{-2}s^{-1}]	Tint [10^{-8}cm^{-2}]
peak CSG	40	20	27.5	162.5	0	1.54	2.45
peak UMG	40	27	29	165	2	1.49	2.48
middle CSG	50	20	30	172.5	0	1.43	2.80
middle UMG	47	29	29	173	1	1.44	2.73
tint CSG	62.5	20	30	182.5	0	1.25	2.94
tint UMG	62	33	30	187	1	1.22	2.90

Table 8. Final optimal designs for the cylindrical moderator

	R_{mod} [mm]	R_{pm} [mm]	H_{pm} [mm]	R_{Be} [mm]	X_{pos} [mm]	Peak [10^{-5}cm^{-2}s^{-1}]	Tint [10^{-8}cm^{-2}]
peak	40	20	27.5	162.5	0	7.54	1.18
middle	50	29	30	172.5	0	7.04	1.41
tint	62	33	30	182.5	0	6.06	1.45

4.2 TUBE MODERATOR

The tube moderator optimization study makes use of the results obtained with the CSG model [4], in which optimal configuration for the *peak*, middle and *tint* FOM have been selected based on a scoping study with two parameters (tube length and annular premoderator thickness). We start by verifying that these three designs are indeed optimal by performing a parameter study around the parameters with the UMG model. Afterwards, a Pareto set optimization is performed varying four parameters (tube length, annular premoderator thickness, beryllium radius, horizontal position of the moderator).

4.2.1 Parameter study around CSG optima

In Table 9 the parameters of the optimal designs are listed. The design parameters have been decided based on the scoping study with a CSG model [4]. Around these designs a parameter study is performed to evaluate the sensitivity of the brightness FOMs *peak* and *tint* for each parameter. Three point detector tallies are considered in calculating the FOMs. In this section, we confirm that these parameters are indeed optimal and that the optimum is not sensitive to a small variation of the parameters.

Table 9. Parameters of optimal configurations (selected with CSG model, see [4] for which a sensitivity study is performed.

	T_{len} **[mm]**	R_{pm} **[mm]**	R_{Be} **[mm]**	X_{pos} **[mm]**	*peak* **[10^{-4}cm^{-2}s^{-1}]**	*tint* **[10^{-8}cm^{-2}]**
peak	125	25	220	0	2.273	5.049
middle	170	27.5	220	0	2.200	5.847
tint	210	27.5	220	0	2.068	6.083

The results for this parameter study are shown in Figures 18, 19, 20 and 21. The stars in these figures are the MCNP results for the optimized designs. The colors green, purple and red are used to indicate the results of the parameter study around the *tint*-, combined and *peak*-optimized designs respectively. The dashed lines indicate the ±1% values around the optimum.

Figure 18 illustrates the effect of the tube length. This parameter is chosen very differently for each of the optimal designs, going from 125 mm for the *peak*-optimized design to 210 mm for the *tint*-optimized design. The simulation is run for 5-mm differences for three steps in each direction, resulting in a 30-mm variation in total around the optimal design. For the tint-optimized design, the optimized metric remains within 1%, indicating a broad optimum. Also for the *peak*-optimized design, most evaluations fall within the 1% band, but for smaller tube lengths, the *peak* brightness decreases slightly more.

To evaluate the effect of the premoderator thickness, we look in 2.5-mm steps from the optimized designs (see Figure 19). A variation of ± 2.5 mm remains within 1%, but a variation of ± 5 mm decreases the brightness by 1-2%.

For the optimized designs, a beryllium radius of 220 mm was chosen. This parameter was not a part of the CSG optimization study. Figure 20 illustrates the effect of decreasing the beryllium radius by 15 mm and 30 mm. A 30 mm difference results in ≈1% change in FOMs. Decreasing the beryllium radius by 30 mm decreases *tint*, but increases *peak*. A small gain in *peak* brightness is expected if the design has a smaller beryllium radius.

The sensitivity of the horizontal moderator position is very similar to what we observed with the cylindrical moderator: a 20 mm variation results in ≈1% decrease in brigthness, see Figure 21.

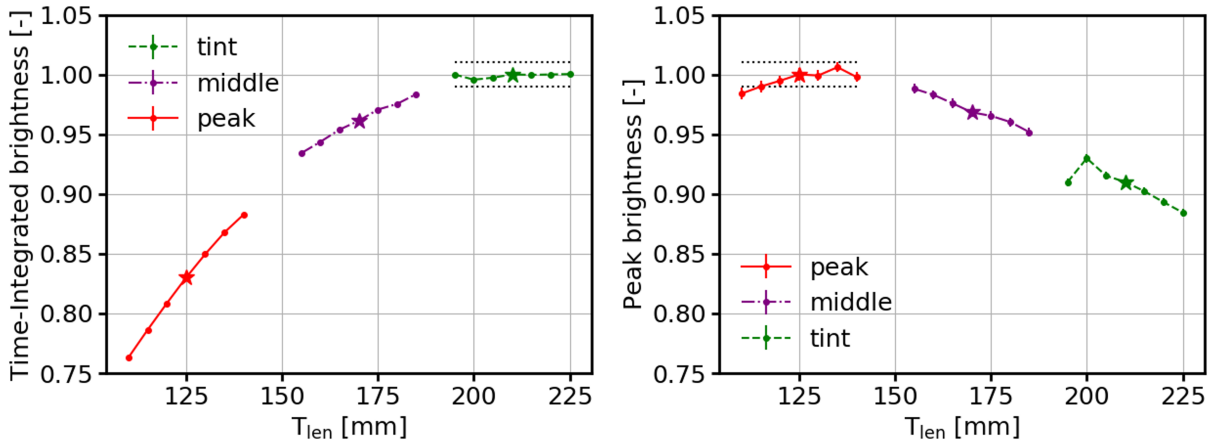

Figure 18. Optimization metrics (relative to value of the optimal design) as a function of tube length. The colors green, purple and red indicate the results of the parameter study around the *tint*-, combined, and *peak*-optimized designs respectively. The stars indicate the value for the chosen optimal designs.

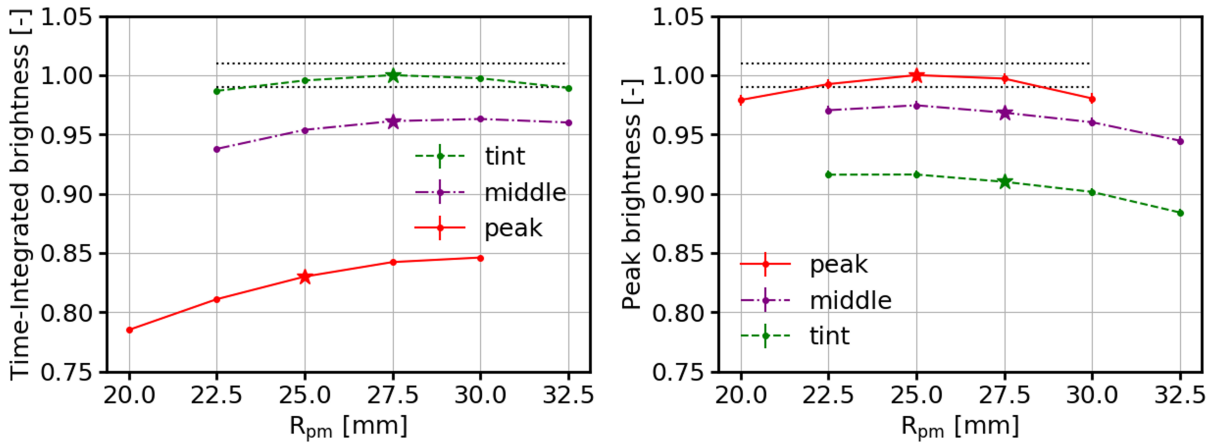

Figure 19. Optimization metrics (relative to value of the optimal design) as a function of premoderator thickness. The colors green, purple and red indicate the results of the parameter study around the *tint*-, combined, and *peak*-optimized designs respectively. The stars indicate the value for the chosen optimal designs.

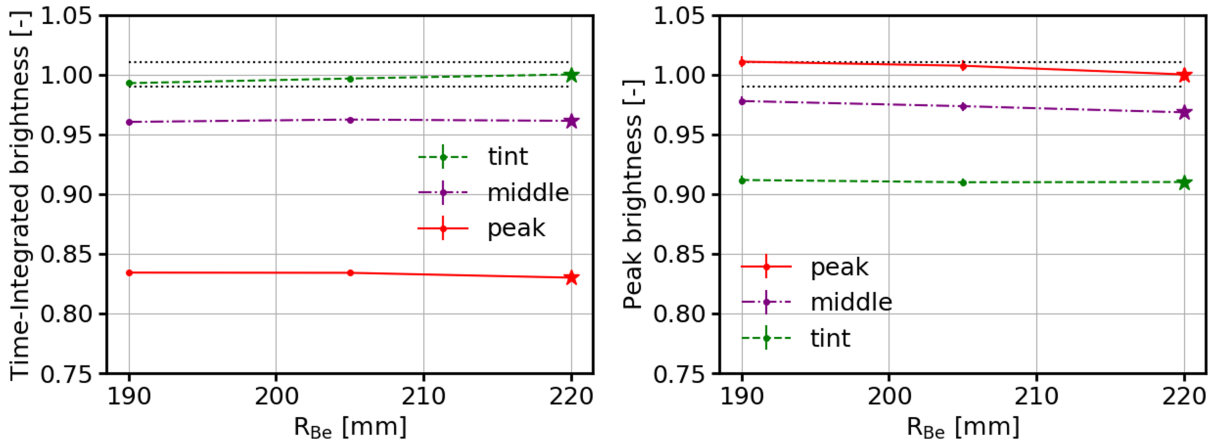

Figure 20. Optimization metrics (relative to value of the optimal design) as a function of beryllium radius. The colors green, purple and red indicate the results of the parameter study around the *tint*-, combined, and *peak*-optimized designs respectively. The stars indicate the value for the chosen optimal designs.

