

Nuclear Fuel Bundle Design Optimization using a Simplex Method

Anders Haulin

Supervisor: Jan Pallon, Ph.D.



LUNDS
UNIVERSITET

Department of Nuclear Physics, Faculty of Engineering
Lund University
August 2014

Abstract

A strategy for nuclear fuel bundle design optimization in boiling water reactors(BWRs), using a simplex method, is presented and examined. The objective of the study is to see how an automatic optimization scheme can create nuclear designs for use in BWR fuel bundles and how the scheme can be improved. Challenges with applying linear simplex optimization to non-linear bundle physics lie in approximating bundle characteristic parameters accurately while keeping calculation times down. Several areas for improvements are found in the calculations and investigation of these find that many can improve accuracy only at the cost of deteriorating calculation time, whereas others such as utilizing symmetries can improve accuracy at no cost. The important R-factor is examined in detail and its implementation in a simplex optimization is studied and improvements to it suggested. It is found that automatic fuel bundle optimization using a simplex method can generate feasible designs and in some cases perform better than layouts made by nuclear designers. With simplex targets and constraints too narrowly specified they will however underperform in areas other than the target function. The reference design entered into the optimization is also shown to be of great importance to performance. It is finally stressed that the optimization framework does have potential to improve the nuclear design process but needs a user with nuclear design experience to operate well and an improved data-management system to be user-friendly.

Acknowledgements

I want to thank Sven-Birger and Simon for the idea of the project and the necessary assistance to complete it, Sven and the BTF department for great support in a stimulating environment and Mikael for sharp insights in optimization as well as new angles and perspectives on the problems I faced.

¹

¹I also want to thank Bitsy for being supportive and generally awesome.

Nomenclature

a	The difference in k_{∞} at 0 GWd/tU burnup which would result from removing all BA from the design.
b	The burnup at which maximum bundle k_{∞} occurs
F_{int}	Internal power peaking factor: the maximum relative pin power in an axial fuel bundle segment.
k_{∞}	The ratio of the neutron production rate and the neutron loss rate in a specific fuel lattice under reflective (or periodic) boundary conditions. Used as a measure of a specific fuel bundle's contribution to the neutron economy of a reactor.
P	Relative power distribution: the power generated in a chosen axial segment of a fuel rod relative to the average power in that axial segment of the bundle.
R	R-factor: a parameter used in <i>CPR correlations</i> as a measure of bundle <i>dryout</i> sensitivity to changes in bundle power. The R-factor is a function of the fuel bundle relative pin power distribution.
BA	Burnable absorber: an element with a high thermal neutron absorption cross section, typically gadolinium, added to nuclear fuel to alter its characteristics.
bundle	Boiling Water Reactor nuclear fuel bundle. An assembly of fuel rods used to fuel the reactor.
burnup	Nuclear engineering term for the amount of energy that has been extracted out of a nuclear fuel.
BWR	Boiling Water Reactor
correlation	A set of equations correlating reactor parameters such as bundle <i>R-factors</i> , bundle power, axial power shape, bundle coolant flow and inlet sub-cooling into a measure of <i>CPR</i> .
CPLEX	A commercial simplex optimization program used in this project.
CPR	Critical Power Ratio: a reactor safety margin stating how much power can increase before inducing dryout conditions.
DOFACT	A program used in this project to calculate bundle R-factors for a dryout <i>correlation</i> .

dryout	A situation when heat flux from a BWR rod reaches a level where boiling water can not cool sufficiently and fuel rod temperature rises swiftly.
enrichment	The weight fraction of ^{235}U to total uranium within a confinement; in this project normally a fuel rod or an axial segment of a fuel bundle
gap	In a simplex optimization the gap is a measure of the discrepancy between the incumbent best solution's objective function value and estimation of global optimum's. <i>Gap tolerance</i> prescribes how small this gap should be for an optimization to terminate.
lattice code	A coarse-grid finite mesh program that calculates burnup-dependent fuel parameters k and P for bundle designs.
LHGR	Linear Heat Generation Rate: a reactor parameter and constraint describing heat generation per fuel rod length, in this project simplified as a constraint on F_{int} .
OPL	Optimization Programming Language, a proprietary programming language used to create optimization models.
PHOENIX	A licensed 2D <i>lattice code</i> used in this project.
PWR	Pressurized Water Reactor
S-matrix	Sensitivity matrix: a set of data describing the effects from changes in fuel bundle composition on fuel bundle characteristic parameters.
Simplex	Simplex method, a deterministic optimization algorithm used for this project.
tolerance	see <i>gap</i>

Contents

1	Introduction	8
1.1	Nuclear Reactor Theory	8
1.1.1	Nuclear Fission	8
1.1.2	Neutron Economy	10
1.1.3	Boiling Water Reactors	11
1.1.4	Critical Power Ratio	11
1.1.5	Uranium Enrichment and Burnable Absorber	12
1.2	BWR Fuel Bundles	13
1.2.1	Mechanical Design	13
1.2.2	Nuclear Design	14
1.2.3	Fuel Bundle Characteristic Parameters	16
1.3	The case for automatic optimization	19
1.3.1	The problem at hand	19
1.3.2	Optimization strategies	20
1.3.3	The simplex method	21
1.3.4	Other work in the field	21
2	Existing Optimization Strategy	22
2.1	Outline	22
2.2	Constraints and objective	24
3	Method	26
3.1	Research Questions	26
3.2	Project Progression	27
3.2.1	Practical Issues	27
3.2.2	Improving the Approximations	28
3.3	Two Optimization Cases	28
4	Apparatus	29
4.1	S-matrix hardware and software	29
4.1.1	Hardware	29
4.1.2	Software	29
4.2	Simplex optimization hardware and software	29
4.2.1	Hardware	29
4.2.2	Software	29
4.2.3	Equipment upgrade	30

5	Results	31
5.1	Evaluating the framework	31
5.1.1	Two-step approach	31
5.1.2	Symmetry issues	32
5.1.3	Average enrichment implementation	33
5.2	Improving approximations	34
5.2.1	k_{∞}	34
5.2.2	Power distribution	39
5.2.3	R-factor	40
5.3	Optimized layout results	50
5.3.1	BA rod position optimization	50
5.3.2	Uranium enrichment and Gd concentration optimization	52
5.3.3	Performance compared to manual designs	54
6	Discussion	62
6.1	Simplex framework	62
6.1.1	Benefits of simplex fuel bundle optimization	62
6.1.2	Simplex method improvements	63
6.1.3	Feasibility and performance of optimization	65
6.1.4	Recommendations and usage guidelines	66
6.2	Project evaluation	67
6.2.1	Goal attainment	67
6.2.2	Method used	68
6.2.3	Areas for further investigation	68
6.3	Project conclusion	70
7	Bibliography	71
A	Additional Data	74

Chapter 1

Introduction

1.1 Nuclear Reactor Theory

The goal of this thesis has been to study the performance and applicability of automatically optimized nuclear designs for Boiling Water Reactor fuel bundles. To explain the setting of the problem a background in nuclear engineering is provided, more detailed background on nuclear physics can (amongst others) be found in Krane [22] and detailed knowledge on nuclear engineering in Lamarsh & Baratta [23].

1.1.1 Nuclear Fission

The underlying principle of all commercial nuclear power is nuclear fission. Nuclei are made up of nucleons combining through interplay between the strong nuclear force supplying an attracting force between all nucleons and the electromagnetic force contributing a repelling force between protons. The general picture of the potential arising from these forces, although complicated in its details, can be displayed in a binding energy diagram (see Figure 1.1). The diagram in Figure 1.1 shows the binding energy per nucleon in the nuclei from hydrogen to uranium. Omitting the local variations caused by quantum effects in nuclear structure we see a general trend of a sharp rise in binding energy per nucleon as light nuclei combine to form larger cores. Fusing smaller cores into larger allows nucleons to bind more strongly to each other thus releasing energy which of course is the fusion power that drives the sun. However, at $A = 62$ the slope of the curve changes as the marginal repulsion effect of increasing proton numbers overtake the marginal attraction of increasing nucleon number. This negative slope gives rise to a binding energy per nucleon in ^{235}U which is 13.7 % lower than that in ^{62}Ni . Although forming elements with heavier cores than these cost energy it is possible and the earth's crust contains substantial reserves of accessible elements with high nucleon numbers [7]. Nuclear fission energy is released as these heavy nuclei split and individual fission products form tighter bound nuclei in lower states of total energy. Following fission of a heavy element there is also usually an emission of free neutrons, this is due to another nuclear structure effect that can be visualized on the line of beta stability, Figure 1.2.

The line of beta stability, Figure 1.2, shows how increasing the proton number in a nucleus requires a marginally increasing number of neutrons to obtain stability against beta decay. Stable nucleons with low mass numbers roughly follow a 1-1




Figure 1.1: Binding energy per nucleon for a large set of isotopes. NASA [5]






Figure 1.2: A large number of isotopes, black dots represent nuclei stable against beta decay. Note the trend towards more neutrons per proton in beta-stable nuclei. Brookhaven National Laboratory [6]


ratio such as in ^{16}O whereas the most abundant form of uranium, ^{238}U , has 1.59 neutrons per proton. This explains why common fission reactions such as Equation 1.1 create lighter elements with neutron abundance and therefore release high energy neutrons.



Although fission can occur spontaneously in heavy elements the most important fission reactions in commercial reactors are the neutron-induced fissions of ^{235}U and ^{239}Pu . These nuclei both have an even number of protons and odd number of neutrons and a significant cross section for neutron absorption into an even-even nucleus¹. This intermediary nucleus is created in a high energy state due to the absorbed neutron; this makes it extremely likely to split. [22]



Often the reaction 1.2 is written in compound form 1.3 to demonstrate how the ${}^{236}\text{U}$ is only an intermediary stage of ${}^{235}\text{U}$ fission.



1.1.2 Neutron Economy

To obtain a sustained nuclear chain reaction useful for power production it is of course a requirement that the neutrons required to induce a fission event are replaced by new neutrons from that fission reaction. In a thermal-neutron induced ${}^{235}\text{U}$ fission there are on average 2.44 new neutrons released [3]. This implies that the total loss of neutrons from phenomena such as leakage, neutron decay and neutron capture in other elements (not inducing fission) must be approximately 59 % for the reaction to be stable. The state of neutron production in a reactor is described by its multiplication factor k : the ratio of the neutron production rate and the neutron loss rate. A reactor with $k = 1$ is said to be *critical* which means that the neutron flux is constant over time as one fission induces exactly one fission. A normally operating reactor as a whole will always be extremely close to $k = 1$, or *criticality*, as any deviations from $k = 1$ will create very fast swings in power multiplying over extremely short neutron regeneration timescales², this would make the reactor hard to control. Since reactors are assumed to be very close to criticality a commonly used unit³ for deviations from criticality (called *reactivity*) is the pcm, pour cent mille or $\frac{1}{100000}$, used to represent small changes in multiplication factor. The k value for the whole reactor is normally denoted k -effective (k-eff) as it covers the balance between neutron production, neutron absorption and neutron leakage for the actual system. To describe the reactivity characteristics of individual fuel assemblies the term k_∞ (k-infinity or k-inf) is used. In this case an infinite reactor core with the specific fuel lattice is created using reflective (or periodic) boundary conditions.

The k_∞ -values for individual fuel bundles and rods may be far above or below criticality, i.e. $k_\infty = 1$, contributing their share to the total k of the reactor.

¹Even-even nuclei, with even numbers of both protons and neutrons are more stable than those with odd numbers due to quantum nucleon pairing effects, it is therefore likelier for an even-odd nucleus to absorb a neutron and become even-even than vice-versa.

²Mean neutron lifetimes in nuclear reactors range from 10^{-4} to 10^{-7} seconds, whereby a small deviation from criticality multiplies extremely rapidly.

³ k is of course dimensionless.

The concept of k_∞ on the bundle-level will be explored in further detail in the next chapter. As a reactor operates it fissions the nuclei loaded into its core for production of energy. In most commercial reactors a batch of fuel is inserted for a *cycle* of 12-24 months and the core is kept closed without any possibility for addition or extraction of material during this time period. Throughout this cycle the reactor will consume, or burn, its fissile ^{235}U leaving a smaller and smaller fraction of this isotope available for fission. This decreasing macroscopic fission cross section implies that the reactivity contribution from ^{235}U would be steadily decreasing, seemingly at odds with stable operation. Four of the factors solving this problem are:

- Neutron-induced transmutation of ^{238}U to ^{239}Pu generating new fissile material.
- Neutron-absorbing movable control rods that take up excess reactivity in the beginning of the cycle to later be withdrawn.
- Addition of a *burnable absorber* (BA), such as gadolinium.
- Variation of core coolant flow to change moderation (void fraction in coolant).

Burnable absorber is of key significance to this project and its properties will be further explained in Section 1.1.5.

1.1.3 Boiling Water Reactors

Among commercial light-water moderated nuclear reactors there are two very prominent classes of designs: pressurized water reactors (PWR) that use highly pressurized light water for cooling and boiling water reactors (BWR) that operate at a lower pressure allowing the reactor water to boil.

There are many differences between the two kinds, notably the number of coolant loops (3 in PWR, 2 in BWR) and the operating pressure (~ 70 bar for a BWR [1] and ~ 159 bar for a PWR [2]). From a nuclear core design viewpoint there are a few principal features that make BWR's interesting.

- The fuel bundle lattice contains inter-assembly gaps with non-boiling water which makes the environment in which neutrons travel more heterogeneous than that of a PWR fuel. Typically BWR designs feature greater numbers of enrichment levels in more complicated patterns than PWR nuclear designs.
- The properties of the coolant water changes axially as it boils, the top region of a fuel assembly will thus operate in higher void than the bottom region. This makes the fuel design in BWRs more 3-D dependent than in PWRs.

These traits of BWRs mean that there is a large scope for fuel designs to optimally utilize uranium and burnable absorber material in the fuel and thus a good prospect for optimization; to deliver favorable properties of the fuel assemblies. [23]

1.1.4 Critical Power Ratio

The large power density of nuclear reactors brings with it several cooling issues related to the properties of the coolant water. In BWRs an important safety issue

concerns Critical Heat Flux (CHF), the heat flux from a fuel rod at which the film of liquid water present on the surface of the rod under good operating conditions can not be sustained and no liquid water is available for surface boiling (and thus cooling). As long as there is enough water to sustain boiling an increase in heat flux will cause an increase in heat transfer through increased boiling. If the heat flux rises too much, the water film can become broken up by bubbles and if the flux gets too intense large regions of insulating void can occur. When this insulating void occurs the heat transfer rate from the rod to its environment can not increase. This dangerous phenomenon is also called *dryout* and when increased heat flux leads to decreased boiling heat transfer the rod temperature can rise dramatically; risking fuel damage or perhaps even core meltdown. Dryout is said to occur when a rod temperature increase of more than 25 degrees per second is observed [20]. The value of the CHF power flux depends heavily on reactor parameters and fuel design. Modern BWR fuel employs several details to distribute water in an optimal fashion and bundle load patterns are often optimized for a good Critical Power Ratio (CPR). CPR is the ratio of CHF and the heat flux in the worst cooling channel across the reactor, stating how much the power can increase without reaching CHF. The interplay between CPR and bundle design through the important R-factor will be explained in greater detail in Section 1.2.3.⁴

1.1.5 Uranium Enrichment and Burnable Absorber

This project deals with optimizing two components over nuclear fuel bundle grids to obtain optimal bundle properties. The components are uranium enrichments. Uranium enrichment refers to the percentage by weight of ²³⁵U in the uranium component in a certain area such as fuel rod, fuel bundle or the entire reactor core. The fuel assemblies considered in this project can have uranium enrichments from natural unenriched (0.71) up to 4.95 % ²³⁵U. Burnable absorber (BA) is an element, typically gadolinium, added to the nuclear fuel to alter the profile of Figure 1.3. Burnable absorbers are isotopes that have very high thermal neutron absorption cross sections and can therefore decrease the fission-inducing neutron flux in their vicinity, see Table 1.1 for gadolinium thermal neutron capture cross sections. However, subsequent neutron absorptions during the reactor cycle eventually transmute the isotope into one with a much smaller absorption cross section. The diminishing number of nuclei of absorbing isotopes means that the negative contribution to k also decreases.

Thus the BA effect on k counters that of the enriched uranium, first lowering the neutron flux but eventually decreasing its absorption effect, becoming “burnt out”. Figure 1.3 shows how introduction of BA into a lattice drastically changes the bundle’s k contribution to a core, its k_{∞} curve, and decreases the difference between

⁴The dryout safety concern in BWRs has a PWR analogy known as departure from nucleate boiling, DNB. Both situations concern a critical heat flux where cooling properties change starkly. The difference between the two lies in the ratios of steam and liquid water in the two reactor types: PWRs have cooling channels of mostly non-boiling water with boiling occurring at fuel rod surfaces whereas BWR cooling channels contain mostly steam with films of water dispersed to cover fuel rods. DNB occurs when bubbles from boiling combine and form films that insulate the fuel rods and the boiling regime changes from nucleate boiling to film boiling. The PWR CPR “equivalent” called DNB ratio plays a similar role as CPR as a parameter to optimize to risks of dangerous steep temperature increases.[26]

Thermal Capture Cross Sections							
Isotope	Abundance	ENDF			RPI		
		Thermal Capture	Contribution to Elemental	Percent	Thermal Capture	Contribution to Elemental	Percent
¹⁵² Gd	0.200	1 050	2.10	0.00430	1 050	2.10	0.00430
¹⁵⁴ Gd	2.18	85.0	1.85	0.00379	85.8	1.87	0.00422
¹⁵⁵ Gd	14.80	60 700	8 980	18.4	60 200	8 910	20.1
¹⁵⁶ Gd	20.47	1.71	0.350	0.000717	1.74	0.356	0.000804
¹⁵⁷ Gd	15.65	254 000	39 800	81.6	226 000	35 400	79.9
¹⁵⁸ Gd	24.84	2.01	0.499	0.00102	2.19	0.544	0.00122
¹⁶⁰ Gd	21.86	0.765	0.167	0.000342	0.755	0.165	0.000372
Gd	—		48 800	100.0		44 300	100.0

Table 1.1: Neutron absorption cross sections for gadolinium isotopes, cross sections in barns. Source: Leinweber et al [21].

minimum and maximum reactivity. The effect of BA is often quantified using two parameters: the decrease in initial k_{∞} arising from the BA in a design compared to an identical enrichment design void of BA (denoted a) and the level of burnup at which peak k_{∞} occurs (denoted b). In BWR fuel BA is usually present in some fuel rods in the lattice and its concentration can vary axially through the bundle. The fuel bundles considered in this project can have BA in any rod position, except for peripheral and part length fuel rod positions, in weight percentages from 2 to 9 %.