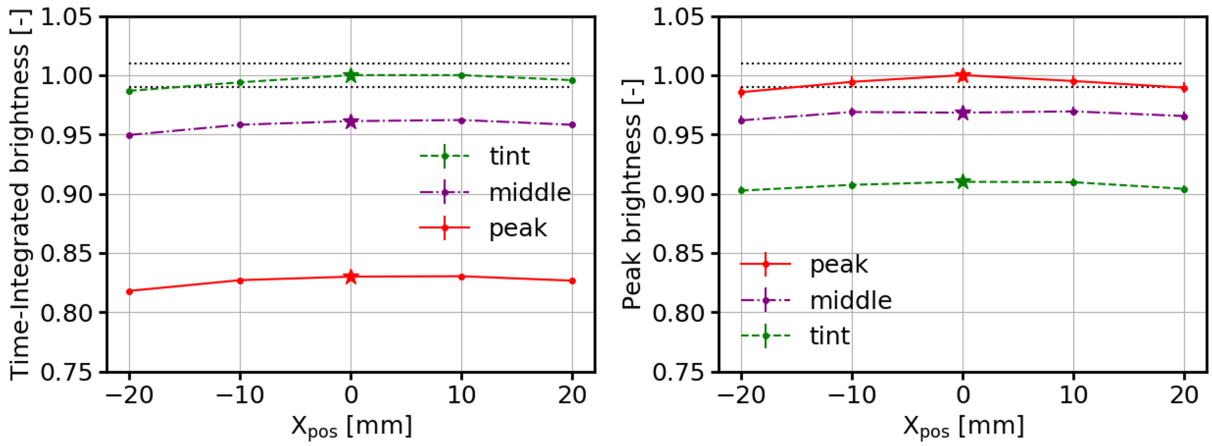

Figure 21. Optimization metrics (relative to value of the optimal design) as a function of the moderator position. The colors green, purple and red indicate the results of the parameter study around the *tint*-, combined and *peak*-optimized designs respectively. The stars indicate the value for the chosen optimal designs.

Contour plots of the brightness metrics as a function of the tube length and the premoderator thickness are shown in Figure 22. The yellow line indicates the contour in which the metric (either *tint* or *peak*) differs by 1% from the maximum. This line is constructed using a polynomial fit built from the CSG data (similar as in Appendix E). The dots represent the parameter study results with the UMG model. They are colored

yellow if their value is within 1% of the value of the optimal design (MCNP evaluation) and black if they differ >1% from this reference. Except for the results at tube length < 125 mm, the results match perfectly. We conclude that the UMG results confirm the results of the CSG results. A wide range in the parameters gives a design that only differs 1% of the optimal design.

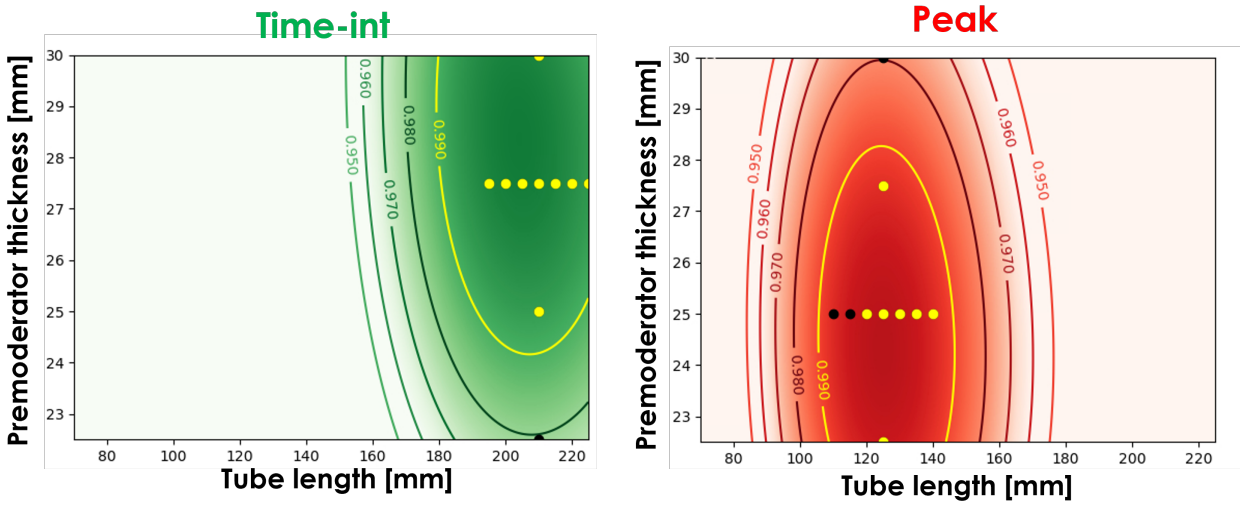

Figure 22. Visualization of the broad optimum for the tube moderator dimensions. The flooded contour and contour lines are constructed using a polynomial fit built from the CSG data (and relative to the maximal value from this model). The dots represent MCNP results from the parameter study with the UMG model. They are colored yellow if their value is within 1% of the value of the optimal design (MCNP evaluation) and black if they differ >1% from this reference.

4.2.2 Pareto set optimization with four parameters

A Pareto set optimization was performed for the tube moderator. The objective function is the following:

$$Objective = w_1 \frac{peak}{2.25 \cdot 10^{-4}} + w_2 \frac{tint}{6.0 \cdot 10^{-8}}, \tag{4}$$

with *peak* and *tint* the brightness metrics obtained from three point detector tallies. In total 113 MCNP evaluations were performed to obtain convergence. The brightness metrics for all 113 evaluated designs are shown in Figure 23. More data plots can be found in Appendix G. The Pareto front is indicated in orange. The green and red crosses indicate the designs that are within 1% of the maximal *tint* and *peak* respectively. Again, it is clear that many designs have brightness metrics within 1% of the optimized designs, indicating a broad optimum with a low sensitivity to many of the design parameters.

Table 10 lists the final designs for the six sets. Sets 5 and 6 resulted in the same design. The chosen parameters are very similar to the those chosen with the CSG model (Table 9). A noteworthy difference is the longer tube length for the *peak*-optimized design: 136 mm instead of 125 mm. The beryllium radius is chosen large for the *tint*-optimized design and small for the *peak*-optimized design.

Table 10. Results of final Pareto set run for lower moderator

Set	w_1	w_2	T_{len} **[mm]**	R_{pm} **[mm]**	R_{Be} **[mm]**	X_{pos} **[mm]**	*peak* **[10^{-4}cm^{-2}s^{-1}]**	*tint* **[10^{-8}cm^{-2}]**
1	0	1	214	27.4	220	8	1.96826	5.8307
2	0.25	0.75	198	27.8	220	7	2.02734	5.7935
3	0.5	0.5	175	27.1	200	0	2.12036	5.621
4	0.7	0.3	161	25.2	184	2	2.16852	5.4764
5	0.85	0.15	136	25	180	3	2.21668	5.1002
6	1	0	136	25	180	3	2.21668	5.1002

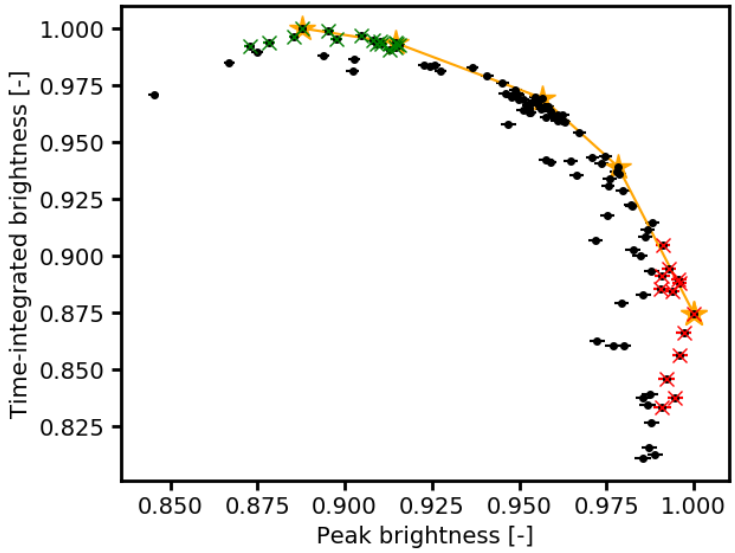

Figure 23. Relative *tint* and *peak* for all MCNP runs for the final lower moderator optimization. The orange line indicates the Pareto front. The green and red crosses indicate the designs that are within 1% of the maximal *tint* and *peak* respectively.

5 CONCLUSIONS

We have selected three optimal designs for both the upper (cylindrical) moderator and the lower (tube) moderator. These designs have been optimized for maximal peak brightness, maximal time-integrated brightness, and an equally weighted combination of both.

The optimization was performed with a new optimization workflow using Dakota, which contains state-of-the-art optimization algorithms. For each design selected by the algorithm in Dakota, a new geometry was created in Creo and SpaceClaim and then exported to Attila4MC to construct an unstructured mesh geometry for MCNP. The brightness metrics were evaluated using an MCNP transport calculation for each design. This process is fully automated and thus requires no user intervention.

Using the Pareto set and EGO optimization methods in Dakota, a series of optimal designs (balancing two objectives) has been obtained with ≈40-110 MCNP evaluations. This is much faster than a typical scoping study for four to six design parameters, which easily requires several hundreds or even thousands of design evaluations.

The most sensitive parameters for the cylindrical moderator are the moderator radius and the premoderator thicknesses (radial, top and bottom). A 5-mm change in any of these parameters results in 1% decrease in brightness around the optimum. The position of the moderator and the beryllium radius are less sensitive: a 20 mm change from the optimal value is needed to decrease the brightness by 1%. The design optimized for time-integrated brightness has a larger radius (62 mm) than the design optimized for peak brightness (40 mm). The same trend is observed for the other parameters: a larger value is selected for a maximized time-integrated brightness.

The wall thicknesses have a large influence on the optimal design parameters and the value of the brightness metrics. For the cylindrical moderator, we compared the results with varying wall thicknesses (based on structural analyses) to those with a fixed thickness. A larger wall thickness, which is required for larger vessels, penalizes the time-integrated brightness: the optimal value decreased by 13% at a drastic reduction of the optimal hydrogen radius from 80 mm to 62 mm compared to a simulation with fixed wall thicknesses.

The tube length is the most crucial design parameter for the tube moderator. A length of 210 mm and 125 mm is selected for the optimal designs that maximize respectively the time-integrated and peak brightness. The premoderator thickness and the beryllium radius are chosen larger the for the design optimized for time-integrated brightness. The designs have a broad optimum and thus have a low sensitivity to the design parameters: a large range of parameters gives a brightness metric close to the optimal value.

6 REFERENCES

[1] B. Adams *et al.*, "Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 6.14 user's manual," Tech. Rep. SAND2021-5822, Sandia National Laboratories, May 2021.

[2] C. Werner *et al.*, *MCNP User's Manual, Code Version 6.2, Los Alamos National Laboratory, LA-UR-17-29981*. 2017.

[3] L. Zavorka, K. Ghoos, J. Risner, and I. Remec, "An unstructured mesh based neutronics optimization workflow," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1052, p. 168252, 2023.

[4] J. Risner, "Optimization studies for the second target station moderator/reflector assembly using an MCNP CSG model," Tech. Rep. S03120000-TRT10000, Oak Ridge National Laboratory, February 2023.

[5] " Creo 7.0.6.0." https://www.ptc.com/en/products/creo, 2022.

[6] "SpaceClaim 2021 R1." https://www.ansys.com/products/3d-design/ansys-spaceclaim, 2020.

[7] *Attila4MC 10.2 Overview of Core Functions, Silver Fir Software, Inc., Gig Harbor, WA, USA, SFSW-UR-2020-OCF102*. 2020.

[8] B. Adams *et al.*, "Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 6.14 theory manual," Tech. Rep. SAND2021-5822, Sandia National Laboratories, May 2021.

[9] B. Adams *et al.*, "Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 6.14 reference manual," Tech. Rep. SAND2021-5822, Sandia National Laboratories, May 2021.

[10] J. Janney, "Moderator reflector assembly thickness vs. size curves," Tech. Rep. S03040100-TR0001-R00, Oak Ridge National Laboratory, October 2021.

APPENDIX A. COMPUTER HARDWARE AND SOFTWARE

APPENDIX A. COMPUTER HARDWARE AND SOFTWARE

Two executables of MCNP have been used on STS Linux computational cluster Saturn. The first one mcnp/dagmc6.2-mtNEW is outdated and does not contain the point detector bug fix (2021). The second one is mcnp/mcnp6.2mod_20220105 is the newest MCNP executable at the time of the analysis and does contain the point detector bug fix. The brightness metrics differ ≈3% between the two executables.

Dakota version 6.14 has been used on Saturn (module dakota/6.14).

APPENDIX B. LOCATION OF COMPUTATIONAL INPUT AND OUTPUT FILES

APPENDIX B. LOCATION OF COMPUTATIONAL INPUT AND OUTPUT FILES

The input and output files are located on saturn in the sts-archive in directory
`sts_archive/S.03.04_Mod/Optimization_2022/`.

APPENDIX C. GEOMETRY PLOTS OF OPTIMAL CONFIGURATIONS OF THE UPPER MODERATOR

APPENDIX C. GEOMETRY PLOTS OF OPTIMAL CONFIGURATIONS OF THE UPPER MODERATOR

For the optimized designs of Table 2, a vertical and horizontal cut through the geometry is respectively shown in Figures 24 and 25. In Figures 26 and 27 compares the *peak*- and *tint*-optimized designs, showing a vertical cut through the geometry, for the models with fixed and variable wall thicknesses. The *tint* and *peak* FOM in these figures are in units of neutrons/cm^2/proton and neutrons/cm^2/proton/shake, with a shake equal to 10^{-8}s.

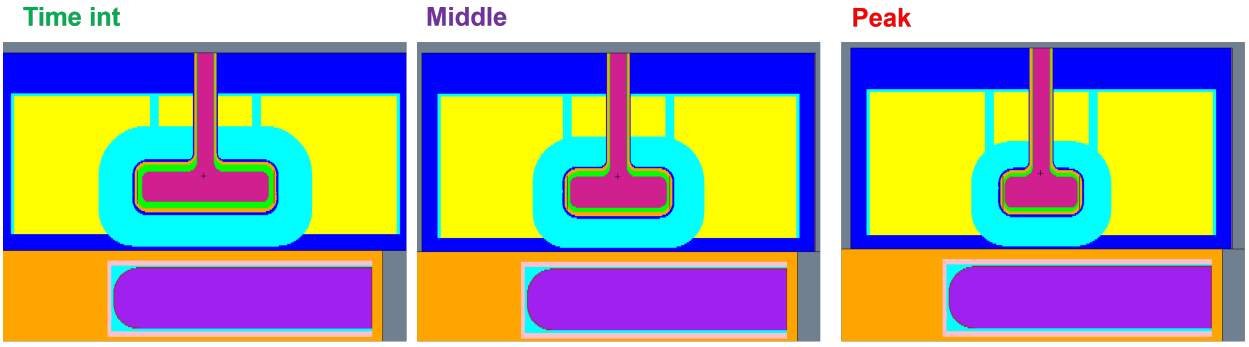

	R_{mod} [mm]	R_{pm} [mm]	H_{pm} [mm]	R_{Be} [mm]	X_{pos} [mm]	Peak Brightness [e-12]	Time-integrated brightness [e-8]
Time int	62	33	30	187	2	1.218 (80%)	2.899 (100%)
Peak	35	26	28	165	1	1.524 (100%)	2.276 (79%)
Middle	47	29	29	173	1	1.435 (94%)	2.734 (94%)

Figure 24. Vertical cut through the geometry of the optimized upper moderator designs.

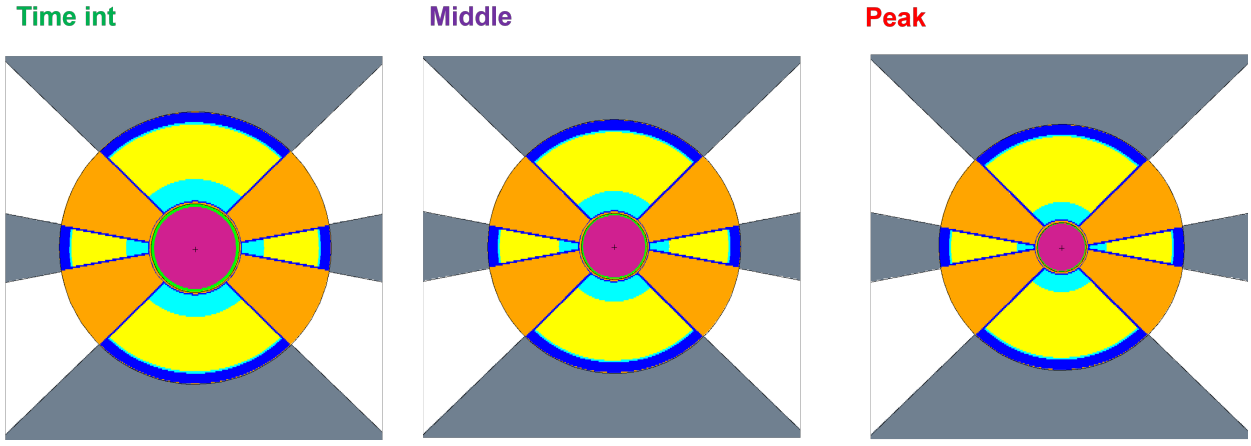

Figure 25. Horizontal cut through the geometry of the optimized upper moderator designs.

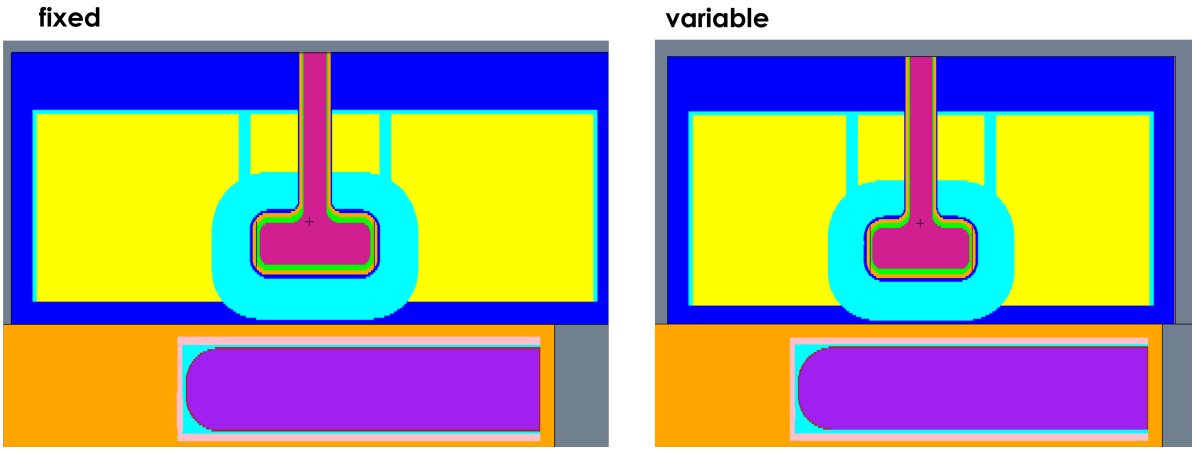

	R_{mod} [mm]	R_{pm} [mm]	H_{pm} [mm]	R_{Be} [mm]	X_{pos} [mm]	Peak Brightness [e-12]	Time-integrated brightness [e-8]
Fixed	40	27	29	200	4	1.494	2.484
Variable	35	26	28	165	1	1.524 (+2%)	2.276 (-9%)

Figure 26. Vertical cut through the geometry of the *peak*-optimized design with fixed and variable wall thicknesses.