1.2 BWR Fuel Bundles

1.2.1 Mechanical Design

The fuel bundles used in a BWR are different from PWR bundles in a number of ways. A typical PWR bundle has 17x17 rods in a quadratic grid [8]. BWR fuel bundles are usually smaller using 10x10, 9x9, or 8x8 grids [10]. The axial and radial void variation in a BWR typically makes the optimization of a BWR bundle a greater challenge involving greater numbers of enrichment grades than in a PWR. BWR fuel bundles have several adaptations to perform better in high void environments: notably mixing vanes, water crosses with non-boiling water and partial length rods. Mixing vanes are features in the mechanical design, such as small wing-like plates, that spread water laterally within the fuel bundle [4]. Water crosses are sections with non-boiling water in the center of the bundles that improve the moderation of the fuel and contribute to a flatter power distribution within the fuel bundle.

The BWR bundles considered for optimization are made up of 10x10 grids with a water cross of four central non-fuel positions for a total of 96 fuel rods. The bundle is considered to consist of four sub-bundles (NW, NE, SE, SW) of 24 rods each; within a sub-bundle rod spacing is almost equal, whereas the rod to rod distances between neighboring sub-bundles are somewhat greater due to the presence of the water cross. A special feature of the assembly designs considered for optimization is that they contain partial length rods, these are rods of approximately 1/3 or 2/3 the length of the fuel assembly. These rods are located in the center and corners of the fuel bundle, as shown in Figure 1.4.

The inclusion of these partial length rods means that the physics will differ

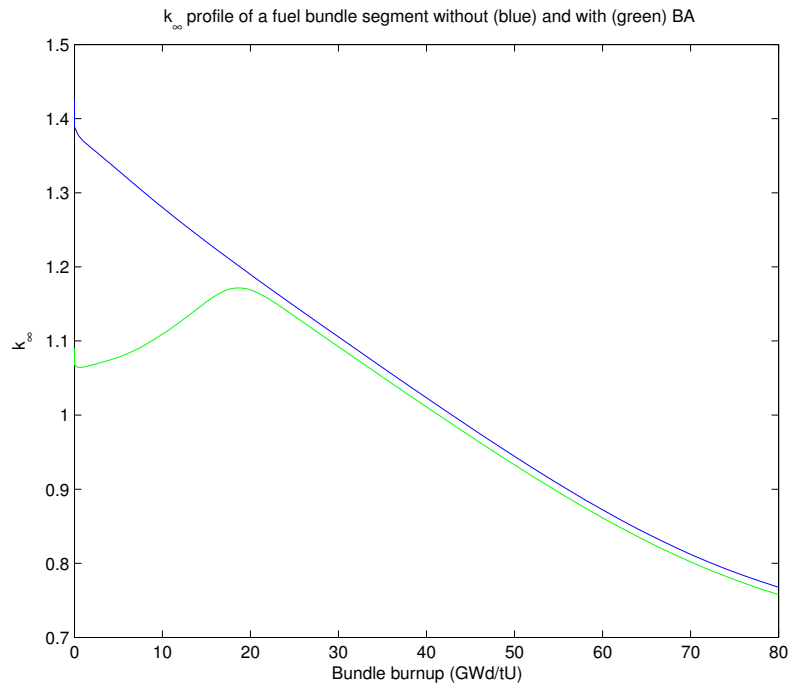


Figure 1.3: k profile for a high-enriched BWR fuel bundle without BA (blue) is contrasted with one adding 12 rods of 8 % Gd (green).

significantly between axial regions with different numbers fuel of rods. To account for this the nuclear design and all calculations are done in three slices: bottom (96 rods), middle (92) and top (84).

Optimization by changes to the mechanical design, such as varying the length of partial length rods, is not considered in this thesis.

1.2.2 Nuclear Design

The nuclear design of a bundle is the specific layout of uranium enrichments and gadolinium concentrations in the bundle. Obtaining the optimal layout of these elements is the overall topic of this thesis and a few words should be said about how to proceed.

Fabrication constraints

As mentioned previously the fuel in a reactor is composed of hundreds of bundles, bundles are in turn assemblies of 96 rods, each full-length rod is 3-4 m long and is made up of cm-long pellets stacked on each other surrounded by zirconium alloy⁵ cladding. In theory it would be possible to individually design each fuel rod using pellets of unique compositions. To connect the optimization problem to the assembly lines and economics of reality some fabrication constraints are however introduced, in Table 1.2 examples of some of these constraints are summarized.

⁵Zirconium has a very low neutron absorption cross section and is therefore used for cladding in fuel assemblies.




Figure 1.4: The layout of the fuel bundles considered for optimization. Westinghouse 2003 [11].

Typically Gd rods have axially varying composition to optimize overall performance but their design regions need not be the same as the slices used in calculation (which are defined by the length of PLRs). To simplify calculations this assumption has however been made in this project.

These constraints apply to each rod and to obtain an economical bundle design some further limitations apply to the bundle optimization as a whole:

- A maximum number of uranium enrichments used.
- A maximum number of Gd concentrations used.
- A maximum number of Gd rod types, a rod design with a specific U and Gd profile.

Optimization goals

Apart from generating the best possible bundle performance an optimization may also be targeted at reducing the complexity and cost of the manufacturing by targeting the number of fuel rod types used or the total number of Gd rods.

Rod types	Axial distribution	Possible concentrations
Uranium rods	Entire rod homogenous	2.4 - 4.8% w/o U235 in 0.2% increments and 4.95%.
BA rods	3 axial zones (bottom, middle, top), with individual Gd & U properties	0% and 2 to 8% Gd in 0.5% increments. Enrichments as for pure uranium rods.

Table 1.2: Examples of fabrication constraints used in optimization.

2	2	3.4	3.4	2.8	2.8	3.4	3.4	2.8	2
2	3.2 & 3 %	4.4	4.4	3.8	3.8	4.4	4.4	3.2 & 3 %	2.8
3.4	4.4	4.4	4.4	3.4	3.4	4.4	4.4	4.4	3.4
3.4	4.4	4.4	3.2 & 3 %	3.8	3.8	3.2 & 3 %	4.4	4.4	3.4
2.8	3.8	3.4	3.8			3.8	3.4	3.8	2.8
2.8	3.8	3.4	3.8			3.8	3.4	3.8	2.8
3.4	4.4	4.4	3.2 & 3 %	3.8	3.8	3.2 & 3 %	4.4	4.4	3.4
3.4	4.4	4.4	4.4	3.4	3.4	4.4	4.4	4.4	3.4
2.8	3.2 & 3 %	4.4	4.4	3.8	3.8	4.4	4.4	3.2 & 3 %	2.8
2	2.8	3.4	3.4	2.8	2.8	3.4	3.4	2.8	2

Figure 1.5: An example of a nuclear design: plain boxes are uranium enrichments and yellow boxes BA rods. with % declaring Gd concentrations.

1.2.3 Fuel Bundle Characteristic Parameters

A fuel bundle design such as Figure 1.5 can be characterized by a number of parameters that describe its performance, there are a three major parameters (and one unit) of interest to the optimizations in this project.

Burnup

Burnup is a nuclear engineering term for the amount of energy that has been extracted out of a nuclear fuel. It is in this project measured in GWd/tU, Gigawatt-days per ton uranium metal. Burnup is used in this project chiefly as the X-axis against which the three parameters listed below are evaluated. The burnup is therefore intuitively seen as an analogy to time, but the time in which bundle will reach a certain burnup is in no way constant but varies with the power level at which the fuel is operated.

Bundle k_∞

The first parameter of interest is the infinite lattice multiplication factor of an axial cross section of an individual fuel bundle, denoted k_∞ . The multiplication factor for an infinite lattice is computed by assuming a mirror symmetry around the fuel bundle in question, thereby effectively computing the k of an infinite reactor built solely with the type of bundle considered. Although this assumption is clearly unrealistic it serves to compute a measure of the contribution to whole-core k -effective from a single bundle. Although the reactor as a unit will of course need to operate at $k = 1$ at all times the properties of individual fuel rods in the core can vary significantly as they have very different burnup (being fresh or re-used), nuclear design and environment. While the k -eff of the core as a whole is controlled with control rods and core coolant flow to be close to unity, the k_∞ of a single bundle will vary much more over time as its Gd, ^{235}U and ^{239}Pu content evolves.

Relative Pin Power distribution

The second interesting parameter is the relative pin power distribution (referred to as P for brevity), the power generated in a chosen axial segment of a fuel rod relative to the average power in that axial segment of the bundle⁶. In bundle evaluations this distribution is obtained using the same mirror symmetry as for k_∞ . This relative power distribution is determined by the geometry of the fuel lattice in combination with the enrichment and BA layout for individual fuel rods. The internal power peaking factor F_{int} : in this report defined as the maximum relative pin power in an axial fuel bundle segment, is an important parameter of bundle performance. In the loaded reactor core the total power peaking factor is limited, i.e. the combination of internal peaking factor, radial (within core) and axial power peaking factors. If the internal peaking factor is high there is less room for radial and axial power peaking in the reactor. Furthermore a high F_{int} means that the power is unevenly distributed over the bundle which implies that the burnup will be inhomogeneous within the fuel bundle. The relative pin power is often limited in optimization by a constraint on Linear Heat Generation Rate, $LHGR$, on whole core level this constraint limits the nominal power in watts per meter of fuel rod but in bundle level optimization this value can not be computed and the limit is instead simplified as one on F_{int} .

R-factor

One important limitation for the core loading is the margin to dryout expressed as Critical Power Ratio (CPR, explained in Section 1.1.4). CPR is calculated using a correlation where important parameters are bundle power, axial power shape, bundle coolant flow and inlet sub-cooling. The inherent property of the fuel is described by the R-factor (sometimes called K-factor) and is used as input to the correlation. The R-factor correlates relative pin power distribution to a measure

⁶The nominal power of bundles depend strongly on their surroundings in the core is therefore beyond the scope of this thesis. The combination of k_∞ and relative power is enough to evaluate the performance of individual fuel bundles.

of dryout sensitivity for the fuel bundle⁷. Keeping the R-factor low⁸ will make it possible to operate the fuel bundle at a higher power level or with a lower coolant flow which has an economical value. Alternatively better CPR performance can be used for more operational flexibility. The importance of the R-factor on critical power ratio is significant: a rule of thumb is that an increase in R-factor of 1 % worsens the CPR by $\sim 2\%$ [12] but the effect can be even greater [13]⁹. The R-factor algorithm combines the internal power distribution with the individual fuel rod dryout sensitivities, normally expressed as additive constants. The additive constants for a BWR mechanical bundle design are obtained through extensive tests in laboratories where fuel bundles are heated electrically in a manner that simulates fission heat generation. By conducting experiments with different internal power distributions and fitting the results into to a basic theoretical dryout correlation the individual fuel rod dryout sensitivities are obtained as additive (fitting) constants. The exact components of such a correlation are specific for every fuel vendor and model but the main equations are similar; the correlation studied in this thesis is based on XL boiling length correlation[13]. The basic formula for the R-factor calculation is common for most XL correlations and here it is given in a general form without weighting factors or additive constants, Equation 1.4.

$$R_{i,s,z} = \left(\frac{T}{S_z} \right)^{1/2} \left[\frac{WI_i \cdot (r_{i,s,z})^{1/2} + SJ_{i,s,z} + SK_{i,s,z}}{WI_i + \sum_{j=1}^{n_j} WJ_j + \sum_{k=1}^{n_k} WK_k} \right] + I_i \quad (1.4)$$

Where $r_{i,s,z}$ is a function of the power in a relative to its sub-bundle, Equation 1.5.

$$r_{i,s,z} = \frac{P_{i,s,z} \cdot S_z}{P_{s,z}} \quad (1.5)$$

And SJ , SK are non-linear functions using the smallest value of nodal power in pin i and a weighted average of side 1.6 or diagonally 1.7 bordering nodal powers.

$$SJ_{i,s,z} = \min \left(\sum_{j=1}^{n_j} WJ_j \cdot (r_{j,s,z})^{1/2}, (r_{i,s,z})^{1/2} \cdot \sum_{j=1}^{n_j} WJ_j \right) \quad (1.6)$$

$$SK_{i,s,z} = \min \left(\sum_{k=1}^{n_k} WK_k \cdot (r_{k,s,z})^{1/2}, (r_{i,s,z})^{1/2} \cdot \sum_{k=1}^{n_k} WK_k \right) \quad (1.7)$$

It should be observed that there are two major non-linearities here, the square-root dependence on nodal power and the two functions of minima. In a linearization scheme this must be accounted for in some fashion, this will be further investigated in the results, Chapter 5.

⁷An example of R-factor effects: assuming the bundle in Figure 1.4 has a homogeneous power distribution it will have a much greater margin to critical heat flux in the corner rods since these have much greater contact with water than the internal rods, the R-factor is therefore used to quantify the effects of bundle power distribution and geometry on core CPR performance.

⁸Specifically keeping the maximum R-factor of the bundle low. The rod with the highest R-factor is said to be R-limiting and minimizing bundle maximum R-factor is a common optimization objective.

⁹Page 298.

1.3 The case for automatic optimization

Optimization has been a topic of research in nuclear engineering for a long time. While whole-core level optimization has been a well investigated territory[14], the aspect of optimizing the design of a single fuel bundle is less well documented¹⁰.

The reasons for this disparity can be.

More variables A core-level optimization typically deals with a few types of fuel assemblies and their placement in a reactor whereas fuel bundle designs must choose values from a large, potentially continuous, set of values for enrichment and BA concentration.

Internal process Bundle optimization is typically done by the fuel vendor before presenting the result to the customer, thus not showcasing the method of optimization as much as the loading pattern which can be done by the utility and the fuel vendor in collaboration.

PWR preference Another possible key factor has been that the design of a PWR fuel bundle is quite straightforward [12], whereas the loading pattern for such a reactor is a harder problem. BWR fuel assemblies are harder to design but the BWR market is only a third of the PWR market [16], making developments less prioritized.

1.3.1 The problem at hand

The task of designing a nuclear fuel bundle is intuitively well suited for automatic optimization: the problem is simply to assign values at every point in a grid in a manner such that one obtains the best possible evaluation of a specified objective function. The design constraints mentioned in the design section imply that a typical design containing 18 Gd rods with three axial slices has 78 uranium rods and 54 individual BA rod segments to optimize (BA nodes will be assigned both a Gd concentration and a uranium enrichment). The fabrication constraints adopted, 14 permitted enrichments and 16 gadolinium concentrations (including 0%) gives over

$$n = \binom{96}{18} \cdot 14^{78} \cdot (14 \cdot 16)^{54} \quad (1.8)$$

or 10^{235} possible layouts. In reality the vast majority of the combinations in Equation 1.8 are easily excluded for a number of reasons. For an isolated fuel bundle there would be no point in breaking symmetry in the design: a skewed layout will exhibit skewed suboptimal power distributions and utilize neutrons inefficiently. If the environment surrounding the quadratic bundle is homogeneous it would indeed suffice to examine an eighth of the bundle. However the environment in the reactor is not homogeneous because of the presence of control rods: in this project these are viewed as blades covering two of the sides in a (for example the North and West

¹⁰A notable exception to this is the N-streaming concept[15] which optimizes an entire loading problem of a BWR by repeatedly targeting the worst rod in the worst bundle according to some optimization target. Although this approach seems to be efficient it is not well documented (publicly) and it may well be based more on pure heuristics than an exhaustive optimization algorithm.

side) which generates a diagonal symmetry. This reduces the 100x100 grid to 55 positions of which 52 are occupied by fuel rods.

After including more realistic fabrication constraints such as:

- Total of 6 enrichments used.
- Total of 2 Gd concentrations used.
- Three types of BA rods, with axial specific Gd and U profiles.
- No BA in peripheral positions of the bundles and in part length rods.

and assuming 10 are Gd rods with two of these on the symmetry line the number of combinations is reduced to:

$$n = \binom{29}{10} \cdot \binom{10}{3} \cdot (12^3 \cdot (12^2 \cdot 6)^2) \cdot 6^{42} \cdot \binom{14}{6} \cdot \binom{15}{2} \quad (1.9)$$

or $5 \cdot 10^{56}$ combinations. The majority of these designs will be incompatible with average enrichment targets and the like, but it still remains a magnificent-sized problem.

The overwhelming issue with optimizing this grid is that the values of interest; such as the internal form factor F_{int} and k_{∞} , depend on the layout in a highly non-trivial manner. A simple function for the exact dependence of the form factor on enrichment or Gd concentration cannot simply be constructed. Without an analytic function from input to output we cannot obtain an analytic gradient, and optimization strategies based on those are therefore useless. In this project the calculations of the bundle parameters have been conducted by a licensed 2D coarse-grid finite mesh program, PHOENIX (referred to as 2D lattice code), that puts up neutron balance equations and solves them using vast cross section libraries to obtain power distributions and k_{∞} as a function of burnup. Using this program it has during the project been possible to calculate around 10 slice designs in parallel with approximately 10 seconds computing time on the cluster (during normal use). It is therefore evident that simply evaluating any significant fraction of the possible designs without a clever strategy is impossible in decent time.

1.3.2 Optimization strategies

Without an analytic gradient to work with the scope for optimization appears rather limited. There are however techniques that can approach a problem such as this. These techniques can generally be divided into deterministic and stochastic methods.

Stochastic methods

Stochastic methods utilize random but clever perturbations with inspiration from evolutionary biology's natural selection¹¹ or condensed matter physics' annealing processes¹² to generate the bundle designs to evaluate. The greatest strengths of stochastic algorithms is that the evaluation part can be done in any way with any type software being chosen and that the objective function can use any result from

¹¹Called genetic algorithms, GA.

¹²Simulated Annealing or SA.

this evaluation¹³. A weakness is however that fuel bundle PHOENIX evaluations have to be done for every considered design in every step and that the program will more or less run blindly. It will not know whether the objective function value being sought is unobtainable or how long calculation time a successful optimization will take [19]. It is also vulnerable to local optima, a perturbation contribution (mutations in the GA example) ensures some departure from local optima but there is no guarantee that a the best possible solution was not a combination of two “inferior” parents that were both discarded on their individual merit [14].

Deterministic methods

To avoid these pitfalls of stochastic optimization this project uses a *deterministic* method, namely the simplex method. This method is a linearization scheme that can optimize any linear problem. Although less exact in each evaluation, it has the great advantage of giving the user the power to choose the balance between optimality and calculation time.

1.3.3 The simplex method

The theory behind the simplex model is covered in the appendix. It has several advantages over stochastic methods in that it does not get trapped in local optima, can quickly determine whether a set of constraints is infeasible and can give an approximation of how close to the global optimum a found solution is. Its key disadvantage is that its assumption of linearity along with the non-linearity of the real world necessitates several linearizations of the problem at hand. The degree of success of using the simplex model on this problem will depend on the validity of the model being used: if large errors propagate in the optimization program’s approximations then the task of finding the design which performs best according to these approximations becomes pointless as its connection to reality is severed. A key task in this thesis has therefore been to evaluate the approximations used in the program and if possible improve them. Better approximations are likely to increase the size or complexity of the problem, thereby prolonging the optimization run-time.

1.3.4 Other work in the field

The simplex method has been implemented previously for whole-core loading pattern optimization, for example by Kim & Kim [17]. In their model core-reactivity is targeted for maximization at end-of-cycle. Including other optimization methods there are several patents regarding whole-core nuclear optimization ([19],[18]), but apart from N-streaming discussed earlier [15], no patents pertaining optimization of individual bundles have been found.

¹³A genetic algorithm applied on a fuel bundle optimization problem would widely follow these few steps, words in italic are biology parallels: 1. Stochastically generate a large number of bundle designs, *a population*. 2. Convert the bundle design to a string of values, *a genome*. 3. Evaluate the designs and calculate their respective fitness values according to an objective function. 4. Discard the worst X percent of the population, *survival of the fittest*. 5. Stochastically combine genomes of the survivors, *breeding*. 6. Impose some random *mutations* on the property strings. 7. Return to evaluation stage with the new generation of designs. After a set number of times or when a desired objective function value is reached, the program can terminate and the best design be presented.

Chapter 2

Existing Optimization Strategy

2.1 Outline

The tool used to improve upon in this thesis is a framework of programs operating in the following manner:

- Reference nuclear designs are entered and many new designs with grid positions individually perturbed in enrichment or BA are generated from them.
- Perturbed designs are sent to a 2D lattice code running on a cluster to perform calculations.
- Result data is extracted from files to form sensitivity matrices for use in a commercial simplex program.
- Sensitivity matrices enter into an optimization module in a commercial simplex program where their values are used as constraints for an optimization BA and enrichments; linking these values to bundle parameters.
- Simplex optimizer evaluates designs and outputs the values of decision variables which, in its approximated model, correspond to the optimal design.

The optimization code written in proprietary Optimization Programming Language (OPL) implemented on IBM's CPLEX optimizer. The idea behind this strategy is that the effect of a number of combined perturbations can be reasonably well approximated by the sum of their individual contributions to any fuel parameter. The optimization strategy used is extremely versatile and it provides an exhaustive search of the solution space thanks to simplex implementation. It can not only produce a solution but also compare its objective function to an approximation of the objective function of the global optimum.

The framework used in the project optimized the design in two distinct steps: the first step took a fixed enrichment distribution reference without BA to optimize BA positions with regard to bundle parameter constraints and optimization objectives, the second step fixed these BA positions to optimize the enrichment in each fuel rod and the Gd concentration in the BA rods. The reason for this division of tasks had been that both approximation errors and computing times became very large in experiments for this larger more general problem. A more detailed justification for this approach will be given in Section 5.1.1. The first step had been evaluated

in some detail and measures had been taken to mitigate approximation errors, the second step was however untested by the intended users (the problems concerning this are explained in Section 3.2.1). The simplex implementation of the two steps in CPLEX was slightly different. The perturbation variables for Gd position and concentration in step 1 were integers taking of dimensions rod number and Gd-concentration, thus taking the value 1 at combinations that were true and 0 where they were false. The second step instead formulated linear decision variables as floats, Δ -enrichment and Δ -Gd, the resulting variables for chosen enrichments and Gd concentration were however limited to certain allowed values, effectively making these variables binary as well.

To represent the actual physics in the reactor a large number of constraints are needed to link the decision variables of enrichment and BA distribution to real parameters such as k_∞ , the fuel rod relative power distribution, internal power peaking factor (F_{int}) and R-factor. The key limitation using a simplex model is that these constraints have to be expressed linearly. In the results section an analysis of the linearization errors associated with these approximations is presented.

Linear perturbation terms are computed by running a large number of possible perturbations in the design through a 2D lattice code obtaining values of the relevant parameters P and k_∞ . From this, sensitivity matrices (S-matrices) are generated: as an example the enrichment S-matrix element for k_∞ perturbation in a certain rod is computed in the following manner.

- A reference design is entered into the 2D lattice code.
- Two designs, where the enrichment in the specific fuel rod is changed to the maximum allowed enrichment and lowest allowed enrichment respectively, are entered.
- k_∞ -values extracted from the 2D lattice code output and the difference between the maximum and minimum cases is computed and divided by the enrichment difference.
- The value $\frac{\Delta k_\infty}{\Delta U}$ is computed and entered into an S-matrix file (in this project an Excel-file).

The S-matrix elements for the relative fuel rod power distribution are calculated in a similar way but here S-matrix elements for one perturbed fuel rod position are calculated for both the specific rod where the perturbation is done and for all other fuel rod positions in the bundle. For a typical fuel bundle there are 52 rods (taking symmetry into account) with 3 axial slices (their number can be chosen freely but 3 are always used in this project since it is the minimum required to correctly model partial length rods) for which perturbations need to be calculated. Computing the S-matrix elements therefore requires 315 simulations for the enrichment alone (a max min run for every position, plus three reference runs). Fortunately the core simulator runs on a cluster and will therefore usually calculate 10-12 cases simultaneously, finishing the entire process in less than 10 minutes. The resulting sensitivity matrix for k_∞ has 52 elements for every burnup step considered. The parameters R and P are rod-wise properties which mean that there will be one sensitivity matrix for every axial segment of every fuel rod. Pin power sensitivity matrices take into account every other rod in the same slice which gives $52 \cdot n_{bu}$ (number of burnup

steps) elements in every matrix. R-factor calculations are done by sub-bundle which means that the size of the S-matrices for R follow the number of independent fuel rods in the sub-bundle in question¹. These figures are indicators of the size of the problem and explain why these simplex optimizations typically have around 100000 variables. The S-matrix files are read by the simplex program (CPLEX) and the sensitivity matrix elements are implemented as linear programming constraints in between “true” decision variables enrichments and Gd concentration on one hand, and the result decision variables such as k_∞ and P^2 .

Upon starting the project the actual mechanics of generating S-matrices and putting them into excel sheets was very cumbersome, but since the data management is not a focus of this thesis no greater effort has been made to alter it.

2.2 Constraints and objective

To connect the decision variables used by CPLEX to each other and the constants describing the problem a long list of constraints is needed, the constraints can be divided into four different classes.

- Constraints representing the approximations for bundle physics parameters.

$$k_\infty = k_\infty^{ref} + \sum_i^{n_U} (\Delta_i U \cdot S_i^U) + \sum_j^{n_{Gd}} (\Delta_j Gd \cdot S_j^{Gd}) \quad (2.1)$$

Meaning that the k_∞ value in the bundle is equal to the reference bundle k_∞ plus the sum of the changes in Uranium enrichments rodwise multiplied by S-matrix elements plus the sum of changes in Gadolinium concentrations rodwise multiplied by their S-matrix elements over the BA rods.

- Constraints describing actual constraints on parameters in the model.

$$\max(P_{i,z}) < Limit \quad (2.2)$$

- Constraints representing fabrication constraints, such as maximum number of enrichments utilized.
- Logical constraints, such as ”one and only one enrichment per rod position”, to eliminate bogus solutions.

The distinction between the first two types of constraints if for convenience, they could very well be combined into the same equations but on a problem of this scale that would make the OPL code much less readable without making gaining anything in optimization time. The *objective* is the decision variable that CPLEX seeks to minimize (or maximize), it must be a real scalar but that can in turn for example be a sum of several variables or a maximum of a vector in some interval. The most common objective in this project was minimization of the bundle R-factor. A glance of the CPLEX progress window is seen in Figure 2.1. An optimization following the

¹R-factor calculations do take into account the power of other sub-bundles, but through a scalar “mismatch factor”, therefore there are either 14 (NW&SE) or 24 (SW) S-matrix elements for every




Figure 2.1: Part of the CPLEX window showing parameters of a running optimization and its progress. The user can view the progress of an optimization with its incumbent solution and an estimated bound of the global optimum, the difference between the two is referred to as the gap.

described framework runs in CPLEX until the ratio of the best found objective function value to CPLEX’s estimate of the global best value comes within a prescribed “gap tolerance”³. When the optimization terminates the design associated with the best objective function value is post-processed and exported, its most important parameters are of course the enrichment and/or BA profile.

burnup step.

²In the CPLEX implementation, variables such as k_{∞} and ΔGd are not fundamentally different; they are both treated as decision variables for CPLEX to assign values to. Through constraint equations using S-matrix values these decision variables are linked by the linear equations formulated.

³During a minimization the best found value will decrease as new better designs are found, and the estimate will increase as more and more branches are evaluated and their actual best values are found.

Chapter 3

Method

3.1 Research Questions

The project broadly followed the topics of a set of research questions.

- Can a linear approximation of the constitutive relations in a nuclear fuel bundle through a simplex optimization generate potential nuclear fuel bundle designs? How does the existing scheme perform and how can it be improved?
 - How do the initial conditions of enrichment and BA weights specified in the reference affect the outcome?
 - How do results from such an optimization compare to “manually optimized” designs made by nuclear engineers using perturb-and-investigate programs?
 - Is it feasible and motivated to incorporate 2nd order terms, or piecewise linear approximations in the optimization method (thereby increasing the size of the simplex problem) and how much would this affect optimization time?
 - Are there major simplifications that can be made?
- Are there optimization problems associated with reactor cycle length?
- How much engineering time can be saved with a functional optimization tool? Which barriers hinder implementation?
- Does the optimization tool have a value beyond associated time-savings in the design phase? Would other parties be interested in applying it?
- Is it possible to reduce the total number of BA rods in a bundle compared to manually optimized designs?
- Can bundle design optimization be coupled with core-loading pattern optimization?
- How should the project proceed?

3.2 Project Progression

The project for this thesis was suggested by a nuclear company with a prototype BWR Fuel Bundle Design Optimization project. The thesis was supposed to investigate the validity of the designs suggested by the program and compare these with designs obtained by nuclear engineers using manual trial-and-error schemes to produce a recommendation for further development of the tool. To pursue the strategy knowledge would have to be gained in the theory and programming of optimization. This was the initial method used but due to the prototype nature of the program a two-prong strategy was used: optimizations would be run from start to finish while error-checking and obtaining a functioning strategy to finally use this strategy to run optimizations on cores where manual designs had already been made. Using these as a benchmark the performance of the optimization scheme could then be evaluated, conclusions could be drawn and recommendations formulated.

3.2.1 Practical Issues

Using the prototype optimization framework proved harder than expected for a number of reasons. Unnecessarily complicated data handling made the system error-prone and although several modifications were made to improve it the focus of the project was not to produce a production-ready program so most of the programs involved were only modified rather than replaced.

During the project it was found that the optimization framework contained some unfinished areas and errors. Some of the issues were:

- Optimization step 1 could functionally optimize BA placements but the enrichment and Gd concentration optimization in step 2 was hard-coded for a single case of BA placements.
- Symmetries were not fully utilized. In the S-matrix generation full bundle (i.e. without diagonal symmetry) calculations were used. The result of an off-diagonal perturbation was then expressed as the sum of the response of the two symmetric positions in lower left half and the upper right half of the bundle. This way some accuracy is lost for positions close to the symmetry diagonal since interaction between the perturbed positions makes for non-additive results.
- No functionality to control average enrichment.
- Some errors were present in R factor calculations, including weighting factors and faulty linearizations.
- k_∞ profile in the enrichment optimization step (step 2) could not be determined but was instead constrained only so that the starting value from step 1 did not change. This made it impossible to define a desired k_∞ profile while also making the problem artificially seem much smaller than it was.
- Pre- and post-processing features in CPLEX not fully utilized.
- The only initially available CPLEX computer was very slow and its program version old, making full sequences of optimizations for relevant burnup lengths nigh-impossible to run.

3.2.2 Improving the Approximations

During the project the main body of work concerned several methods of improving the accuracy of the model to evaluate how optimality could be improved. The results of these investigations are covered in the results section and they included.

- Piecewise linear approximations of characteristic parameters.
- Adapting input reference cases to match the desired design as much as possible, without loss of generality.
- Quadratic approximations for bundle design parameters.
- Shadow-factor correction for k_∞ approximations in step 1 extended to step 2.
- Omitting R-factor calculations entirely, instead replacing them with relative pin power targets.

3.3 Two Optimization Cases

To test the application feasibility of the optimizer a complete study on two reactors was performed. The design characteristics for the two reactors are very different, making them good candidates for an optimization benchmark.

Type	Short-cycle reactor	Long-cycle reactor
Power Output	857 MWe	1102 MWe
Bundle average enrichment ¹	3.41 w/o U-235	4.15 w/o U-235
Cycle length	12 months	24 months
Bottom Slice enrichment	3.62	4.41
Middle slice enrichment	3.64	4.48
Top slice enrichment	3.57	4.44
Enrichments used	2.0, 2.8, 3.2, 3.4, 3.8, 4.4	2.8, 3.8, 4.2, 4.6, 4.95
Gd positions per slice	8, 8, 8	15, 15, 14
Gd concentrations (%)	3	5, 7
Optimization objective	Minimize F_{int}	Minimize R_{max}

Table 3.1: Design features for two commercially offered fuel designs used as optimization criteria.

The bundle design guidelines in Table 3.1 come from core-level optimizations in real fuel designs which make them suitable for use as benchmark cases for the automatic optimization. The bundle enrichments used in the designs were set as constraints to the optimizer, along with the number of gadolinium rods and concentrations. To obtain the right k_∞ profile the properties a and b , explained in Section 1.1.5, were extracted from the designs and set as constraints for the optimizer. The objective variables of the designs were set as target variables in the automatic optimizations.

Chapter 4

Apparatus

4.1 S-matrix hardware and software

4.1.1 Hardware

The nuclear calculations for bundle parameters were done on a local UNIX cluster, generally it allocated 10 or more cores to the perturbation computations necessary for S-matrix generation.

4.1.2 Software

Apart from the licensed 2D lattice code PHOENIX a number of other programs were used on the cluster: a local program that took 2D pin powers to calculate R-factors as well as MATLAB for simulations of R-factor behavior and enrichment inputs. The output files from the simulator were read and S-matrix elements calculated by locally developed FORTRAN scripts, which were eventually re-written. The output files from these scripts were later exported into Excel file format using Excel macros, these Excel files could in turn be read by CPLEX for optimization input.

4.2 Simplex optimization hardware and software

4.2.1 Hardware

Initially the CPLEX optimizations were run remotely over the Atlantic on a 10 year old desktop computer with a 2 GHz dual-core Intel processor and 2GB RAM.

At the end of the project a new remotely operated computer was installed, this machine had two processors with 4 dual-thread 2.4 GHz processors each for a total of 16 CPLEX threads on 32 GB RAM.

4.2.2 Software

The proprietary simplex optimization suite CPLEX was used throughout the project. CPLEX is an IBM program with an associated proprietary optimization programming language OPL. The codes available at the start of the project were written in OPL and executed in CPLEX 12.2.

4.2.3 Equipment upgrade

Along with the new hardware that became available late in the project a newer version (12.6) of CPLEX was installed: this version had improved heuristics and when supplied the same problem as the old one it reported a different size of the resultant matrix. The impact from upgrading hardware and software was immediate and positive bordering on painful. The computation time needed to find one possible solution (although not the optimal by far) for an enrichment optimization problem decreased from 17 hours to 4 minutes. Before this upgrade it had not been possible, using the available framework, to run optimizations on full cycles and experiments were limited to investigation of approximations rather than comparisons to manual designs. After the new machine was installed it became possible to optimize enrichments for full cycles and to consider any extensions of the project calculating parameters more efficiently.

Chapter 5

Results

5.1 Evaluating the framework

A considerable part of the project concerned the evaluation of the framework previously developed described in chapter 2 and specifically the effects of the linearization employed in the optimization. This had not before been well investigated and a review was necessary to improve the framework.

5.1.1 Two-step approach

One important study was to investigate whether the two-step approach taken in previous development was indeed necessary. It was found that perturbation matrices for all elements change dramatically when BA is present (see Table 5.1) so the fine-tuning of enrichments and BA concentrations was easier to do with BA positions locked¹.

S-matrix case	Without neighboring BA	With 6 % neighboring BA
Δ_{k_∞} , enrichment change	0.00275	0.00388
Δ_P , enrichment change	0.456	0.392

Table 5.1: Impact of adding BA in one rod on enrichment S-matrices in a diagonally neighboring rod. The effect on perturbation calculations in neighboring rods is significant in both k_∞ and P , justifying the separation of optimization into two steps. Full data for perturbations with and without BA is found in Figure A.1

The data in Table 5.1 concerns S-matrices for k_∞ and relative power P and how these change. The second column values concern perturbations in uranium enrichment in a *gadolinium-free* environment, whereas the third column values are from the same perturbations in an environment where gadolinium is present in a

¹The effects of enrichment perturbations in BA rods could be accounted for by calculating them for situations without BA but the effect on neighboring rods would be too big to accommodate. Although the absolute errors on k_∞ are small it must be noted that this parameter is scalar and contributions from all rods are summed up giving a prospect of large error propagation. With this justification a division of tasks was implemented but of course it did mean that the risk of sub-optimization increased.

neighboring rod. It is clear that the introduction of gadolinium changes the S-matrix elements for enrichment perturbations quite substantially, and this warrants the division of tasks used in the optimization: where BA placements are handled by themselves so as to not affect the accuracy of calculations concerning enrichment variations.

Second step issues

As aforementioned the second CPLEX optimization step concerning enrichments and Gd concentrations had not previously been well-tested and a substantial effort (re-writing a few hundred lines of code) had to be made to generalize the data handling and constraints since these had been hard-coded for a reference case with a certain set of BA rod types and positions. Introducing variables for BA rod types and positions required some work but did neither change the result nor the execution of the hard-coded reference case, therefore full optimizations could continue for any BA geometry using the re-written code representing the same fundamental equations. Something that did change the computation time was a change in how the k_∞ profile in the second step was limited, the code received at the onset of the project limited k_∞ at 0 burnup to within a very small (100 pcm) interval around that of the reference design, which basically prohibited the second step from doing any major changes to the design as these would violate this condition. This criterion did not reflect the nuclear design process and was instead replaced by criteria on two parameters: the difference in k_∞ at 0 GWd/tU burnup which would result from removing all BA in the design, called a , and the burnup at which maximum bundle k_∞ occurs, called b . These values, along with the average enrichment, are required to fully constrain the k_∞ of the bundle and were therefore approximated in the simplex model. Rewriting the program so that k_∞ followed conditions on the a and b parameters meant that many more designs became possible and the calculation time increased by more than one order of magnitude.