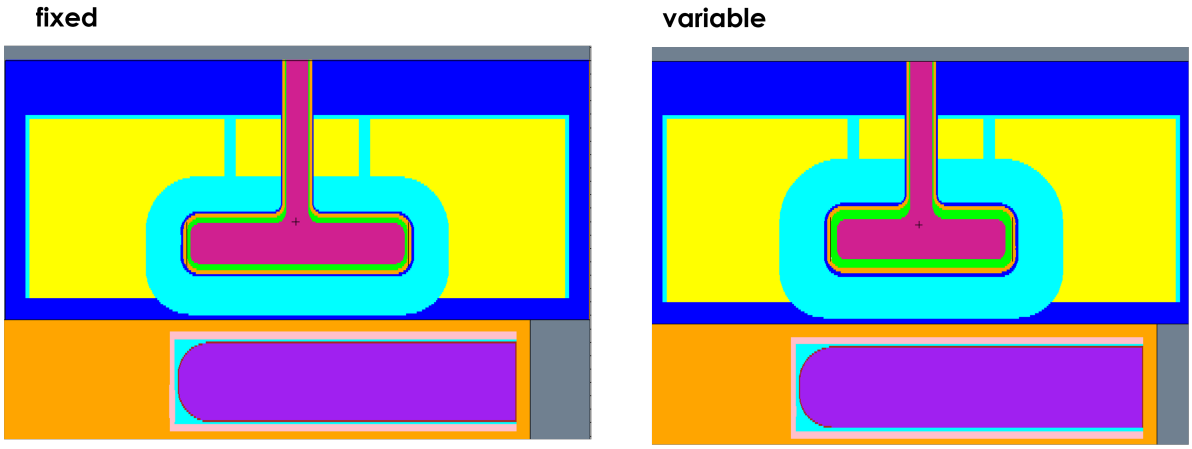

	R_{mod} [mm]	R_{pm} [mm]	H_{pm} [mm]	R_{Be} [mm]	X_{pos} [mm]	Peak Brightness [e-12]	Time-integrated brightness [e-8]
Fixed	80	26	28	200	1	1.260	3.312
Variable	62	33	30	187	2	1.218 (-4%)	2.899 (-13%)

Figure 27. Vertical cut through the geometry of the *tint*-optimized design with fixed and variable wall thicknesses.

APPENDIX D. DATA PLOTS FOR UPPER MODERATOR

APPENDIX D. DATA PLOTS FOR UPPER MODERATOR

Figures 28, 29, 30, 31 and 32 present the data from 196 MCNP runs of the upper moderator with different parameter combinations. In each graph, all data is plotted against one parameter, regardless of the value of the other parameters. These 196 runs come from several (sometimes incomplete) Dakota runs, both optimization and parameter study. The *tint* and *peak* FOM in these figures are in units of neutrons/cm^2/proton and neutrons/cm^2/proton/shake, with a shake equal to 10^{-8}s.

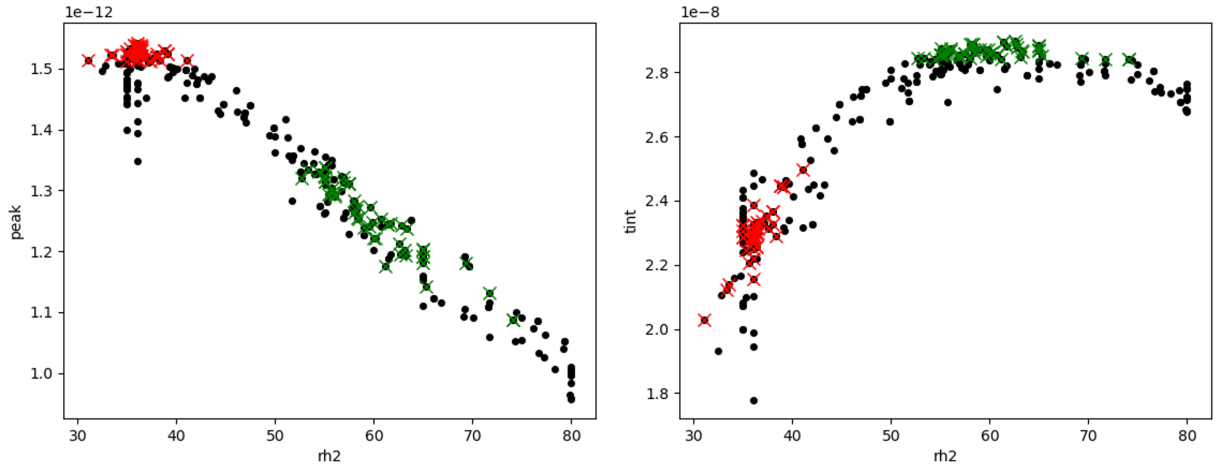

Figure 28. Brightness metrics *peak* and *tint* of 196 MCNP runs as a function of the moderator radius. The green and red crosses are within 2% of the maximal *tint* and *peak* respectively.

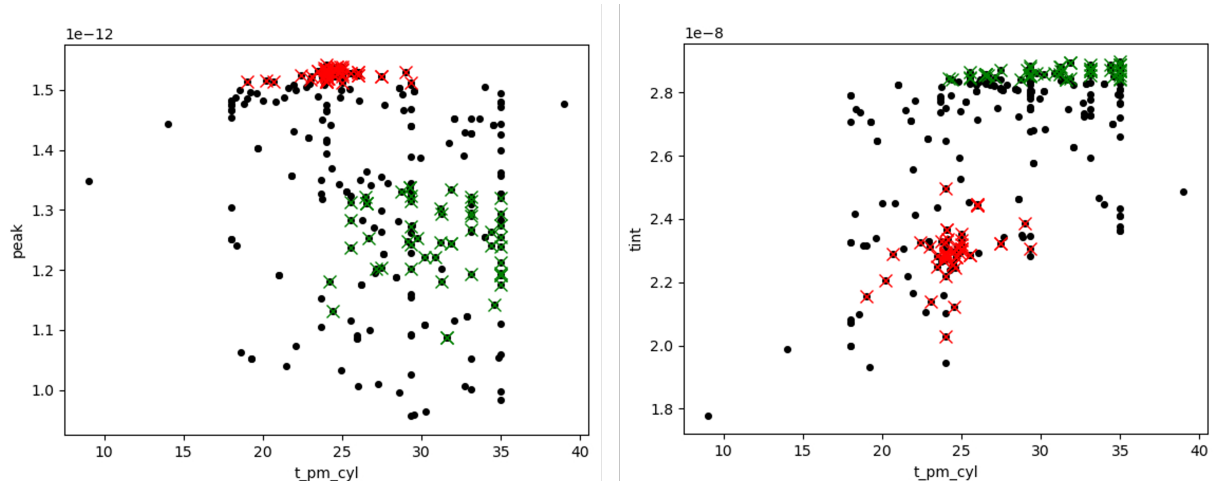

Figure 29. Brightness metrics *peak* and *tint* of 196 MCNP runs as a function of the premoderator thickness radially and on top. The green and red crosses are within 2% of the maximal *tint* and *peak* respectively.

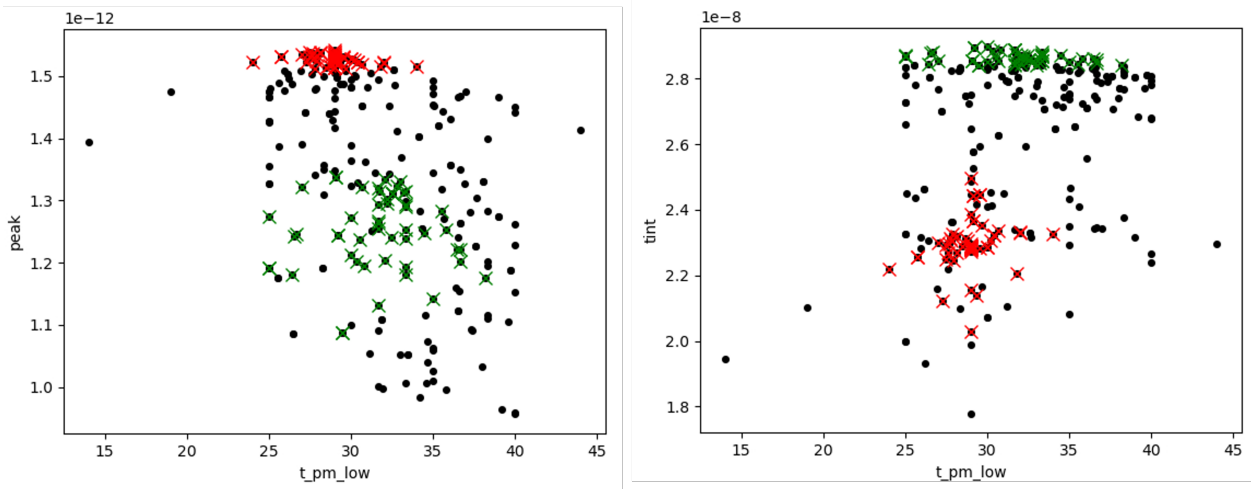

Figure 30. Brightness metrics *peak* and *tint* of 196 MCNP runs as a function of the premoderator thickness at the bottom. The green and red crosses are within 2% of the maximal *tint* and *peak* respectively.