5.1.2 Symmetry issues

The existing optimization model uses diagonal symmetry, as described in the theory section. However, the perturbation matrices were generated with full-bundle calculations making perturbations only in the lower-left half. The results from these perturbations in the lower and upper half was added, i.e. the S-matrix value for a change in rod position (I,J) is set to the sum of the result in positions (I,J) and (J,I). For k_∞ which does not have a position dependence the consequence will be that the results of a perturbation will simply be doubled to account for the effect on the upper half.

The hypothesis in this project was that this was an unnecessary simplification; actually making the calculation process more cumbersome and less accurate.

Consider first Gd perturbation in rod position (9,2): an assumption of symmetry (maintained throughout this project and justified in the Section 1.3.1) means that the same perturbation should be present in (2,9). These rods are far apart from each other and their interaction can be assumed to be small. However, consider instead a perturbation at (3,2), implying the same perturbation in (2,3): these two rods border each other diagonally and it should be expected that the drop in neutron flux and its effect on k_∞ from these two changes will not be twice that of an individual change.

The result of these perturbations can be seen in Table 5.2 where a BA-free reference with $k_\infty = 1.37557$ has been perturbed symmetrically (on both sides of the NW-SE diagonal line of symmetry) and asymmetrically (only below the line of symmetry) and S-matrices have been calculated from this; the asymmetric perturbations were multiplied by two as in the original optimization code. If all rods were independent of changes in other rods then the effect of symmetric perturbation would simply be twice that of the asymmetric and the S-matrix elements would be equal.

Type perturbation	Far apart	Bordering nodes
Symmetric	1.30054	1.31372
Asymmetric	1.33847	1.34145
S-matrix symmetric	-0.07503	-0.06185
S-matrix asymmetric	-0.07422	-0.06824
Error from asymmetric perturbation	-1.08 %	10.3 %

Table 5.2: Effects from adding burnable absorber in mirror rods close together or far apart and calculating k_∞ at 0 GWd/tU with symmetric or asymmetric perturbation. BA-free reference = 1.37557.

The effect of a perturbation, one that will always be symmetric, is not well represented by doubling asymmetric perturbations for k_∞ . The effect on bordering rods (mirror rods that also border each other, such as 3,4 and 4,3) is expected as the effect on k_∞ from addition of two Gd rods close to each other is not quite twice as strong as just applying symmetry in the equation, due to a shielding effect. The effect on the far separated nodes is less obvious, it seems the effect of two Gd perturbations is greater than the sum of its parts². Regardless of the reason for it: it is evident that an asymmetric perturbation calculation is not a good representation of the problem at hand.

The data-handling and S-matrix calculations were changed to account for symmetric perturbations and introduced into the framework. This removed previous programming that accounted for whether a rod was on the diagonal or not and whether it should be doubled, since all rod perturbations became equal when perturbations were done symmetrically.

5.1.3 Average enrichment implementation

Average enrichment in a bundle depends on both uranium enrichments and Gd concentrations, see Equation 5.1. The latter affect average enrichment since a rod with high Gd content contains less uranium and its enrichment thus weighs less on the average (with for example 9 % mass fraction Gd_2O_3 only 91% of the fuel rod weight is related to uranium).

$$\epsilon = \frac{\sum_i^n \epsilon_i \cdot l_i \cdot (1 - c_i^{Gd})}{\sum_i^n l_i \cdot (1 - c_i^{Gd})} \quad (5.1)$$

ϵ is the average enrichment in the bundle while ϵ_i denotes the enrichment in rod i , c_i^{Gd} its Gd concentration and l_i its length. In this equation Gd concentration effects on the density of the rods have been neglected. In a simplex implementation the

²It could be related to asymmetric designs having lower k_∞ than symmetric ones

average enrichment should be present as a constraint on the problem; this can be visualized by defining a decision variable depending on enrichment and Gd concentration and limiting this to be within a tolerance interval of deviation around the desired enrichment. Since this equation is not linear it has to be modified to be able to implement in a simplex optimizer, in this project the easiest route was taken: simply using the reference Gd concentration in the denominator omitting the dependence on changes in BA content, thus the problem of dividing decision variables with each other was eliminated and a linear implementation obtained³.

5.2 Improving approximations

The lion's share of the project was devoted to improvement of the approximations and methods used in the optimizer to account for the characteristic bundle parameters. Given the constraints imposed by linearity two main types of errors were expected to be encountered: local non-linearities, that the effect of changing a parameter in a node was not linear, and cluster effects, that the effect of several changes was not equal to the sum of its parts. The tools available to the simplex program at first seemed limited and two "strategies" were considered: looping the optimization⁴ and improving the framework and approximations.

Looping

Looping the optimization is a very crude way of improving results. It entails running an optimization with low tolerance, accepting the solution as a reference, calculating new perturbation matrices and running the optimization again. This can be repeated as many times as necessary until the discrepancy between PHOENIX results and model results diminishes below an accepted value. A major risk is that this partition leads to sub-optimization by locking on to a poor path. In reality the only increase in accuracy from repeating the optimization will come from the smaller absolute changes done in the CPLEX step. It was found that the main gain from running an optimization twice was that the second run can be done without the result and reference average enrichments and gadolinium levels differing very much, that the initial guess becomes "better". This can however be done much more easily before the first step even starts, see Section 5.2.1. Any repetition of the framework was therefore discarded, especially considering that the CPLEX run-times were often measured in hours.

5.2.1 k_{∞}

The most thoroughly examined fuel design parameter was k_{∞} . Due to its scalar nature that takes contributions from an entire slice into account it is more easily graphed than the fuel rod properties R-factor and relative pin power.

The effects on k_{∞} from uranium and BA changes are quite different in nature; Figure 5.1 shows the bundle k_{∞} as a function of varying uranium enrichment in one

³It was late in the project suggested that accounting for BA effects on average enrichment could be omitted altogether and possibly speed up optimization, this was not done to keep results comparable.

⁴This was not a very clever or promising strategy, but was investigated nonetheless.




Figure 5.1: Bundle k_{∞} effect from enrichment changes, 8 samples.

fuel rod from 2.4 to 4.95 %. As can be seen these effects are reasonably linear.

Burnable absorber is harder to model using linear assumptions; Figure 5.2 shows the effect on k_{∞} from a BA addition of 2-8 %. Note that there is not only a strong curvature but the curvature undergoes sign changes during the burnup interval studied. This is due to the depletion of BA, where its effect virtually disappears. Different BA concentrations have different depletion times, at a certain burnup a low concentration of BA may be almost fully depleted whereas for a high concentration a large part of its effect still remains. The effect on k_{∞} from BA perturbations is much less linear than from enrichment changes, this increases the risk of poor approximations in a simplex implementation.

Choosing a reference case

The reference case plays a great role in the accuracy of parameter estimations. Figure 5.3 shows the real and estimated values of k_{∞} in two step-2 optimizations both targeting a slice enrichment of 4.48 % ^{235}U : the top one uses a reference design with an average enrichment of 4.40 % ^{235}U whereas the one at the bottom uses 3.95 % ^{235}U . It is to be expected that identical perturbations on different reference designs will generate different S-matrix elements, in Table 5.3 it is shown how sensitivity to BA changes differ depending on its enrichment level. The higher reference enrichment in the bottom case means that its k_{∞} is greater which could explain a greater S-matrix element. Although the differences are small it must be noted that these are effects from individual rods on a bundle parameter and thus the combined error can become significant.

The k_{∞} approximation errors from two optimizations can be seen in the two right panes of Figure 5.3; the errors have different signs but a very characteristic form present in all optimizations featuring BA changes. The optimization in the top design featured a +1 % Gd change in 3 rods and a -1 % change in 11 rods. The

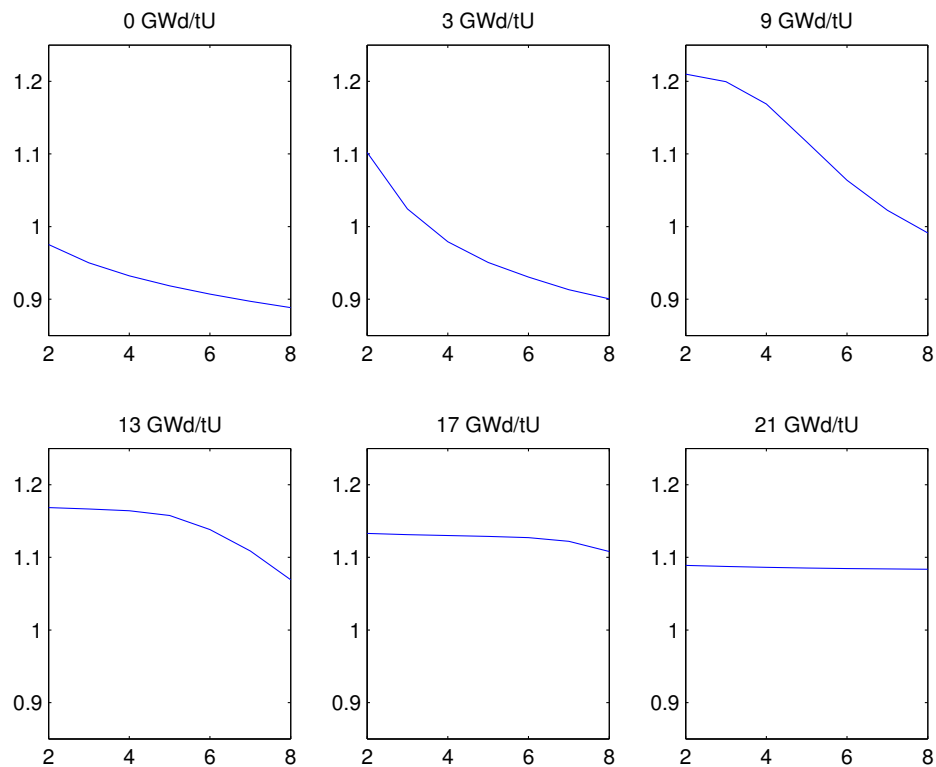


Figure 5.2: Bundle k_{∞} (y-axes) effect from Gd concentrations (in % on x-axes) changes, 7 samples in each subplot. The subplots represent different burnup steps studied, from 0 GWd/tU (fresh bundles) to 21 GWd/tU (Gd virtually burnt out) and exhibit different types of curvature.

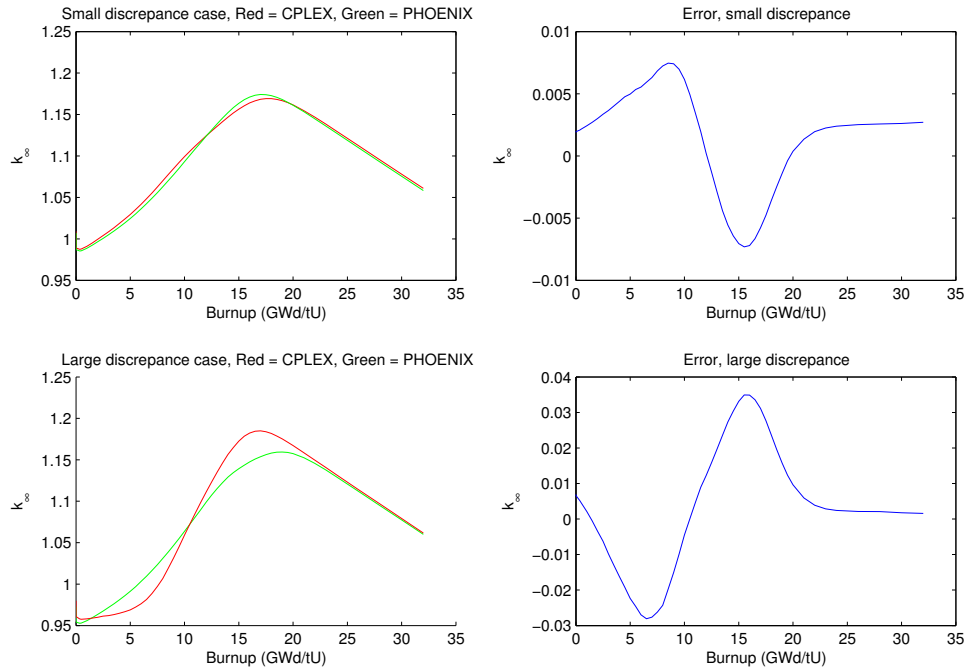


Figure 5.3: The level of discrepancy between reference and target designs plays a large role in the accuracy of CPLEX approximations. Note that the scale of the errors in the bottom case is 4 times larger than the top case.

Enrichment	k_{∞} at 4% Gd	k_{∞} at 8% Gd	S-matrix
3.89% ^{235}U	0.86308	0.86031	-69.25 pcm/%Gd
4.32% ^{235}U	0.94929	0.94635	-73.5 pcm/%Gd

Table 5.3: k_{∞} (at 0 burnup) and corresponding S-matrices from BA perturbations in rod (2,2); note that the rod perturbations are the same but the difference in enrichment background creates a difference in the S-matrix.

optimization in the bottom design featured a +2 % Gd change in 14 rods and a -2% change in 5 rods for a large total BA addition.

Combining the knowledge from Figure 5.2 and the aggregate changes in Gd concentration it can be argued that the majority of this error comes from BA. It can be seen in the bottom pane that CPLEX at first overestimates the effect of the BA addition by showing a much lower k_{∞} value than the actual, an explanation for this could be the marginally decreasing effect that rod increases in BA contribute (increasing BA in one rod keeps k_{∞} low for longer but the effect of several increases is not the sum of its parts). This initial overestimation of BA effect can then bring with it an underestimation of remaining Gd concentration later in the cycle. The error in the top pane has a mirrored shape since the BA change is now a negative one, the error is smaller corresponding to a smaller aggregate BA change. Both CPLEX approximations converge on the true value further along in the burnup sequence, hinting that the enrichment component of the error is not as significant as the BA component. The four times larger error in the bottom panes of Figure

5.3 can however be abated by better choosing a reference enrichment to avoid large enrichment changes and their associated changes in BA concentration. A better fitting reference design means that the optimization generates smaller changes in second step and smaller errors in characteristic parameters.

Piecewise linear approximations

Modern simplex optimizers such as CPLEX or Gurobi allow the use of piecewise linear constraints. In the eyes of the optimizer the calculations representing k_∞ and other bundle characteristic parameters are merely constraints tying together decision variables (enrichments and BA concentrations) to other decision variables (k_∞ , F_{int} and R). Therefore a piecewise calculation could be used to represent any bundle physics; of particular interest for this project was the effect on k_∞ from BA changes in step 2⁵.

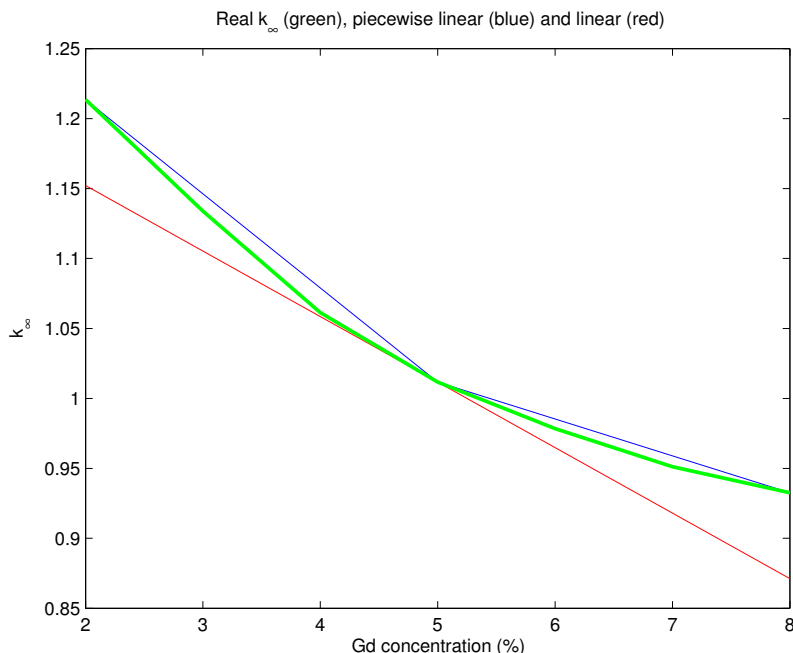


Figure 5.4: k_∞ values from 7 PHOENIX samples (green), linear interpolation between extreme points around the reference (red) and piecewise on each side of the reference (blue).

The appeal of piecewise linear approximations is displayed in Figure 5.4 where it can be seen that the error arising from linearizing typical BA behavior can be significantly decreased by using a piecewise approximation. This implementation better utilizes the available data, note that no new PHOENIX calculations are necessary⁶ since the perturbation scheme calculates the two extreme points for every fuel rod position, along with a reference case used as a midpoint⁷. This gain does

⁵ k_∞ -errors can disqualify an optimization and BA changes are their most likely culprit.

⁶Additional intermediary calculations in the S-matrix machinery are however necessary, but none involving the 2D lattice code.

⁷If the curve is to be divided by a piecewise linear function of more segments than two, new PHOENIX calculations would be required.

however come at the cost of approximating every decision variable under piecewise linear constraints as two separate variables⁸: in our case one positive and one negative. This increase in variables could prove to be a problem for CPLEX and in this implementation it does turn out to be so: calculation time increased 3-fold while reaching a tolerance 3 times greater than using linear assumptions.

Shadow factors for aggregate BA concentration

Another approach taken to account for non-linear effects of changes in BA concentration was to introduce a shadow factor counting the aggregate Gd amount in a bundle. Such a factor would then multiply the perturbation element from BA and decrease or increase its effect. This method was used in the first optimization step to account for cluster effects of the total amount of BA rods. There the correction for total number of BA rods was added to the BA effect of every rod; since decision variables cannot be multiplied with each other (a non-linearity) the binary decision variable for Gd instead had an extra dimension, so that it was true only for a combination of rod position, Gd concentration and total number of Gd rods⁹. This would be possible to implement in step 2 as well but the decision variable for Gd changes would need a new dimension with one row for every distinction level in total BA amount. This situation is similar to that of piecewise linear constraints but harder to implement which meant that after the piecewise implementation showed unacceptable increases in computation time it was abandoned.

5.2.2 Power distribution

The power distribution calculations showed some behavior similar to k_∞ with burnable absorber positions being most susceptible to errors. Figure 5.5 shows how the introduction of BA into 10 positions affect the accuracy of S-matrix calculation compared to 2D lattice code output.

Figure 5.6 shows the same differences as Figure 5.5 but later in the cycle where k_∞ peaks, it is evident (and not very surprising) that the large approximation errors in pin powers for a fresh bundle decline greatly as the total effect of BA on the bundle fades. This time in the cycle is the time where approximation error are most important since it is the time at which total power output in the bundle will peak. The errors in pin power estimation will be of great importance to the alternative R-factor optimizations discussed in Section 5.2.3.

Step 2 optimizations exhibit smaller errors since the Gd rod positions are fixed and thus avoiding strong coupling interaction non-linearities. In an initial state such as Figure 5.7 the largest errors arise in the corner rods where also the power is initially the highest (due to the greater thermal neutron flux).

Indeed correcting the errors for the local power in the rod shows that the relative errors are all below 2 %, this can be seen in Figure A.3. The maximum absolute errors in the middle slice without rods in the corners are smaller, as are the errors later in the cycle at high burnup, see Figure A.4.

⁸Or rather as one variable with an extra 0, 1, ... , n_{pieces} dimension added representing which piece the change in question is along.

⁹Additional logical constraints then had to link the total number of Gd rods to individual true/false values.

Bottom Slice, 0 GWd/tU

-0.008	-0.004	-0.005	-0.007	-0.011	-0.012	-0.006	-0.002	-0.006	-0.011
-0.004	0.092	0.009	-0.008	-0.007	-0.007	0.039	0.007	0.05	-0.008
-0.005	0.009	0.03	-0.004	-0.013	-0.01	0.001	0.005	0.007	-0.005
-0.007	-0.008	-0.004	-0.009	-0.012	-0.009	-0.011	-0.003	0.037	-0.01
-0.011	-0.007	-0.013	-0.012	0	0	-0.013	-0.01	-0.009	-0.009
-0.012	-0.007	-0.01	-0.009	0	0	-0.015	-0.008	-0.008	-0.01
-0.006	0.039	0.001	-0.011	-0.013	-0.015	-0.009	-0.006	-0.003	-0.011
-0.002	0.007	0.005	-0.003	-0.01	-0.008	-0.006	0.029	0.009	-0.005
-0.006	0.05	0.007	0.037	-0.009	-0.008	-0.003	0.009	0.092	-0.005
-0.011	-0.008	-0.005	-0.01	-0.009	-0.01	-0.011	-0.005	-0.005	-0.007

Figure 5.5: Pin power errors at 0 burnup, CPLEX-value minus true PHOENIX value after a step-1 optimization. The 10 BA placements coincide with the green areas of large positive errors.

Bottom Slice, 7.5 GWd/tU

0.001	0.003	-0.002	0.003	-0.001	-0.001	0.001	0.001	0.004	-0.001
0.003	-0.004	0.001	-0.003	-0.002	-0.003	-0.006	0.002	0.001	0.003
-0.002	0.001	-0.006	0.002	0	0	-0.001	0	-0.005	0.002
0.003	-0.003	0.002	0.004	0	-0.003	0.001	0.001	-0.005	-0.002
-0.001	-0.002	0	0	0	0	0.004	0.001	0	0.003
-0.001	-0.003	0	-0.003	0	0	0.002	0.001	0	0.002
0.001	-0.006	-0.001	0.001	0.004	0.002	0.003	-0.002	0.002	0.004
0.001	0.002	0	0.001	0.001	0.001	-0.002	-0.006	-0.004	0.002
0.004	0.001	-0.005	-0.005	0	0	0.002	-0.004	-0.003	0
-0.001	0.003	0.002	-0.002	0.003	0.002	0.004	0.002	0	0

Figure 5.6: Pin power errors at 7.5 GWd/tU burnup, CPLEX-value minus true PHOENIX value at the burnup point where maximum k_{∞} appears. This late in the cycle the errors from BA introduction have declined considerably (compared to Figure 5.5, note that the color scales in the figures are not the same and that absolute error magnitudes have decreased).

The overall behavior of pin power approximations is acceptable and the strategy of summing contributions from every rod upon every rod is working, albeit requiring a lot of memory for perturbation matrices.

5.2.3 R-factor

The R-factor calculations proved to be the hardest part of the project to evaluate, errors proliferated in the codes and many weighting factors interplayed to obscure the source of every error. It was time-consuming work to obtain a licensed software, DOFACT, of the right version that could calculate the R-factor. A MATLAB script was also written to do these calculations and extend the cases for which R-factors could be estimated to those with relative power distributions from other sources than the 2D lattice code. Comparisons between R-factor values from simplex approximations, DOFACT and MATLAB were done to gauge the accuracy of optimization.

Bottom Slice, 0 GWd/tU

-0.022	0.003	0.008	0.001	0.001	0.002	0.001	0.008	0.003	-0.022
0.003	-0.001	0.000	-0.001	0.003	0.003	-0.001	0.000	0.000	0.004
0.008	0.000	0.002	0.002	0.000	0.000	0.003	0.001	0.001	0.009
0.000	-0.001	0.002	0.000	-0.003	-0.004	-0.001	0.003	0.000	0.001
0.001	0.003	-0.001	-0.003	0.000	0.000	0.000	0.000	0.004	0.002
0.001	0.002	0.000	-0.004	0.000	0.000	0.001	0.000	0.003	0.002
0.000	-0.001	0.002	-0.001	0.000	0.001	-0.002	0.002	0.000	0.000
0.007	0.000	0.000	0.003	0.000	0.000	0.002	-0.001	0.001	0.007
0.002	-0.001	0.000	-0.001	0.003	0.003	-0.001	0.000	-0.001	0.002
-0.023	0.003	0.008	0.000	0.000	0.001	0.000	0.007	0.002	-0.023

Figure 5.7: Pin power errors at 0 GWd/tU, CPLEX-value minus true PHOENIX value after a step-2 optimization. This second optimization step exhibits smaller errors than the first step seen in Figure 5.5, notable however are the corner rods where the approximation yields comparatively large errors.

The result of an optimization on the R-factor can be seen in Figure 5.8.

The four cases in Figure 5.8 are R-factor from DOFACT (green), R-factor from CPLEX approximation (red), R-factor from MATLAB script using 2D lattice code pin powers (blue) and R-factor from MATLAB script using linearized pin powers.

The inclusion of the MATLAB values is a form of sanity check: to assure that the right version of the DOFACT program was used corresponding to the correlation equations provided for simplex approximations. By comparing the values from CPLEX and MATLAB using pin-powers linearized in CPLEX one can see if potential errors arise all across in the linearization scheme or only in the specific R-factor calculations. In Figure 5.8 the four curves follow each other nicely and the error in the CPLEX R-factor is on average smaller than the R-factor analytically calculated from the linearized pin power approximated by CPLEX, this could be due to larger linearization errors in P than R .

An optimization with a slightly different objective run on the same perturbation matrices from the same reactor did however show a different behavior. Figure 5.9 demonstrates how the R-factor approximated by CPLEX differs significantly from the true DOFACT value. By looking closely at which rods exhibited the maximum values in these two curves it was found that the limiting rod in CPLEX's implementation was not the same as the one in DOFACT and that the limiting rod in CPLEX was one that had a higher additive constant than its neighbors but lower pin power.

Minimum-value simplification

The linearization scheme for the bundle R-factor was thus found to be faulty in a non-trivial manner. The equations for the R-factor, 1.4 to 1.7, were expressed in the introduction and of special importance is the neighbor weighting functions like

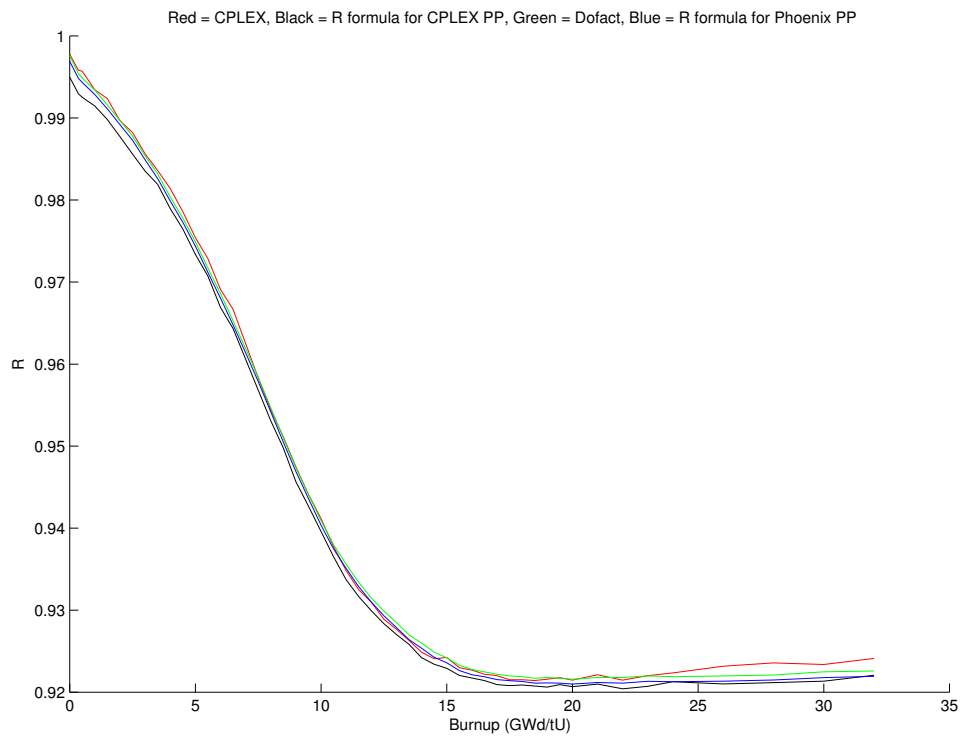


Figure 5.8: R-factors: Red = CPLEX, Black = R formula for CPLEX approximated pin powers, Green = DOFACT, Blue = R formula for PHOENIX pin powers.

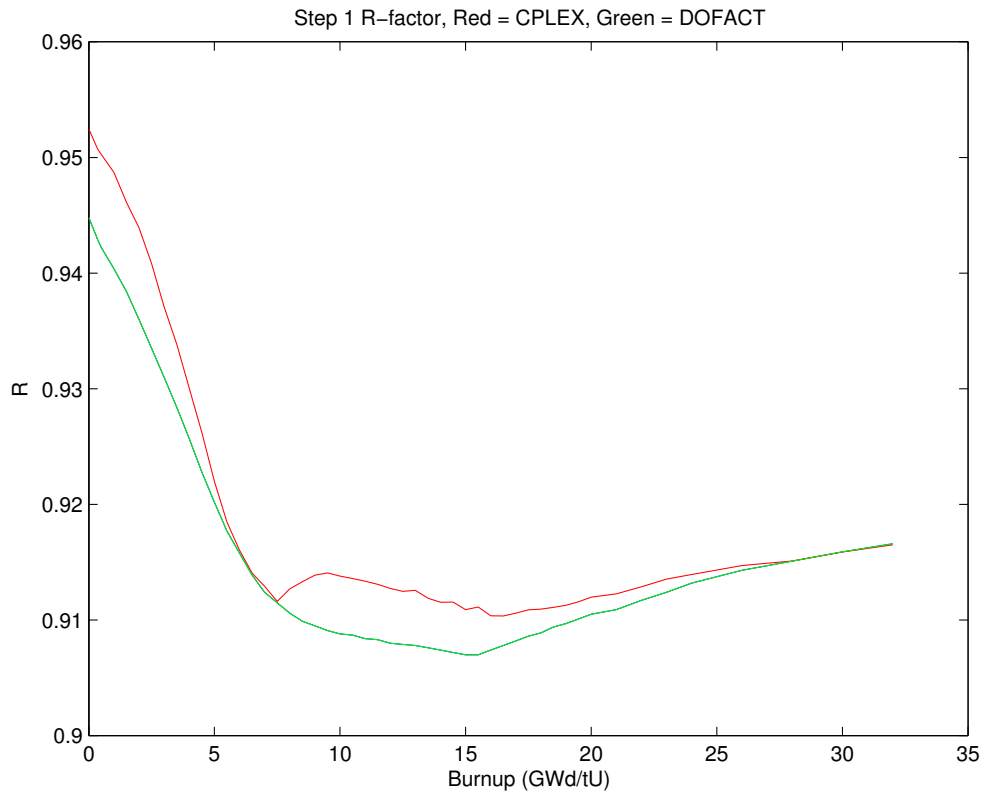


Figure 5.9: R-factors: Red = CPLEX, Green = DOFACT.

5.2 (same as Equation 1.6).

$$SJ_j = \min \left(\sum_{j=1}^{n_j} WJ_j \cdot (r_{j,s,z})^{1/2}, (r_{i,s,z})^{1/2} \cdot \sum_{j=1}^{n_j} WJ_j \right) \quad (5.2)$$

In this equation SJ_j is an input term used in the calculation of the R-factor (Equation 1.4), $r_{j,s,z}$ is the relative power of a node (Equation 1.5) and WJ_j is a weighting factor for neighboring rods that varies by the fuel and the dryout correlation used. Equation 5.2 includes a minimum value calculation, a non-linearity. In words it says: if a node has a higher relative power than a weighted average of its neighbors' powers, choose the first sum and if it is lower than the weighted average of its neighbors, choose the second. The physical reason for this is to limit the effect of high isolated pin powers in Equation 1.4 to make the model comply better with observed CPR performance.

This nonlinearity was eliminated in the existing model by arguing that the interesting parameter is the limiting bundle R-factor R-max and this value will be found in a node that is hotter than its neighbors, otherwise it wouldn't be R-limiting. Thereby the first sum should always be chosen as long as minimizing R-max remains the objective. This had been stated in several reports throughout the project but it can be proven faulty by noting that the R calculation in Equation 1.4, apart from nodal power, also has a nodal dependency on the additive constants I_j .

In any power distribution there will be a number of rods where the relative power is lower than average power in the neighboring rods, situations will therefore occur when the simplex algorithm overshoots the actual R value. The rod in question gets a false contribution to its R-factor by using the neighbor average pin power function instead of the lower value of its own. Since the minimum-function occurs separately for side and diagonal neighbors; errors can arise in both functions independently. Figure 5.10 shows how this approximation is very inaccurate for low-power rods, such as BA rods early in the cycle. These large *overestimates* can be explained by the minimum value simplification but any *underestimates* as can be seen to a smaller degree must have other causes arising from the rest of the linearization simplifications.

This large error in very low-powered rods does not affect the interesting parameter R-max. However the original approximation is faulty and dangerous since it rather often happens that a node with lower power than its neighbors can be R-limiting. This is due to the influence of additive constants describing individual fuel rod dryout sensitivities¹⁰. A fuel rod position having a greater additive constant than its neighbors can be R-limiting even at a lower pin power than the weighted average of its neighbors. Some areas of the fuel are especially sensitive to this as they have significantly larger additive constants than their neighbors.

The result of this error is what is seen in Figure 5.9 where the true maximum R-factor is compared to the approximated value. The part of the error attributable to the faulty elimination of the min-value clause is by nature positive so the actual R-factor will be lower than the optimizer suggests (other errors positive and negative can arise from other parts of the linearization). The fact that this is conservative does however not eliminate the problem, since the error will vary between designs and the simplex algorithm won't find the best designs with the lowest R-factor.

¹⁰The exact values of these additive constants are not important, the significant detail is only that their values within the sub-bundle vary enough to enable significant errors in the R-approximation.




Figure 5.10: R-factors by rod number for an initial state of 0 burnup: Red = CPLEX, Green = DOFACT. Note the large positive discrepancies in low power nodes where BA is present (shown by vertical lines).

The minimum value non-linear equations can actually be implemented in CPLEX but subsequently optimizing R-max then involves finding a maximum value amongst many minimum values, which is hard and requires branching for every case [25]. An attempted optimization using this method yielded a worsened actual (calculated by DOFACT) R-factor than the simplification, after 16 hours of computation instead of the already substantial 2 hours used by the simplified method.

Omitting R-factor approximations

As has been demonstrated in the previous sections, approximating k_{∞} and F_{int} for use in a linear optimizer is quite a lot easier than linearizing R. Linearizing the first two parameters from PHOENIX outputs brings errors, especially for large changes in enrichment or Gd concentration, but the R factor is itself a derivative of pin power determined by non-linear equations. These non-linearities are the square-roots and discrete minimum conditions in the dryout correlation. Approximating this function to accommodate it in CPLEX brings with another layer of linearization errors, and computational challenges.

Late in the project an alternative approach was evaluated. This was that the R-factor depends on the pin power distribution in such a way that there exists at least one pin power distribution that will produce a uniform R-factor across the bundle. To find this power distribution for a geometry of interest a script was developed that calculated the R-factor for a uniform power distribution and then gradually increased power in low R nodes and decreased in high R nodes, all the while maintaining the average relative power at 1. After several loops the power distribution stabilized at one corresponding to a flat R distribution, this was done

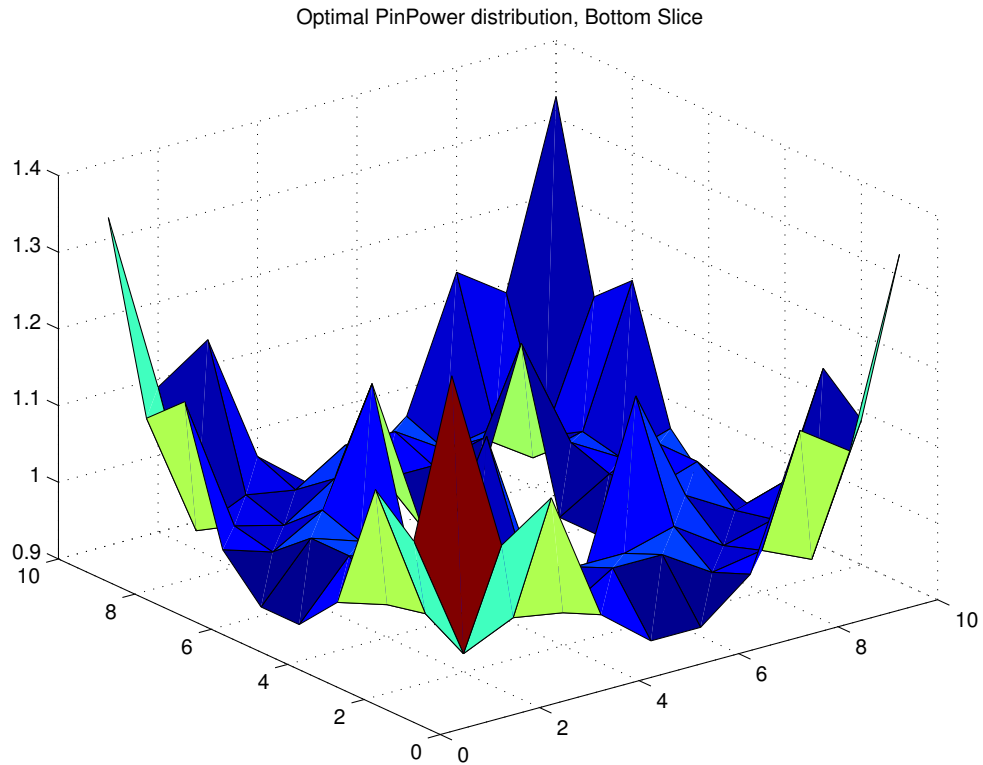


Figure 5.11: A pin power distribution that generates a flat R-factor of 0.8864, bottom segment.

in conjunction for all three axial slices and the bottom slice component can be seen in Figure 5.11.