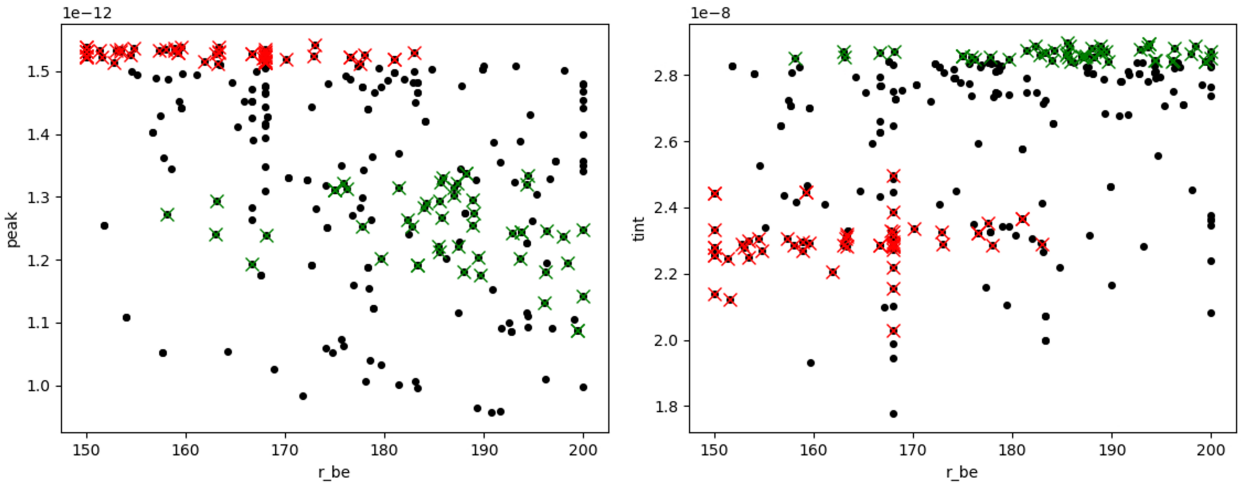

Figure 31. Brightness metrics *peak* and *tint* of 196 MCNP runs as a function of the beryllium radius. The green and red crosses are within 2% of the maximal *tint* and *peak* respectively.

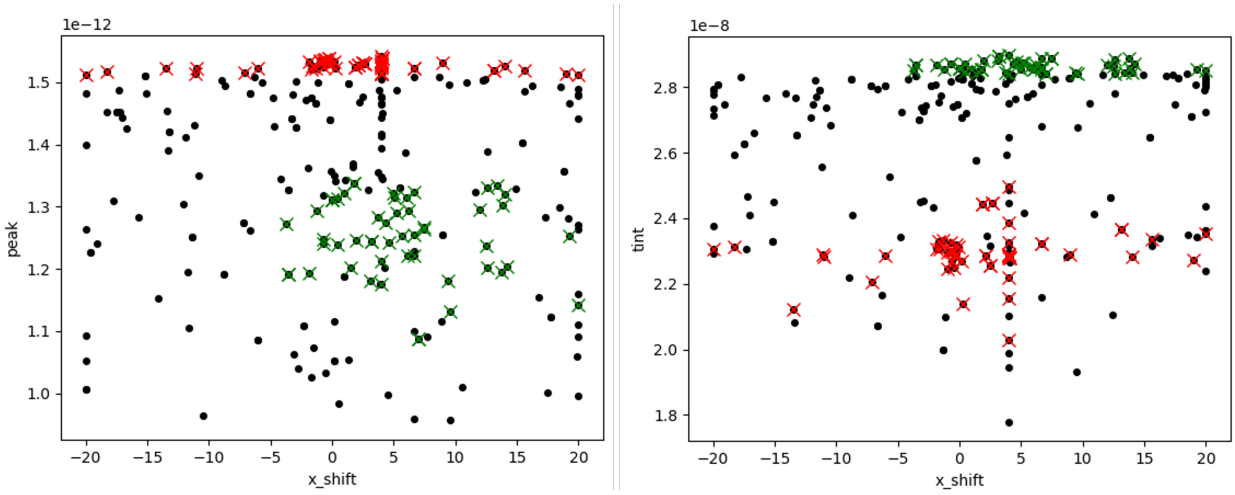

Figure 32. Brightness metrics *peak* and *tint* of 196 MCNP runs as a function of the horizontal moderator position. The green and red crosses are within 2% of the maximal *tint* and *peak* respectively.

APPENDIX E. POST-PROCESSING OF SCOPING STUDY WITH CSG MODEL AND COMPARISON WITH UMG

APPENDIX E. POST-PROCESSING OF SCOPING STUDY WITH CSG MODEL AND COMPARISON WITH UMG

An extensive scoping study was performed using a CSG model and varying the moderator radius R_{mod}, the bottom premoderator thickness H_{pm}, and the beryllium radius R_{Be} [4]. The tally results from that study were used to construct a polynomial fit with Dakota. A subset of the results was selected to obtain a polynomial fit that matches the results well (based on χ^2-value and sensitivity plots around the optima). Only the results with $R_{mod} \geq 37.5$ mm and $peak > 0.7$ (relative) have been used. This fit captures the behavior very well around the *tint*-optimized design, but it does not capture the behavior close to the *peak*-optimized design, as it is close to the domain boundaries and as such fewer samples are available. This fit is sufficiently good as a tool to compare the results and see trends, but it was not used to select an optimum design.

Figures 33 and 34 show contour plots of respectively *tint* (green color) and *peak* (red color) obtained with the polynomial fit based on the CSG data. The dots show the locations in the domain where an MCNP evaluation has been performed with the CSG model. The dot is colored green or red if respectively *tint* or *peak* is within 2% of the maximal value. Comparing the colored dots with the contour line of 0.98 shows an excellent match. These contour plots provide a visual aid to gain more insight in the influence of the parameters on the brightness metrics.

The dark shade in Figures 33 and 34 indicates the location of the optimum. For both *tint* and *peak*, the choice of the moderator radius R_{mod} is crucial. For a fixed H_{pm} and R_{Be}, *tint* and *peak* vary ≥20% within the chosen bounds (35 mm $< R_{mod} <$ 100 mm). This variation is much smaller for H_{pm} and R_{Be}. The contours are very elongated along H_{pm}, indicating a small sensitivity within the chosen bounds (20 mm $< H_{pm} <$ 35 mm). The effect of R_{Be} is seen from comparing the different plots (for the four different values of R_{Be}). Figure 33 shows a much larger and darker (thus higher value) optimum for *tint* for the largest R_{Be} = 18.25 cm compared to the smallest R_{Be} = 15.25 cm. The opposite is true for *peak* in Figure 34, in which darker shades (higher values) are observed for the smallest R_{Be} = 15.25 cm. In general, we see that many designs give a high value for *peak* and/or *tint*. A small difference in the dimensions of the moderator will only have a small effect on the brightness. A broad range of parameters gives results close to the optimum.

Figure 35 compares the Pareto front with optimal designs obtained with the UMG model (dashed lines and stars) and the CSG model (full lines and diamonds). The lines are obtained with the polynomial fits, while the stars and the diamonds are the three chosen optimal designs for respectively the UMG and CSG model evaluated with MCNP. The brightness metrics are shown relative to the maximal *peak* and *tint* obtained with the MCNP evaluations of the CSG model. While there are slight differences between the UMG and CSG results, the overall behavior matches very well. Looking at the lines with optimal choices of R_{mod}, H_{pm} and R_{Be}, both models give very similar optimal designs. This comparison provides a high degree of confidence in the validity of the results.

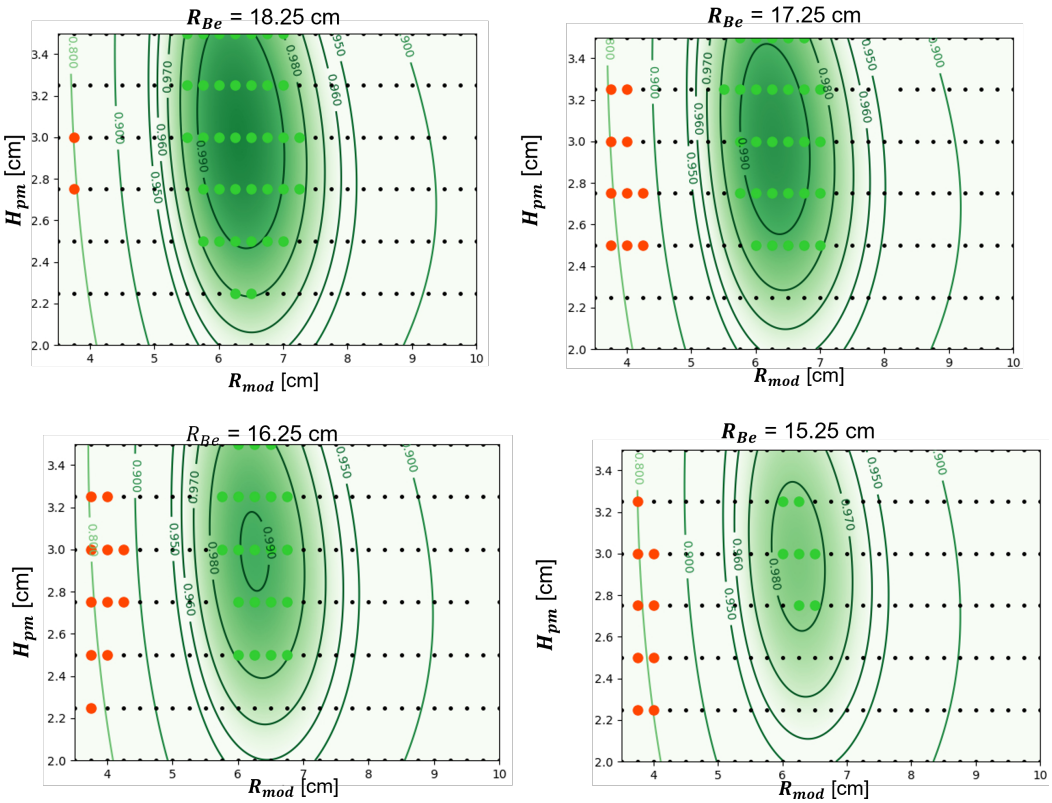

Figure 33. Contour plots of the polynomial fit of *tint* based on MCNP data with the CSG model. The dots show the locations in the domain where an MCNP evaluation has been performed with the CSG model. The dot is colored green or red if respectively *tint* or *peak* is within 2% of the maximal value