Whether this distribution is unique, or whether the uniform R-factor associated with it has a lower maximum value than all other distributions (if or not there exists a distribution with a lower R factor in all nodes), cannot be said with absolute certainty because of the many logical conditions in the equation. This was however investigated using stochastic perturbations and despite trying several thousand alterations none yielded a lower maximum R-factor. The R factor obtained with this method can therefore at least be said to be very good with its maximum of 0.8864 and approaching its associated power distribution can be a feasible method of R optimization. An issue with rod-by-rod target compared to a full evaluation of R including neighboring rod effects is that the margin present in neighboring rods cannot directly be utilized, the rod with the largest excess relative pin power might be "saved" by lower values in neighboring rods so that it is not actually the R-limiting rod in the bundle. An optimization targeting relative power deviations will however target this rod for decreased enrichment instead of the actual R-limiting rod, thus hampering the efficiency of the optimization. A graphical representation of deviations from optimal power distribution (stochastically generated in MATLAB) and associated R-factors is found in Figure 5.12. The X-axis represents the maximum deviation from the optimum pin power distribution, this is a parameter that a simplex method can optimize against, thereby bringing its value as low as possible. An opti-




Figure 5.12: A set of stochastically perturbed pin power distributions graphed by their maximum pin power deviation and their R-factor.




Figure 5.13: A set of stochastically perturbed pin power distributions graphed by their maximum axially weighted pin power deviation and their R-factor.

mization on deviation can be likened to a vertical line in Figure 5.12 being brought to the left, it will not see any difference between designs along such a vertical line, but they might have very different R-values. Since the R-factor in a rod is a result of an axial weighting from its slices the deviation from optimum can be corrected for the lengths of the rods and their weight in the CPR correlation (top slices tend to be weighted more strongly than bottom ones). The correlation could be improved even more by computing a sum of deviations in the axial segments of every rod. The optimizer could then compensate between segments of a rod allowing a colder-than-optimal segment of a rod to balance a hotter-than-optimal one. Figure 5.13 shows the deviations in Figure 5.12 corrected for rod-wise axially weighted sums of deviations. This should increase the correlation found in figures 5.12 and although the distribution gets a bit more concentrated the effect is marginal in these purely stochastic deviations (an optimizer can however seek out the cases where positive and negative cancel each other). This sum can, just as individual deviations, be targeted for minimization in the simplex implementation; the optimizer will then indirectly minimize the R-factor.

For an optimization strategy targeting relative power distribution to be perfectly effective there needs to be a bijective correspondence between the properties of maximum pin power deviation and maximum R-factor. As can be seen in figures 5.12 and 5.13 this is not the case (at least partly due to not utilizing the margin from neighboring rods that was previously explained) but it is clear that small power

deviations have a smaller distribution of corresponding R-factors. Furthermore it is possible (although not in this project's form of implementation) to post-process a number of the best solutions from the optimizer and calculate their real R-factor at a low computational cost, thus choosing the best design from a set with similar power deviations.

Implementing this optimal power distribution $P_{z,i}^*$ in the CPLEX optimizer was done by defining a decision variable as in Equation 5.3.

$$(P_{z,i,m} - P_{z,i}^*) - dev < 0 \quad (5.3)$$

Where m is the burnup step considered for minimizing R , n_z is the number of slices and dev is the decision variable targeted for minimization. Depending on constraints this method could produce solutions more or less close to the optimal value. With a correction, such as in Equation 5.4, the optimization could account for different weighting factors between axial slices (but not the interaction between different axial segments of the rod) and target the largest .

$$(P_{z,i,m} - P_{z,i}^*) \cdot Wz_z - dev < 0 \quad (5.4)$$

This manner of R-optimization yielded a substantial improvement in optimization speed compared to R-approximation, but the R-factor actually worsened, see Table 5.2.3. To make the optimization to R from relative pin power distribution as good as possible the minimization target was set as the weighted sum of the axial deviations in each rod, Equation 5.5, this implementation corresponds to the distribution in Figure 5.13.

$$\left(\sum_z^{n_z} (P_{z,i,m} - P_{z,i}^*) \cdot Wz_z \right) - dev < 0 \quad (5.5)$$

One could expect that excess power in the bottom of the rod could then be offset by a lower power in the top, or vice versa, paving the way for a lower R-limitation. Although this assumption turned out to be true the time required to optimize it also exploded: see Table 5.2.3.

Optimization target	Tolerance	Time required	Deviation	Real R
Equation 5.3	0.01 %	3 h 59 min	0.0747	0.9130
Equation 5.4	0.01 %	47 min	0.026 ¹	0.9166
Equation 5.5	0.1 %	32 h 51 min	0.0463	0.9056
R approximation	0.01 %	4 h 30 min	-	0.9093

¹ 0.12406 without weighting factor, as figure 5.12 is displayed.

Table 5.4: The two pin power strategies compare to the R approximation strategy, note that the deviations are different since the second strategy computes the sum of three deviations in a rod. The deviation values correspond to vertical lines in figures 5.12 and 5.13 and corresponding R-factors should be possible to find on this line.

The comparison in Table 5.2.3 to the design obtained by R approximation shows that the pin power deviation minimization can indeed generate a superior design but at a high cost in calculation time. The reasons for this have not been firmly

established but it could be due to the greater number of variables considered in the optimization.

In the calculations of Table 5.2.3 several constraints applied to the optimization besides the optimization target function; one of these was a constraint on LHGR (see Section 1.2.3): effectively one on F_{int} on the bundle level. This constraint depended on burnup and in the optimization interval it was as low as 1.19 while the optimum corner rod power was 1.3276. This constraint inhibits the optimization and makes it harder to reach optimal power profile in Figure 5.11. By forcing rods to remain colder than their optimum it is possible that the coupling between optimizing power profile and optimizing R-factor weakens. A possible way of improving this is to incorporate the limitation on LHGR in the generation of an optimal power profile, thus this profile will depend on the constraints and needs to be re-generated if the constraint or targeted burnup interval changes. This condition did not however change the optimum power profile much outside the corner rods, and the "penalty" to the R-factor was small as it rose from 0.8864 to 0.8870. It became clear that including the constraint did not however improve the optimization, the optimizer reached $R = 0.9058$, or just above the previous best value, but CPLEX was not even approaching termination after 48 hours and 7583022 iterations so the idea was discarded.

5.3 Optimized layout results

The previous sections concerned the methods used in optimization and the improvements that could be done to these. The questions of errors prevalent in the CPLEX optimization are important but the main objective or the optimization framework is to speed up the fuel design process and therefore a study comparing automatically optimized designs to those made by nuclear engineers using experience, trial and error was conducted. The method and the cases for this study were described in Section 3.3. The alternative methods for calculating the R-factor discussed in Section 5.2.3 were not implemented in these optimizations as a time-efficient method for this had not been found.

5.3.1 BA rod position optimization

The first step in these optimizations was to generate a reference design to start the process from. This reference design had previously been one with maximum-enriched in all non-peripheral rods but since it was found that large deviations in enrichment increased the errors of the linearization model the reference was instead calculated to be as close as possible to the desired enrichment. For the short-cycle reactor this was done by simply taking the design from the manual design in Figure 5.14, replacing its BA rods with pure uranium and adjusting the layout so that all interior rod positions had the same enrichment and total bundle average enrichment was preserved. This design shown in Figure 5.15 was then perturbed with Gd in concentrations of 2, 3, 4 and 6 % in every permitted position. This amounted to 116 designs per bundle segment.

The procedure was identical for the long-cycle reactor although the reference design was instead Figure 5.16 and Gd perturbation concentrations were 4, 6 and 8 %.

2.00	2.00	3.40	3.40	2.80	2.80	3.40	3.40	2.80	2.00	3 % Gd
2.00	3.20	4.40	4.40	3.80	3.80	4.40	4.40	3.20	2.80	
3.40	4.40	4.40	4.40	3.40	3.40	4.40	4.40	4.40	3.40	
3.40	4.40	4.40	3.20	3.80	3.80	3.20	4.40	4.40	3.40	
2.80	3.80	3.40	3.80			3.80	3.40	3.80	2.80	
2.80	3.80	3.40	3.80			3.80	3.40	3.80	2.80	
3.40	4.40	4.40	3.20	3.80	3.80	3.20	4.40	4.40	3.40	
3.40	4.40	4.40	4.40	3.40	3.40	4.40	4.40	4.40	3.40	
2.80	3.20	4.40	4.40	3.80	3.80	4.40	4.40	3.20	2.80	
2.00	2.80	3.40	3.40	2.80	2.80	3.40	3.40	2.80	2.00	

Figure 5.14: The manually optimized design for the short-cycle reactor, bottom segment. Numbers represent enrichment levels, white boxes are ordinary uranium rods and blue boxes are BA rods with 3.2 % enrichment and 3 % Gd.

2.00	2.00	3.40	3.40	2.80	2.80	3.40	3.40	2.80	2.00
2.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	2.80
3.40	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	3.40
3.40	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	3.40
2.80	4.00	4.00	4.00			4.00	4.00	4.00	2.80
2.80	4.00	4.00	4.00			4.00	4.00	4.00	2.80
3.40	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	3.40
3.40	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	3.40
2.80	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	2.80
2.00	2.80	3.40	3.40	2.80	2.80	3.40	3.40	2.80	2.00

Figure 5.15: Gd-free reference design for 1st step of optimization of the short-cycle reactor, bottom segment. Average enrichment in the bundle segment is 3.63 %

2.80	2.80	3.80	3.80	3.80	3.80	3.80	4.60	3.80	2.80
2.80	3.80	4.95	4.2 & 7 %	4.95	4.95	4.95	4.95	4.2 & 5 %	3.80
3.80	4.95	4.2 & 7 %	4.95	4.95	4.95	4.2 & 7 %	4.95	4.95	4.60
3.80	4.2 & 7 %	4.95	4.95	4.95	4.95	4.95	4.2 & 7 %	4.95	4.60
3.80	4.95	4.95	4.95	-	-	4.95	4.95	4.95	3.80
3.80	4.95	4.95	4.95	-	-	4.95	4.95	4.2 & 7 %	3.80
3.80	4.95	4.2 & 7 %	4.95	4.95	4.95	4.95	4.2 & 7 %	4.95	4.60
4.60	4.95	4.95	4.2 & 7 %	4.95	4.95	4.2 & 7 %	4.95	4.2 & 7 %	4.60
3.80	4.2 & 5 %	4.95	4.95	4.95	4.2 & 7 %	4.95	4.2 & 7 %	4.20	3.80
2.80	3.80	4.60	4.60	3.80	3.80	4.60	4.60	3.80	2.80

Figure 5.16: The manually optimized design for the long-cycle reactor, bottom segment. Numbers without percentages represent enrichments, colored boxes represent BA rods with percentages stating Gd concentration.

The perturbations generated a set of step 1 S-matrices which were run in the CPLEX optimization suite. The optimization was subject to numerous constraints including those detailed in Table 3.1 and properties derived from the final designs, namely the peak k_{∞} burnup, b and the BA effect on initial k_{∞} , a . Table 5.5 shows the details of the step 1 optimizations. The documentation for manual design of the short cycle reactor stated that the design was optimized towards minimizing F_{int} . Such a strategy can be implemented in the simplex optimizer but the resulting designs proved very poor with regard to R and it proved more efficient to include R in the optimization target than trying to constrain it under fixed values, a large improvement in R-factor was then obtained from a much smaller sacrifice in F_{int} .

For the short-cycle reactor a CPLEX approximated R-factor along with its real DOFACT value can be seen in Figure 5.17, the CPLEX approximation shown in red is the curve approximated by CPLEX for its winning design. An evaluation of this suggested design shows that its real R-factor is actually lower than the CPLEX approximation through the majority of the cycle.

That the CPLEX approximation in Figure 5.17 for R overshoots the real value can be explained by the errors observed from the faulty but necessary simplification of the minimum function SJ and SK in the R approximation.

5.3.2 Uranium enrichment and Gd concentration optimization

The BA placements from step 1 were taken as references for a step 2 optimization where BA positions were fixed and perturbations ran for every position in the grid. Uranium perturbations were done at 2.4 and 4.95 % ^{235}U in all positions and Gd perturbations in nodes containing BA. Gd perturbations were done at slightly different percentages to account for the differences between 12 and 24 month cycles:

Step 1 Procedure	Short Cycle Reactor	Long Cycle Reactor
Calculation time for S-matrices	15 minutes	15 minutes
CPLEX target function	Sum of max R and F_{int} from 10 GWd/tU onwards	R factor late in cycle
Notable constraints	LHGR, k_{∞} profile	LHGR, k_{∞} profile, types of Gd rods
CPLEX optimization time	8 min 25 s	2 min 12 s
Value of CPLEX target	$R=0.91651$, $F_{int} = 1.141$	$R = 0.921033$
Concentrations used (%)	3	6 & 8
Number of BA rods	10	17

Table 5.5: Results of Step 1 Optimization




Figure 5.17: R-factor for the step1 optimization of the short-cycle reactor. Red = CPLEX, Green = DOFACT.

the first using 2 and 4 % (3 % rods were selected by step 1) and the second using 4 and 8 %. The conditions and constraints on the optimization from step 1 were slightly altered to include average enrichment constraints and the different methods for calculating BA effect on initial k_∞ (a^{11}) and burnup at peak k_∞ (b) burnup. A total of 336 (short-cycle reactor) and 352 (long-cycle reactor) PHOENIX jobs were submitted and S-matrices extracted.

Step 2 Procedure	Short Cycle Reactor	Long Cycle Reactor
Calculation time for S-matrices	10 min	10 min
CPLEX target function	Sum of max R and F_{int} from 10 GWd/tU onwards.	R factor late in cycle, around peak k_∞ .
Notable constraints	LHGR, bundle enrichment,	LHGR, bundle enrichment,
CPLEX optimization time	4 h 12 min	15 min
Value of target	$F_{int}=1.113$, $R_{max}=0.928$	0.9104
Enrichments used	2.4, 2.8, 3.2, 3.8, 4.0, 4.4	3.2, 4.0, 4.4, 4.6, 4.95
BA concentrations used	2 %	5, 7 %

Table 5.6: Results of Step 2 Optimization

5.3.3 Performance compared to manual designs

Short cycle reactor

The optimization for the short cycle reactor was targeted at the sum of R and F_{int} since it had been found that the optimizer could yield a large gain in one parameter at a small sacrifice in the other. Although the optimization target in the automatic and manual processes were not the same it became clear that the compound objective led to a good balance between the parameters with improved F_{int} compared to the manual optimization.

The reason for not minimizing the maximum values over the entire cycle is that the severity of poor R and internal form factors depends on the power of the bundle: R-factor couples CPR sensitivity to bundle power and F_{int} is a relative power measure. The values of these parameters are therefore most important when bundle power is high, which in turn is determined by the peak k_∞ burnup value from k_∞ optimization. It is also useful that narrowing the burnup interval over which optimization is done decreases the number of points CPLEX needs to evaluate and therefore the run-time.

¹¹The step 1 optimization uses fixed enrichments and has no BA in the reference. In that case the a -value is simply the initial k_∞ of the reference minus that of a considered design. When enrichments change together with BA in step 2 the value can not be computed in the same manner since the BA-free k_∞ changes too.

2.40	2.40	3.20	3.20	2.80	2.80	3.20	3.20	2.40	2.40	2 % Gd
2.40	3.20	4.00	4.40	4.40	4.40	4.40	4.00	3.20	2.40	
3.20	4.00	3.20	4.40	4.00	4.00	4.40	4.40	4.00	3.20	
3.20	4.40	4.40	3.80	4.00	4.00	3.80	4.40	3.80	3.20	
2.80	4.40	4.00	4.00			4.00	3.80	4.40	2.80	
2.80	4.40	4.00	4.00			4.00	3.80	4.40	2.80	
3.20	4.40	4.40	3.80	4.00	4.00	3.80	4.40	4.40	3.20	
3.20	4.00	4.40	4.40	3.80	3.80	4.40	3.20	4.00	3.20	
2.40	3.20	4.00	3.80	4.40	4.40	4.40	4.00	3.20	2.40	
2.40	2.40	3.20	3.20	2.80	2.80	3.20	3.20	2.40	2.40	

Figure 5.18: Automatically optimized design for the short-cycle reactor.

The short cycle reactor was found to perform well against the manually optimized reference design on several criteria, the internal form factor in Figure 5.20 was kept lower during peak k_{∞} by sacrificing some initial and final margin. For the bottom slice studied the bottom slice internal form factor at 0 burnup increased from 1.106 to 1.186, at peak k_{∞} decreased from 1.11 to 1.079 and at 40 GWd/tU it increased from 1.046 to 1.055.

The R-factor was included in the optimization and this made a mark since it managed to perform better than the manual design throughout the cycle, see Figure 5.21.

3.20	3.20	4.60	4.00	3.20	3.20	4.00	4.95	4.60	4.00
3.20	4.6 & 5%	3.20	4.95	4.6 & 5%	4.6 & 7%	4.60	4.6 & 5%	4.00	4.60
4.60	3.20	4.6 & 5%	4.95	4.60	4.60	4.95	4.95	4.4 & 7%	4.95
4.00	4.95	4.95	4.95	4.40	4.95	4.95	4.95	4.60	4.40
3.20	4.6 & 5%	4.60	4.40	-	-	4.60	4.95	4.6 & 5%	3.20
3.20	4.6 & 7%	4.60	4.95	-	-	4.95	4.95	4.6 & 5%	3.20
4.00	4.60	4.95	4.95	4.60	4.95	4.4 & 7%	4.95	4.95	4.40
4.95	4.6 & 5%	4.95	4.95	4.95	4.95	4.95	4.95	4.6 & 5%	4.95
4.60	4.00	4.4 & 7%	4.60	4.6 & 5%	4.6 & 5%	4.95	4.6 & 5%	4.00	4.60
4.00	4.60	4.95	4.40	3.20	3.20	4.40	4.95	4.60	4.00

Figure 5.19: Automatically optimized design for the long-cycle reactor.