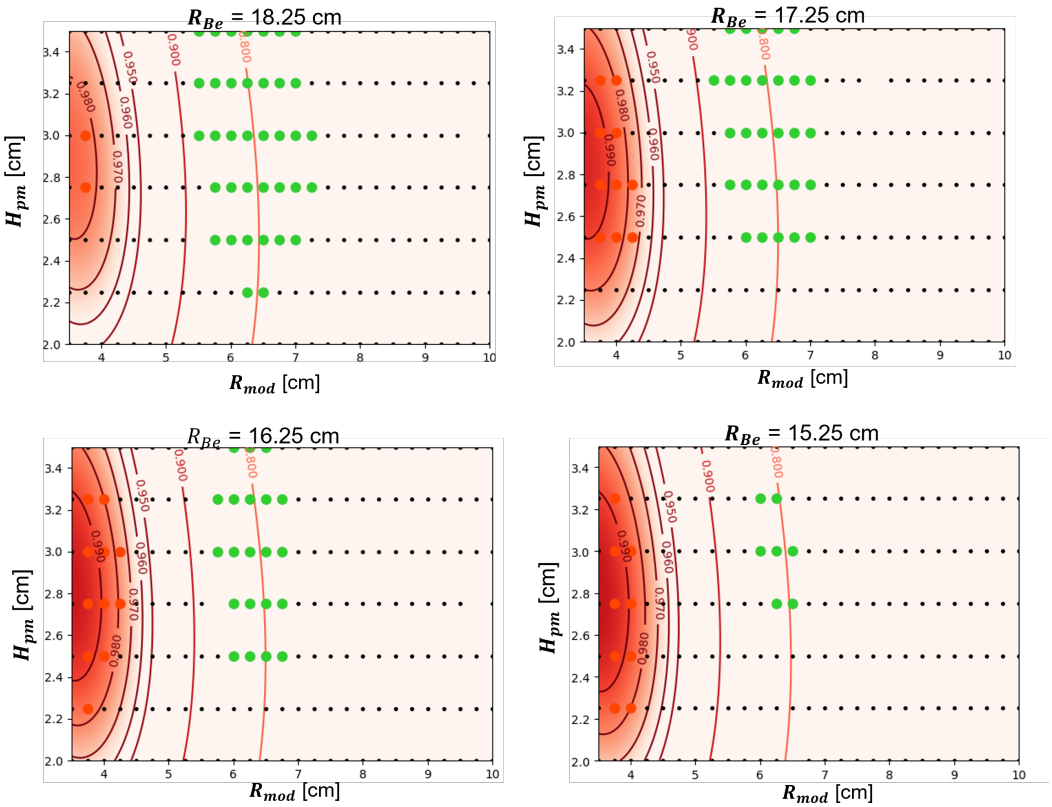

Figure 34. Contour plots of the polynomial fit of *peak* based on MCNP data with the CSG model. The dots show the locations in the domain where an MCNP evaluation has been performed with the CSG model. The dot is colored green or red if respectively *tint* or *peak* is within 2% of the maximal value.

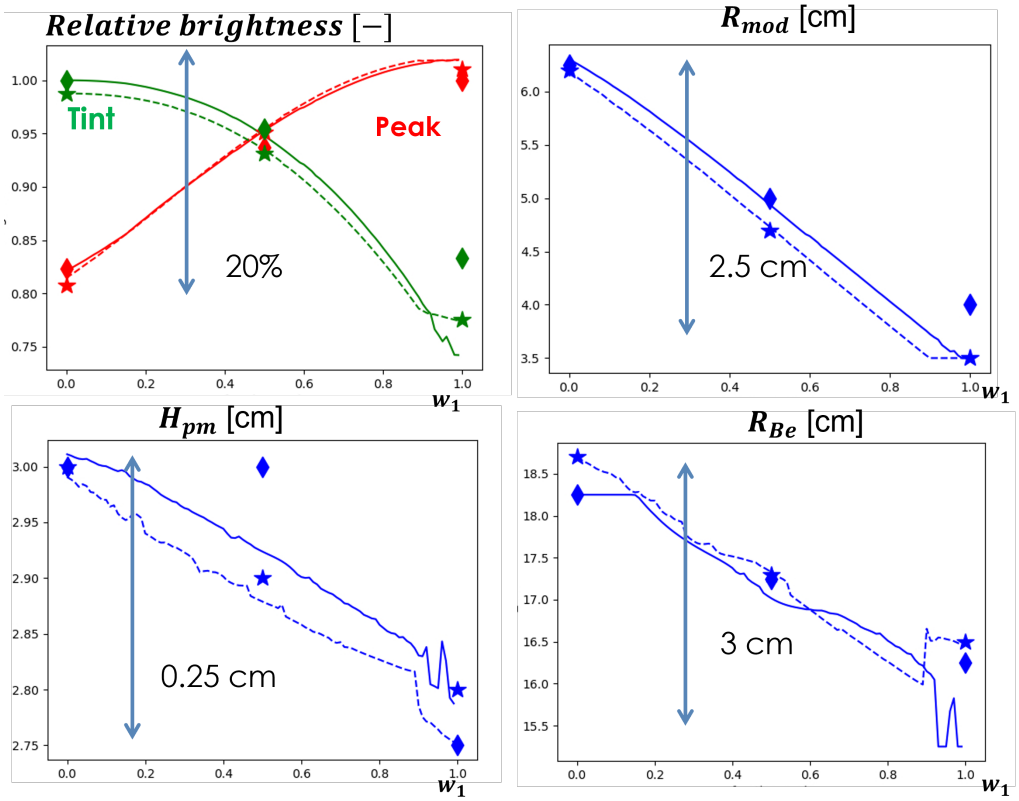

Figure 35. Objective functions and the parameters along the Pareto front obtained with the UMG model (dashed lines and stars) and the CSG model (full lines and diamonds).

APPENDIX F. UNCERTAINTY FROM THE NUMERICAL PROCESS

APPENDIX F. UNCERTAINTY FROM THE NUMERICAL PROCESS

In this appendix, the effect of statistics on the solution is evaluated by separating the contribution from the optimization algorithm and the effect of the MCNP precision. The optimization algorithm is a statistical method because its final result is dependent on the initial set of runs (statistically sampled in the latin hyper cube) to build the surrogate model for the EGO algorithm. This test is performed using the surrogate model of the cylindrical moderator, see Section 4.1.1.

We examine the uncertainty by analyzing the probability distribution of solutions obtained by repeating the optimization algorithm 1000 times with different seeds. The optimization is run for the polynomial cubic fit model of the cylindrical moderator. This approximate model has two advantages: it runs much faster than MCNP and it is deterministic (and thus provides a solution without relative error). A histogram of the resulting optimal R_{mod} is shown in the left graph of Figure 36. For each set (optimized for a different weighted sum of *tint* and *peak* as defined in the Section 4.1), the black lines represent a gaussian fit for the distribution. Except for Sets 5 and 6, in which the lower bound of R_{mod}= 35 mm interferes, the gaussian distributions approach the results well. In the right graph of Figure 36, the mean and standard deviation σ (errorbar) of the gaussian distributions of each set are shown as a function of the weight w_1. The grey errorbars indicate the 3σ error bars. The squares indicate the most extreme values obtained in 1000 runs. In Figure 37, these graphs are shown for all parameters and the objective functions. Of course, these graphs look very similar to those in Figure 7. The valuable addition is the error bars, which provide the expected accuracy from the optimization algorithm.

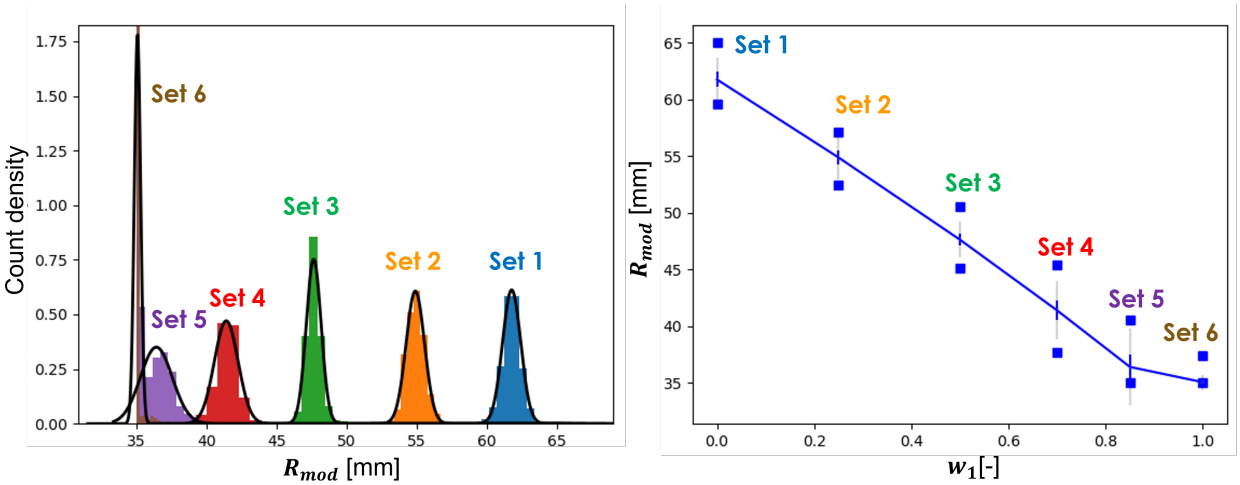

Figure 36. Left: Histogram of R_{mod} values (and gaussian fits) obtained by running the Pareto set optimization algorithm 1000 times on the cubic polynomial surrogate model. Right: Mean and standard deviation σ (error bar) of the R_{mod} values as a function of w_1. The grey errorbars indicate the 3σ error bars. The squares indicate the most extreme values obtained in 1000 runs.