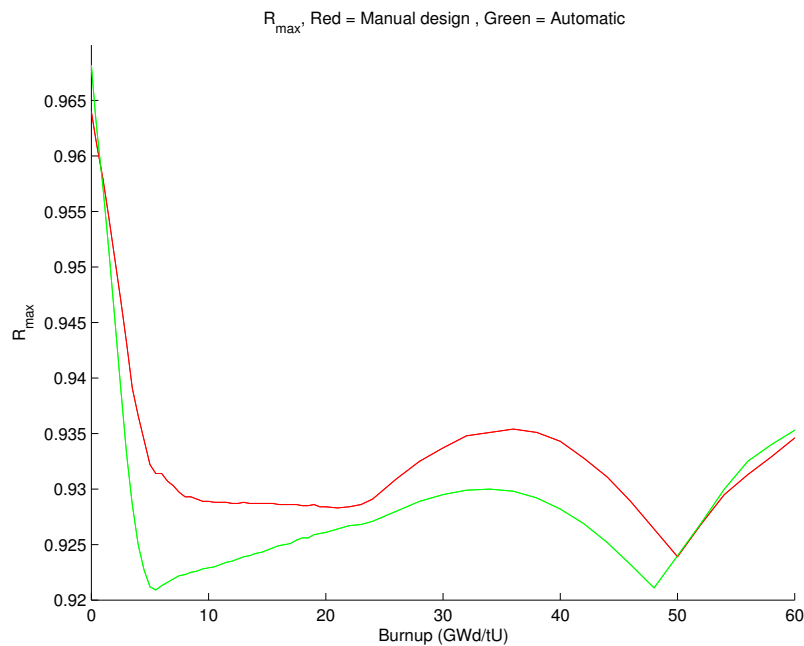


Figure 5.21: Short-cycle reactor maximum R-factor, red = manual design, green = automatic design.

The k_{∞} profile was a bit off from the desired peak k_{∞} burnup as seen in Figure 5.22, this proved to be the result of a combination of lax CPLEX constraints and linearization inaccuracy. The sought profile of the manual optimization peaked at 7 GWd/tU whereas the simplex approximation was calculated to do so at 6.5 GWd/tU¹², the real value for this peak however turned out to be 6.0 GWd/tU after design was evaluated in the 2D lattice code.

¹²A tolerance of 0.5 GWd/tU above or below the sought value was used, this can be omitted




Figure 5.20: Short-cycle reactor bottom slice maximum relative pin power, red = manual design, green = automatic design.




Figure 5.22: Short-cycle reactor k_{∞} profiles for all slices, reference = crosses, CPLEX = solid.

Long cycle reactor

For the long-cycle reactor the optimizer had a harder task complying with the constraints. This manual optimization was well-done and small inaccuracies propagated in the model to impair the results of the automatic optimization. Figure 5.23 shows that the optimized design "wins" in the target interval around 20 GWd/tU, where the R-factor decreases from 0.922 to 0.912; but performs worse in other parts of the cycle, for instance at 0 burnup where it increases from 1.020 to 1.053.




Figure 5.23: Long-cycle reactor maximum R-factor, red = manual design, green = automatic design.

Figure 5.24 demonstrates clearer how the gain in R-performance has come at a large sacrifice in pin power distribution, which is much more unequal than in the manual case: initial F_{int} increases from 1.354 to 1.617 and F_{int} at peak k_{∞} from 1.118 to 1.147.




Figure 5.24: Long-cycle reactor internal form factor from an optimization targeting only R_{max} , red = manual design, green = automatic design.

The result of this optimization showed the importance of balanced constraints, as the maximum relative pin power in Figure 5.24 is quite unacceptable. Attempting to correct the pay-off between F_{int} and R another optimization was therefore run, targeting the sum of R in the target interval and F_{int} (just as in the short-cycle reactor optimizations) over a larger burnup interval, with the R contribution targeted twice as heavy as F_{int} part by defining the target variable t as $t = \frac{F_{int}}{2} + R$.




Figure 5.25: Long-cycle reactor maximum R-factor from an optimization with a compound optimization target, red = manual design, green = automatic design.

The results of this alternative strategy can be seen in figures 5.25 and 5.26. It is clear that no gain in one parameter came without sacrificing the other but it shows to the feasibility of using weighted sums of parameters as optimization targets¹³. The role and potential benefits of unconventional compound objectives is discussed in Section 6.1.4.

¹³If there is time; a greater number of parameters evaluated increases the run-time significantly, especially if they have to be evaluated over longer burnup intervals.




Figure 5.26: Long-cycle reactor internal form factor from an optimization with a compound optimization target, red = manual design, green = automatic design.

For the long-cycle reactor neither the optimization targeting only R nor the one using a compound R and F_{int} target could generate any overall improvements to the manual design in the manner that the short cycle reactor optimization could (at least on a first glance).

Time savings

It should be noted that the manual design in the long-cycle reactor case did indeed require significant design time and represents exactly the kind of case an optimization program is requested for. Actually quantifying the time taken to design a fuel bundle is hard since the process is intertwined with the general fuel bid process but estimates on tough cases range as high as one engineering-week.

Chapter 6

Discussion

6.1 Simplex framework

6.1.1 Benefits of simplex fuel bundle optimization

In the results section a great number of obstacles hindering ideal implementation of a bundle optimization program have been presented and focus has been put on the errors prevalent in the approximations and how to reduce them. A conclusion that has been drawn is that a linear perturbation optimization scheme such as the one examined will be sensitive if it incurs large changes in the design variables (compared to the reference used as input). The investigations of the optimization scheme have shown that these errors are hard to combat and that the optimizations may suffer from them in an unpredictable manner, sometimes more and sometimes less. The relative strength of the optimization program has not however been found to vary between reactors in any significant way, that is, longer cycles are harder to optimize both manually and automatically.

There are of course great potential benefits to be reaped from automatically optimizing bundle designs functionally: better designs can be found with less time being spent on trial-and-error during the bundle design process (a cost estimate for this is found in Section 6.1.4). It should be noted that automatic optimizations requiring a lot of computing time can be beneficial even for nuclear designs where a skilled engineer could have found an equivalent design in a *shorter* time; the key is that the program decreases the total amount of time spent *by the engineer* on the bundle design. This requires a good user interface and integration as well as run-times that are compatible with normal work: i.e. absolutely nothing longer than overnight optimizations. The prototype optimizer studied and developed on is not however at such a stage yet, especially regarding user interface.

Simplex optimization has a beneficial side effect besides actual optimization in that it can determine limits of possible solutions. If a set of parameters, such as limits on internal power peaking factor and number of BA rods, are not compatible the optimization program will in very short time realize it is infeasible and cancel optimization. This feature is not possible in stochastic optimization and it can help the nuclear designer when balancing constraints for bundles to use in a whole core optimization.

An optimization of specific concern to the nuclear company was the idea of optimizing the BA layout to minimize the total number of Gd rods in a bundle

to improve production, logistics and customs expenses. This optimization was attempted by taking a manual design, removing Gd and running a step-1 optimization on the problem. This approach did however not yield any conclusive results due to the strict constraints used, lack of time and that a good implementation of this would require combination of step 1 and 2¹.

6.1.2 Simplex method improvements

Several strategies were concocted to improve the simplex framework available at the beginning of the project. Most of these were investigated in the results section but others were not possible to implement in the scope of this project.

k_{∞}

Several improvements were attempted for k_{∞} . The symmetric perturbations that were implemented proved superior to asymmetric perturbations for the second step of optimization. Using reference cases with carefully adjusted average enrichment proved useful for the optimization. In the second step it was seen that large changes in BA concentration affected the optimization accuracy negatively; that the first step with a correct enrichment profile could better pick BA concentrations which the second step did not have to change could therefore improve the accuracy of the second step and thus the quality of the optimization.

These improvements were implemented into the optimization framework along with more accurate optimization constraints on k_{∞} .

Piecewise linearly constrained variables were interesting to use but made the optimizations too slow to function. A better implementation of these constraints or any other part of the simplex implementation (if it increases the speed of the process) could make these a viable option for improved accuracy.

A shadow-factor was not implemented in the second optimization step since it was expected it would also slow down the process further still. A major concern for step-2 optimizations is however that the k_{∞} profile tends to become shifted horizontally from its real value (creating a faulty estimation of b -value). This is thought to be due to an inability to factor in the group effects of several interacting changes in BA concentration; so a shadow factor accounting for this could be a way to improve the estimation.

The approximation of k_{∞} curves is still inaccurate when BA-concentration changes by large amounts, increasing the concentration in many rods will in PHOENIX calculations shift the peak burnup (b -value) to higher values but the sum of several S-matrix elements will not reproduce this effect. A possible explanation is that the extra gadolinium introduced with one perturbation will be quickly burnt out when all other rods have lower concentrations and the resulting perturbation elements will diminish quickly beyond the reference design b -value, this does not happen when several rods change together and the result will be that the linear approximation puts the peak at a too low burnup. This situation will reverse if an aggregate BA decrease is introduced by the optimization.

¹One would like to run a BA-placement optimization on many enrichment profiles to find a minimum number of BA rods.

Relative pin power distribution

The accuracy of relative pin power approximation was found to behave in a similar manner as that of k_∞ approximation. BA changes typically induced greater errors than enrichment changes, especially when these were clustered. Figure 6.1 shows the burnup evolution of pin power errors, it is found that these diminish strongly as BA becomes burnt out. Due to the matrix nature of pin powers it was harder to visualize the accuracy and implement alternative formulations in CPLEX; thus the behavior of relative pin powers was not as rigorously studied as that of k_∞ .

R-factor

The impact of the previously unknown linearization errors of the minimum-functions in the equations for R have been shown in Section 5.2.3. Allowing this minimum value to be explicitly calculated in the simplex optimization proved infeasible as it greatly increased the optimization time.

Pin power target instead of R-factor approximation The comparison to the design obtained by R approximation shows that the pin power deviation minimization can indeed generate a superior design but at a high cost in calculation time. The reason for this has not been firmly established but the most reasonable idea is that the optimizer when optimizing an approximation of R can omit more variables as it identifies the limiting rod at a burnup time, the same not being possible for the pin power distribution target. It is also striking how summing the deviations from different axial slices drastically increases the calculation time.

Adjusting the target power distribution to account for another constraint on peak rod power did not improve results. This should be expected since the change in the optimum power profile resulting from the LHGR constraint is marginally different in positions other than the limited rods; all the non-limited positions simply shift slightly to higher powers and the gains from recalculating are slim (the optimization only targets relative pin power values *higher* than the optimal distribution).

Although the pin power distribution target relieves the problem of the minimum-value simplification used in the R-approximation it is not without inaccuracies: the analysis of power-distribution errors in Section 5.2.2 showed that these were reasonably small in the optimizations studied. The calculation of deviations from an optimum does however place a much greater demand on accuracy in *every* rod than a simple limitation on F_{int} ; which is the normal case of pin power relevance. The R-factor optimization is also specifically targeted at the peak k_∞ burnup (since this is where the R-factor is most important) where the errors in approximation tend to be the greatest. Figure 6.1 shows the maximum pin power errors approximated by CPLEX in the step 2 optimization leading to the first value in Table 5.2.3: average enrichments have changed by less than a 0.02 of a percentage in all three axial slices and about half BA positions increased from 6 to 7 % with the other half decreasing from 6 to 5 % so the optimization step has not incurred any large changes in composition. This only emphasizes the need for a deviation optimizer to post-process multiple solutions with low deviations to find the best R-factor as such an optimization will have uncertainties both in how great the deviation from optimum is and what R-factor this deviation corresponds to.




Figure 6.1: Maximum pin power errors in CPLEX approximation. Blue = bottom, red = middle, green = top.

6.1.3 Feasibility and performance of optimization

The optimizations on real fuel bids show that while it is feasible to use a simplex method to optimize fuel bundle designs it is hard to naively use an automatic tool and expect to "beat" years of experience in fuel design. It is evident however, that such a tool can generate good designs if the problem is approached with some experience and caution for error in approximations. It also shows that weighing optimization parameters such as different target functions against each other² and adjusting constraints affect the quality of the solution, thus in no way decreasing the need for nuclear design expertise.

Impact of poor approximations

It seems obvious that if the objective of an optimization is to minimize the form factor and the approximation of this property being used by the optimization engine is faulty; then the optimized design will also be faulty. However if the error imposed by the linearization is virtually consistent for all reasonable designs then this error might not limit the optimization tool. After all the goal of the program is to find

²The compound objective used in the optimizations on both the long-cycle and short-cycle reactors were useful in this benchmark as the concern was merely to generate a design that was as "good as possible" from several perspectives. The documentation from the manual designs did not state exactly what constraints had been used in the design, and using the entire F_{int} -curve of a manual design as a constraint gives an unfair advantage to the manual designs since most improvements most come at a (small or large) price in some other parameter or burnup point. When one has hard constraints on parameters, for example a F_{int} -curve, it is logical to optimize a single variable, usually R , and the constrained variable will most likely border its constraint at exactly one point. Thus a compound objective might be better in these kinds of benchmarks than in reality, where constraints are hard and "only" have to be met.

which is the best design; its actual exact properties can be quickly and accurately obtained by 2D lattice code calculations.

Examples of errors affecting optimizations 1. The results in Table 5.2.3 suggest that different errors in power approximation might have affected the values of R , it is unreasonable that a non-corrected deviation should perform better than one corrected for axial weighting factor. 2. After analyzing the results of the long cycle reactor optimization it was found that the Excel file entries for the reference enrichments in 4 of 96 positions had not been entered correctly. This mishap should be expected to have hampered the performance of the optimization since it led CPLEX to calculate differences starting from incorrect assumptions. This might have contributed slightly to the significantly less impressive results generated in this optimization, the difference has however not been computed. These errors were easily created when using the prototype framework since there were many manual steps and no safeguards or reminders for mistakes.

6.1.4 Recommendations and usage guidelines

Some lessons learned through this thesis will be important to keep in mind for further development of simplex fuel bundle optimization:

- The user can make the optimization much stronger by providing good constraints and reference designs. These could partially be computer-generated, to account for average enrichment.
- Automating the comparisons of real bundle parameters from the 2D lattice code to the linear approximations used in the simplex program can yield much better and more thorough understanding of what makes the optimizations work and what doesn't. This was a step that should have been taken early in this project but wasn't due to lack of programming expertise.
- Keeping the errors from above comparisons visible so that all optimizations can be gauged on their accuracy is important. Including a plugin for external R-calculations would also be necessary since these aren't done by the 2D lattice code.
- The time that can be saved with an optimization program is a very important parameter that hasn't been accurately estimated during this project, a good measure of this is essential for decision-making.

During the project a consultant has been working in parallel to create a production version of the optimization program with a new model running on a cluster with significantly faster and more stable S-matrix generation and data-handling. This program could improve the prospects for automatic optimization greatly and potentially include some more accurate models. The problems encountered with the R-factor model are important to keep in mind and continued investigation of implicit R-calculation from pin power distribution, by someone with good optimization knowledge, could be a way forward. This would also future-proof the optimization for fuel types that do not use XL-correlations and thus have significantly different R-factor equations.

Cost-benefit analysis

To sum up the recommendations for the optimization project a crude cost-benefit analysis is useful, the payback time for a project like this is hard to estimate and will partly depend on whether the optimized fuel bundles will be surpass and be more desirable than manual designs. In this project the manual designs were however not surpassed in quality, but it seems viable that an optimizer could eventually design equivalent fuel.

For equivalent fuel a rough calculation of payback time for the nuclear company can be made with the following assumptions:

- It will take one engineering year to develop an optimizer from idea to production and support it over its lifetime. Cost of additional software, such as commercial simplex optimizer license, costs an equivalent 50 engineering hours.
- Fuel bundle nuclear design generally takes between 4 and 40 hours [24], so a design time of 10 hours for 12-month and 25 hours for 24-month cycle reactors is assumed.
- A good optimization program can cut the fuel bundle nuclear design time by three-quarters.
- There are 84 BWRs worldwide [16], assumed to be evenly divided between 12 and 24 month cycle operation³.
- The BWR fuel market share of the company is 30%.
- Long-term internal rate of return for large nuclear companies is 5 % [27].

The payback time from time savings alone would then be 15 years. Uncertainties are many, such as interest rates and the future size of the BWR market, but the most important ones are the actual time that will be saved in design and time needed to develop and maintain the program. The profitability of automatic optimization will however not only come from time savings and market share (through superior designs) but will depend on how its positive effects are measured. Positive externalities from automatic optimization could include better integration with loading pattern optimization, reduced number of Gd concentrations or enrichments and potentially better customer satisfaction from quicker design processes.

6.2 Project evaluation

6.2.1 Goal attainment

The overall goals of the project were: investigating optimizations on nuclear fuel bundles, coming up with and evaluating possible improvements to simplex optimization code and determining the usefulness and benefits of automatic optimization on nuclear fuel bundle designs. These main tasks have all been thoroughly labored and

³If 24-month cycle optimization is assumed to exactly twice as hard as 12-month, the relative abundance of reactor cycle lengths does not matter.

although many results were either expected or erratic there were several notable findings that put the light on some unknown difficulties and can enable improvement of the framework. Although the project as a whole has not generated any dramatic breakthroughs in BWR bundle optimization it has been shown that an optimization scheme based on linear perturbations can generate useful fuel assembly designs when the user is aware of the limits introduced by linearization.

6.2.2 Method used

The method used in the project, of focusing strongly on improvements and error patterns in the optimization framework, was a product of the nature of both the project and the equipment used. The fundamentals of the optimization framework were already in place but poor organization of it (see Figure 6.2) meant that it was very hard to work with. The focus on sources of various errors was therefore necessary to being able to evaluate automatic optimization's performance against manual designs. Furthermore the equipment available proved to be a major challenge since licensing issues limited the use of the CPLEX optimizer to a very old version running on a very old computer. This meant that around half the project was spent not able to do full-cycle optimizations since the calculation time became far too long. After a long wait a new computer with CPLEX was introduced and the optimizations duly became up to 300 times faster as explained in Chapter 4.2.3. A smaller focus on improving optimizer performance could have enabled other questions to have been investigated and thus a more interesting paper.

6.2.3 Areas for further investigation

The project has borne several ideas that have not been possible to try out but could very well be interesting topics of related research.

Stochastic optimization

In Section 1.3.2 a brief explanation of stochastic optimization algorithms was given, it was also argued that these schemes such as these are hampered by the inability to estimate if a set of constraints is at all possible to fulfil, or to provide any clue whether the optimization will run for X seconds or Y hours.

A more thorough investigation including an actual implementation on this problem would however be greatly beneficial to understanding the problem and finding the optimal optimization strategy for BWR bundle designs. The problems surrounding linearizations of the bundle parameters in Section 5.2 show that there could be a real benefit in calculating bundle parameters such as R and k_∞ with 2D lattice codes for every evaluation instead of relying on linearizations. A stochastic approach would also require a skilled developer to create a good initial population of bundles, as only a minuscule proportion of random designs would meet constraints on parameters such as k_∞ . Although a stochastic scheme could work without any "smart" reference input it might require a very long calculation time to find "the right track" since evaluating a design takes around 10 seconds in one processor core.