The solutions obtained with MCNP are expected to have a higher statistical uncertainty because the brightness metrics themselves have a statistical precision. To evaluate the added effect of this uncertainty,

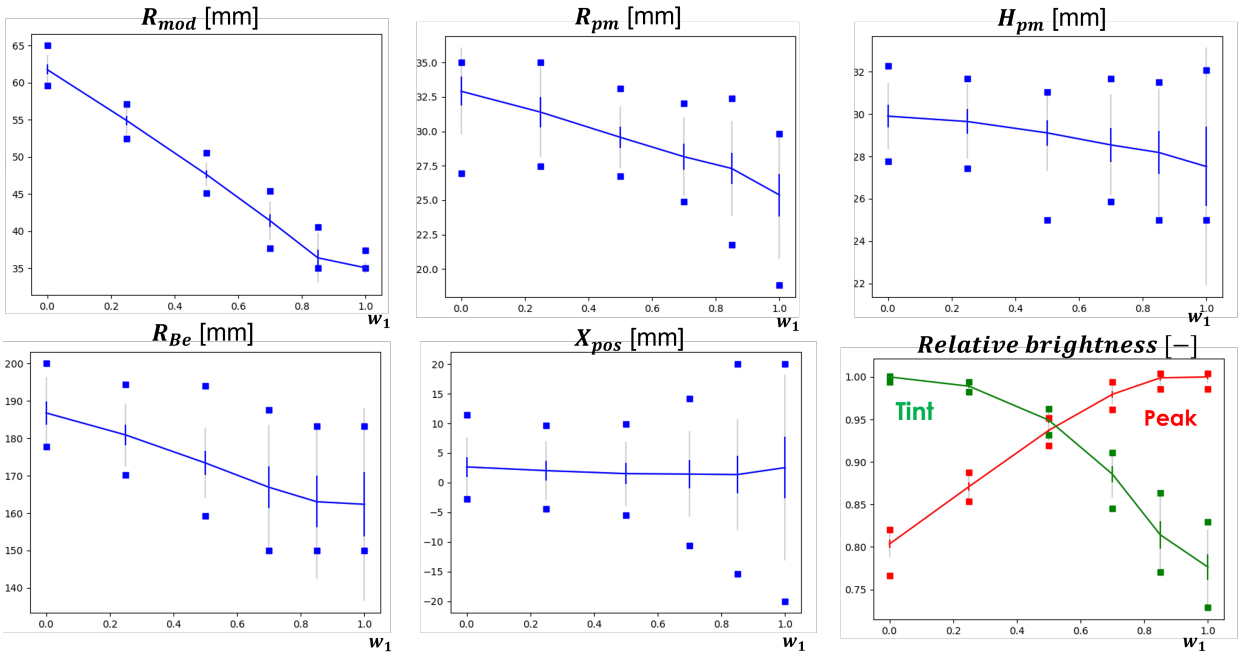

Figure 37. Mean and standard deviation σ (error bar) of all parameters and objectives as a function of w_1. The grey errorbars indicate the 3σ error bars. The squares indicate the most extreme values obtained in 1000 runs.

the analysis is repeated for a noisy surrogate model. To obtain a solution with statistical noise, as is the case with MCNP, a random normally distributed number with the appropriate standard deviation is simply added to the deterministic solution for each evaluation of the polynomial model. Finally, the analysis is repeated once more with a larger precision, which corresponds to 5x the MCNP precision. In Table 11, a summary is given for the range of standard deviations obtained with both noise levels. The standard deviations that have been added to the polynomial fit (1x noise) correspond to the estimate of the statistical uncertainty that we calculated for most of the MCNP simulations.

The variation in optimal parameters obtained with the EGO algorithm is correlated with the sensitivity of each parameter. For R_{mod}, R_{pm} and H_{pm}, the standard deviation is limited to 1.5 mm. This is significantly less than 5 mm, which is the change needed to reduce the optimum by 1%. With noise included, this value increases a little, but not drastically. Even with 5x the noise, the standard deviation remains significantly lower than the the parameter change that causes 1% metric change. The same is true for R_{Be} and X_{pos} for which the sensitivity is smaller. This results in a larger variation in the optimal configurations, but lower than the sensitivity. The maximal values in the objective functions *tint* and *peak* are also estimated accurately. Especially for single-objective runs (set 1 and 6), the standard deviation of respectively *tint* (set 1) and *peak* (set 6) remain well below 1%. The increase in the standard deviation when noise is added, is very limited. This demonstrates that the algorithm is working properly for this problem, even with statistical noise on the simulation results.

Table 11. Variation of parameters and objectives in the optimization process, with and without added noise, compared to the sensitivity of the parameter.

Parameter	Standard deviation from optimization no noise [mm]	Standard deviation from optimization 1x noise [mm]	Standard deviation from optimization 5x noise [mm]	Sensitivity around optimum [mm]
R_{mod}	0.5-1 mm	0.8-1.5mm	1.5-2.5 mm	≈5 mm for 1%
R_{pm}	0.5-1.5 mm	1.2-1.6mm	2.0-3.0 mm	≈5 mm for 1%
H_{pm}	0.5-1.5 mm	0.6-1.6 mm	1.2-2.5 mm	≈5 mm for 1%
X_{pos}	2-5 mm	2-6 mm	4-10 mm	≈20 mm for 1%
R_{Be}	3-8 mm	4-9 mm	7-13 mm	≈20 mm for 1%
Objective	**Standard deviation from optimization no noise**	**Standard deviation from optimization 1x noise**	**Standard deviation from optimization 5x noise**	**MCNP precision (relative error) 1x noise**
tint	0.1-1.5%	0.1-2%	0.3-3.5%	0.06%
peak	0.3-0.5%	0.5-0.7%	1.0-1.8%	0.3%

APPENDIX G. DATA PLOTS FOR LOWER MODERATOR

APPENDIX G. DATA PLOTS FOR LOWER MODERATOR

Figures 38, 39 , 40 and 41 show the data from 113 MCNP runs of the lower moderator with different parameter combinations. In each graph, all data is plotted against one parameter, regardless the value of the other parameters. These 113 runs come from several (sometimes restarted) optimization runs. The *tint* and *peak* FOM in these figures are in units of neutrons/cm^2/proton and neutrons/cm^2/proton/shake, with a shake equal to 10^{-8}s.

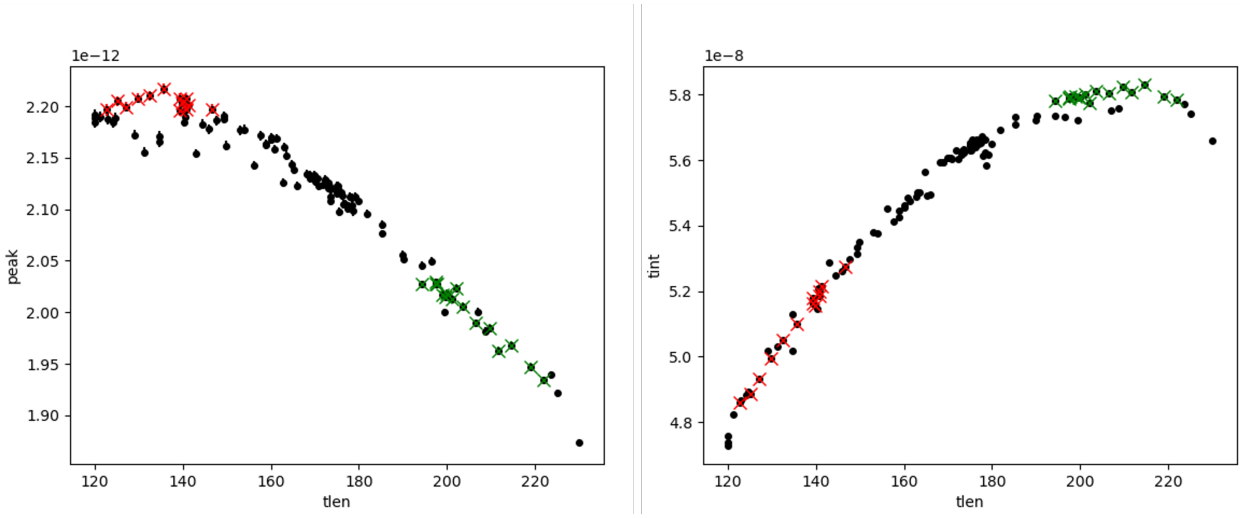

Figure 38. Brightness metrics *peak* and *tint* of all 113 MCNP runs as a function of the tube length. The green and red crosses are within 1% of the maximal *tint* and *peak* respectively.

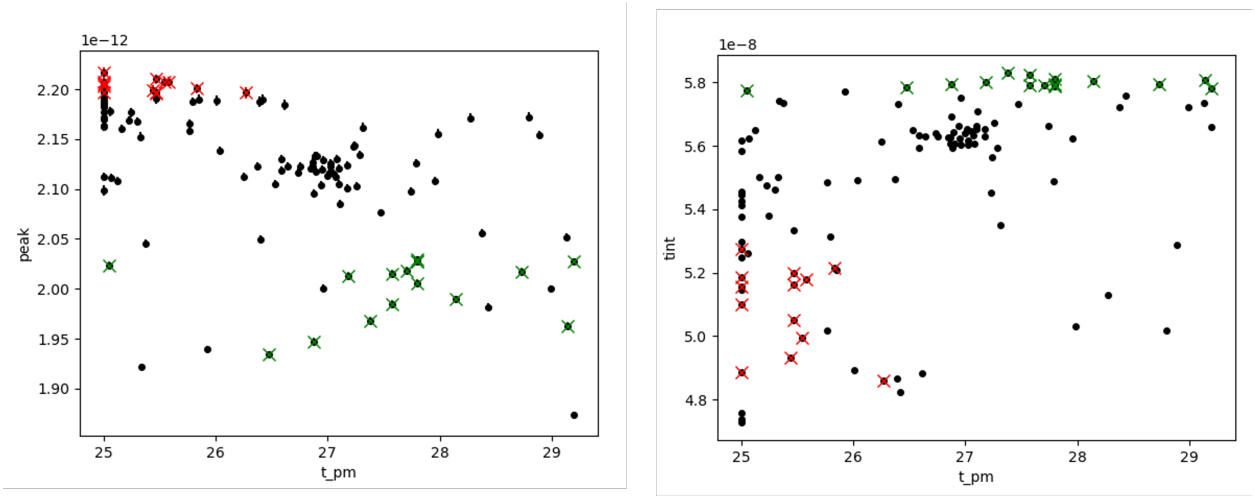

Figure 39. Brightness metrics *peak* and *tint* of all 113 MCNP runs as a function of the premoderator thickness. The green and red crosses are within 1% of the maximal *tint* and *peak* respectively.

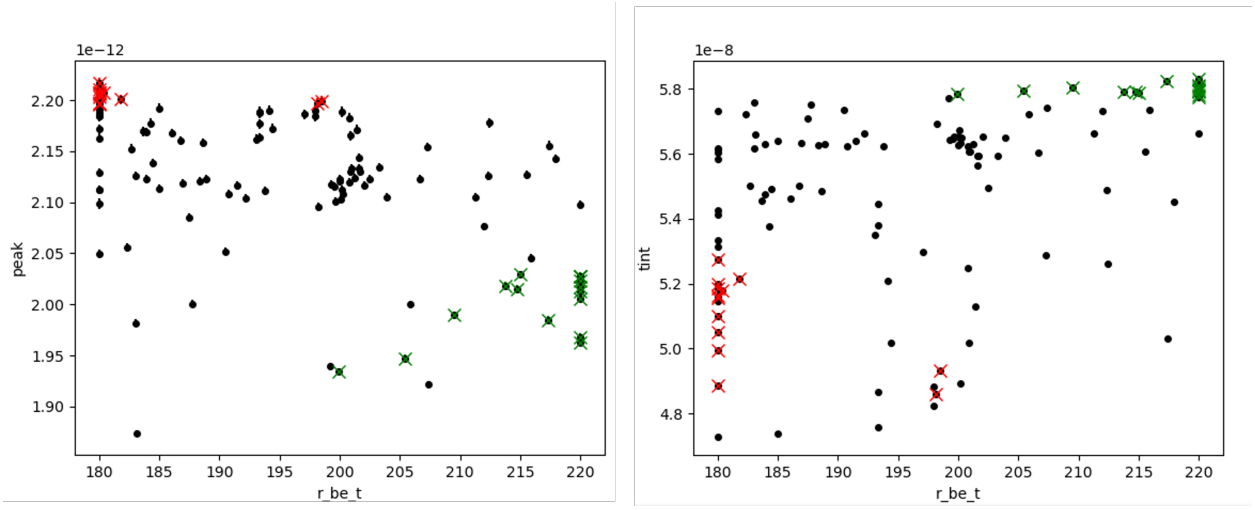

Figure 40. Brightness metrics *peak* and *tint* of all 113 MCNP runs as a function of the beryllium radius. The green and red crosses are within 1% of the maximal *tint* and *peak* respectively.

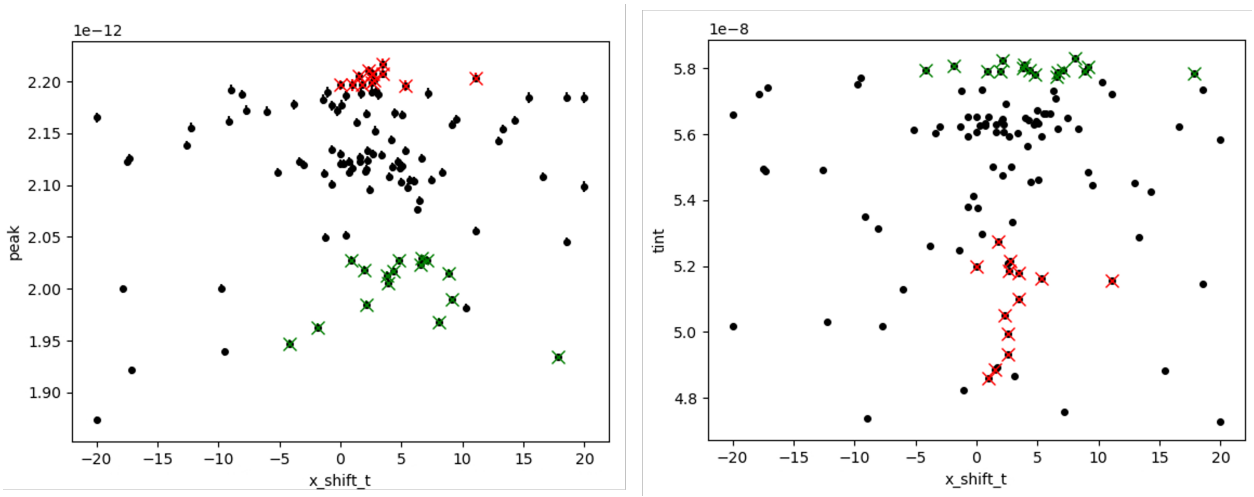

Figure 41. Brightness metrics *peak* and *tint* of all 113 MCNP runs as a function of the horizontal position of the moderator. The green and red crosses are within 1% of the maximal *tint* and *peak* respectively.

APPENDIX H. EXAMPLE INPUT FILES FOR DAKOTA

APPENDIX H. EXAMPLE INPUT FILES FOR DAKOTA

This appendix details two input files for Dakota. More details can be found in the Dakota Reference Manual [9].

The first input file shows an example of a Pareto set run. The environment block defines the name of a tabular output file and points to the first method. In the `pareto_set` method block, the weigths for each set are specified. The EGO algorithm is used to optimize each set. In the EGO algorithm, you can choose to import previous samples (commented out in this example with the # sign). Additional keywords to control the convergence of the algorithm can also be specified. In this example, six variables are specified with their upper and lower bounds. Two objective functions need to be returned. In the interface, a driver script `'run4b'` is called with bash. This script reads the parameters from `'params.in'`, runs the simulation code, and provides the results in `'results.out'`. If a failure is detected, two zeros are returned.

The second input file performs a centered parameter study on a polynomial cubic surrogate model. Instead of providing an interface with the simulation, the model is based on the data provided in the txt-file `'UM_ego_4par_2obj.txt'`. The surrogate model is exported in an algebraic file for inspection later. In addition, several metrics are specified that we use as a first evaluation of the quality of the surrogate model.

```
environment
  top_method_pointer = 'PS'
  tabular_data
    tabular_data_file = 'UM_pareto_set.dat'

method
  id_method = 'PS'
  pareto_set
    method_pointer = 'EGO'
    weight_sets = 0.00000e+00 1.72473e+07
                  8.20210e+10 1.29355e+07
                  1.64042e+11 8.62366e+06
                  2.29659e+11 5.17420e+06
                  2.78871e+11 2.58710e+06
                  3.28084e+11 0.00000e+00

method
  id_method = 'EGO'
  efficient_global
    max_iterations = 10
    # import_build_points_file = 'UM_ego_5par_2obj_nps5e6.txt
    #   freeform
    convergence_tolerance = 1e-5
    # x_conv_tol = 0.1

model
  single
```

```
variables
  continuous_design = 6
    lower_bounds   35       18          25          25         150       -20
    upper_bounds   80       35          40          40         200        20
    descriptors   'rh2'  't_pm_cyl' 't_pm_low'  't_pm_top'  'r_be' 'x_shift'

interface
  fork
    analysis_drivers = 'bash run4b'
      parameters_file = 'params.in'
      results_file = 'results.out'
   failure_capture recover 0 0

responses
  objective_functions = 2
   sense = 'maximization'
  descriptors = 'peak' 'tint'
  no_gradients
  no_hessians
```

```
environment
  tabular_data
    tabular_data_file = 'surrogate_centered_param.dat'

method
   centered_parameter_study
    step_vector = 1 1 1 1
    steps_per_variable = 2 2 2 2

model
  surrogate global
   polynomial cubic
   import_build_points_file = 'UM_ego_4par_2obj.txt'
     freeform
   export_model format algebraic_file
      filename_prefix = 'poly_cubic'
   metrics = 'root_mean_squared' 'sum_abs' 'mean_abs' 'max_abs' 'rsquared'
      cross_validation folds = 5

variables
  continuous_design = 4
    lower_bounds   35       18          25      170
    upper_bounds   90       25          32      200
    descriptors   'rh2'  't_pm_cyl' 't_pm_low' 'r_be'
```

```
responses
   response_functions = 2
    descriptors = 'peak' 'tint'
   no_gradients
   no_hessians
```