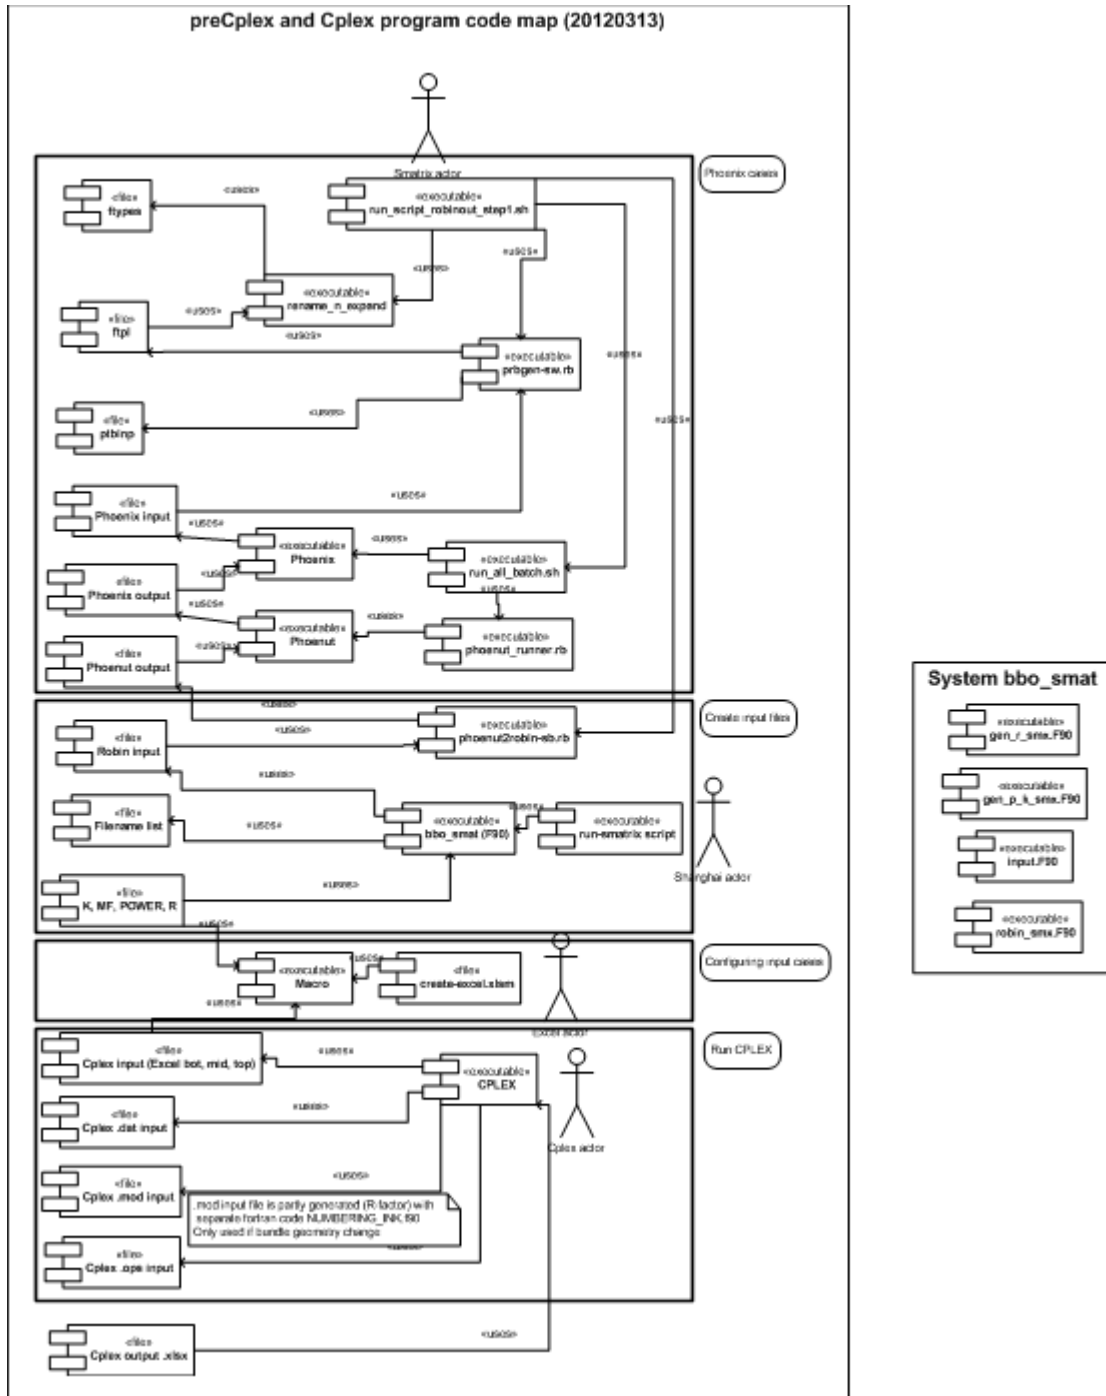


Figure 6.2: A representation of the state of the S-matrix generation for the optimization framework at the beginning of the project. Process was done three times for the different components of optimization. Names of modules are not important but the number of these and their interconnectivity posed a problem during the project, stick men represent instances of user intervention. [24]

Loading-pattern optimization

Integration into automatic whole core loading-pattern (LP) optimization is currently infeasible, but coupling synergies exist. As was mentioned in the introduction, nuclear fuel optimization most often constitutes two separate steps: fuel bundle design and core loading pattern, of which only the first is addressed in this project. There are several reasons for the separation of tasks: The problem of deciding which rod to use at every lattice position in every bundle in a reactor, laterally and axially, is very, very large. The program used to generate perturbation matrices addresses only 2-D single-bundle physics. A completely different interface would be required to generate perturbations on core-level. An interesting topic for research would be to start from a whole core loading pattern problem and use an optimization to provide constraints on the fresh bundles that should be entered. From there another optimization stage could create those bundles and maximize a whole-core parameter, such as CPR.

6.3 Project conclusion

The project has provided an interesting insight into nuclear power and optimization. The nature of the work has meant that the optimization and programming aspects of the problem have constituted the majority of the work with the nuclear physics mainly providing the backdrop. A clearer plan at the onset of the project and greater awareness and criticism before starting it could have enabled more original work to be done and clearer conclusions to be drawn.

Chapter 7

Bibliography

- [1] *Status report 97 - Advanced Boiling Water Reactor (ABWR)*.
IAEA, 2011
<https://aris.iaea.org/sites/..%5CPDF%5CABWR.pdf>
Retrieved on 2014-06-07
- [2] *AP1000 Design Control Document*.
NRC, unknown year
http://www.nrc.gov/reactors/new-reactors/design-cert/ap1000/dcd/Tier%202/Chapter%205/5-1_r15.pdf
Retrieved on 2014-06-07
- [3] *Nuclear Data for safeguards*.
IAEA unknown year
<https://www-nds.iaea.org/sgnucdat/a6.htm>
Retrieved on 2014-06-08
- [4] Zoran Stosic, 1999
Study on Thermal Performance and Margins of BWR Fuel Elements.
7th International Conference on Nuclear Engineering.
<http://www.jsme.or.jp/monograph/pes/1999/ICONE7/PAPERS/TRACK07/FP7283.PDF>
Retrieved on 2014-06-08
- [5] *Imagine the Universe*
Dr. James C. Lochner
2005
<http://imagine.gsfc.nasa.gov/docs/teachers/elements/imagine/Cosmic.pdf>
Page 16.
Retrieved on 2014-06-08.
- [6] *Chart of Nuclides*
National Nuclear Data Center, Brookhaven National Laboratory.
<http://www.nndc.bnl.gov/chart/help/index.jsp#colorcode>
Retrieved on 2014-06-08.
- [7] *Uranium 2011: Resources, Production and Demand*. OECD & IAEA
2012

ISBN 978-92-64-17803-8

<http://www.oecd-nea.org/ndd/reports/2012/uranium-2011-exec-summary.pdf>

- [8] *17x17 Next Generation Fuel (17x17 NGF) Reference Core Report*
Barsic, Conner et. al.
March 2008.
<http://pbadupws.nrc.gov/docs/ML0810/ML081010603.pdf>
Retrieved on 2014-06-09.
- [9] *Upgrade of the FRIGG test loop for BWR fuel assemblies*
Olov Nylund
ABB Review 6/1997
[http://www05.abb.com/global/scot/scot271.nsf/veritydisplay/a3a21d156256c229c1256ebd00308051/\\$File/47-53%20ENG%209706.pdf](http://www05.abb.com/global/scot/scot271.nsf/veritydisplay/a3a21d156256c229c1256ebd00308051/$File/47-53%20ENG%209706.pdf)
Retrieved on 2014-06-09.
- [10] *Fuel Design Evaluation for ATRIUM TM 10XM BWR Reload Fuel AREVA NP Inc.*
Submitted to the NRC April 2010
<http://pbadupws.nrc.gov/docs/ML1011/ML101100643.pdf>
- [11] Westinghouse 2003.
Westinghouse BWR ECCS Evaluation Model: Supplement 3 to Code Description, Qualification and Application to SVEA-96 Optima2 Fuel. Page 14.
<http://pbadupws.nrc.gov/docs/ML0312/ML031220224.pdf>
Retrieved on 2014-06-09
- [12] Personal Source: Sven-Birger Johannesson
Principal Engineer, BWR Core Design
June 2014.
- [13] Butler, Walter R.
GENERAL ELECTRIC BWR THERMAL ANALYSIS BASIS (GETAB): DATA, CORRELATION AND DESIGN APPLICATION
Submitted to NRC September 1974
<http://pbadupws.nrc.gov/docs/ML1022/ML102290144.pdf>
- [14] Yamamoto, A.
A Quantitative Comparison of Loading Pattern Optimization Methods for In-Core Fuel Management of PWR
Journal of NUCLEAR SCIENCE and TECHNOLOGY, Vol. 34, No. 4, p. 339-347 (April 1997)
- [15] Asgari et al.
Application of the N-StreamingSM Concept to Peach Bottom 2 Cycle 17
PHYSOR-2006, ANS Topical Meeting on Reactor Physics
<http://mathematicsandcomputation.cowhosting.net/PHYSOR-2006/C154.pdf>
Retrieved on 2014-06-09

- [16] Euronuclear.org
Nuclear power plants, world-wide, reactor types
Updated 18 January 2013.
<http://www.euronuclear.org/info/encyclopedia/n/npp-reactor-types.htm>
Retrieved on 2014-06-09.
- [17] Kim, T.K. & Kim, C.H. *Mixed Integer-Programming for Pressurised Water Reactor Fuel-Loading-Pattern Optimization*
Nuclear Science and Engineering: 127, 346-357 (1997)
- [18] US Patent 20040220787 A1
Russell, W.
2003
Method and arrangement for developing core loading patterns in nuclear reactors
- [19] US Patent 7224761 B2
Popa, F.
2007
Method and algorithm for searching and optimizing nuclear reactor core loading patterns
- [20] Personal Source: Fredrik Waldemarsson
Engineer, Thermal Hydraulics
Interviewed 2014-07-11.
- [21] Leinweber, G. et. al. *Neutron Capture and Total Cross-Section Measurements and Resonance Parameters of Gadolinium*
Nuclear Science and Engineering: 154, 261–279 (2006)
http://devinbarry.com/pubs/Journal%20Articles/NSE_154_Gd.pdf
- [22] Kenneth S. Krane
Introductory Nuclear Physics
Wiley, 1987
- [23] John R. Lamarsh & Anthony J. Baratta
Introduction to Nuclear Engineering
Prentice Hall, 2001
- [24] Personal Source: Simon Walve, M.Sc. Senior Engineer BWR Core Design
2012
- [25] Personal Source: Mikael Call, Ph.D. Senior software developer.
2014
- [26] Sveriges Mekanstandardisering
Kärnenergiordlista, Glossary of nuclear energy
Tekniska nomenklaturcentralen, 1990
- [27] Thomas Piketty
Capital in the Twenty-First Century
The Belknap Press, 2014

Appendix A

Additional Data

Perturbation case without nearby BA	No BA in rod (2,2)	S-matrix, no BA
Δ_{k_∞} , enrichment change, rod (2,2)	1.37533 \rightarrow 1.37682	0.00149
Δ_P , enrichment change, rod (2,2)	0.750 \rightarrow 1.254	0.504
Δ_{k_∞} , enrichment change, rod (2,3)	1.37372 \rightarrow 1.37647	0.00275
Δ_P , enrichment change, rod (2,3)	0.709 \rightarrow 1.165	0.456
Perturbation case with nearby BA	6 % BA in rod (2,2)	S-matrix, 6 % BA
Δ_{k_∞} , enrichment change, rod (2,2)	1.33940 \rightarrow 1.34032	0.00092
Δ_P , enrichment change, rod (2,2)	0.268 \rightarrow 0.403	0.135
Δ_{k_∞} , enrichment change, rod (2,3)	1.33714 \rightarrow 1.34102	0.00388
Δ_P , enrichment change, rod (2,3)	0.613 \rightarrow 1.005	0.392

Table A.1: Effects on S-matrix elements from introducing BA. The differences in the third column values show the impact of BA addition on k_∞ and P in neighboring rods. This impact on S-matrices justify the two-step approach taken, where BA is first added and new S-matrices calculated with Gd concentration fixed.

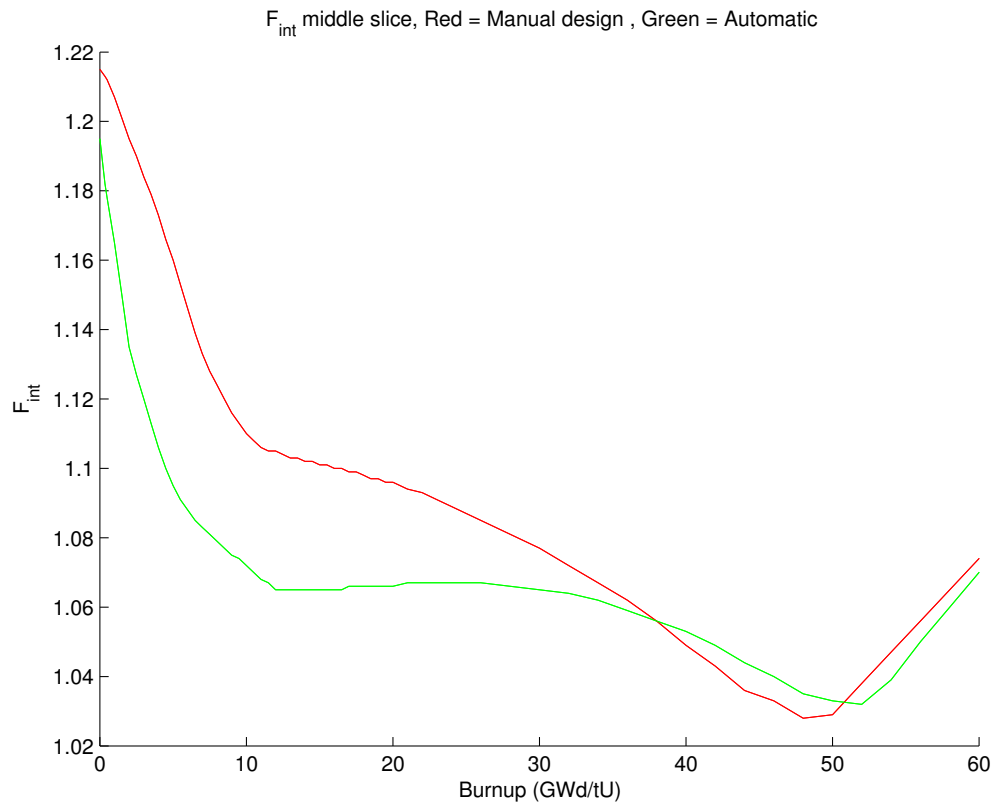


Figure A.1: Short-cycle reactor middle slice maximum relative pin power, red = manual design, green = automatic design. In the middle slice the automatic optimization performs better than the manual design throughout most of the cycle, better than bottom slice comparison shown in Figure 5.20.

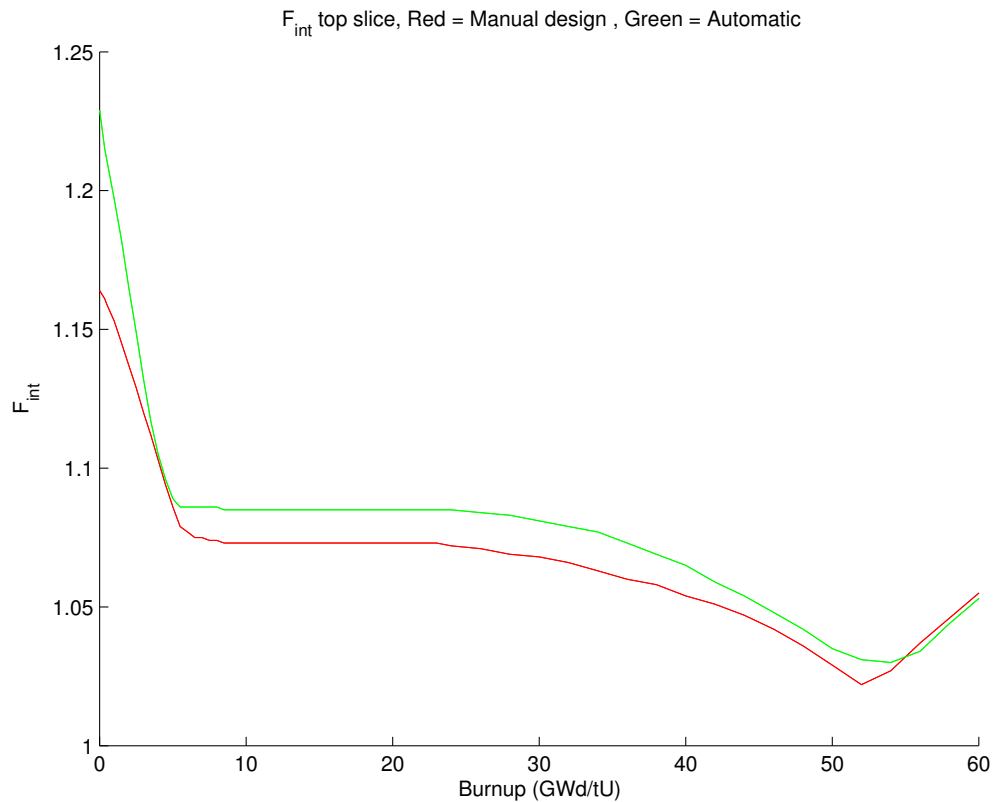


Figure A.2: Short-cycle reactor top slice maximum relative pin power, red = manual design, green = automatic design. In the middle slice the automatic optimization performs worse than the manual design throughout most of the cycle, especially early in the cycle. The manual design however had excess margin so together with Figures 5.20 and A.1 this shows that gains in automatic optimization performance can come from better using all available margins to improve performance in the "worst" rod segment.

Bottom Slice, 0 GWd/tU

-0.014	0.002	0.006	0.000	0.001	0.001	0.000	0.005	0.002	-0.014
0.002	-0.003	0.000	-0.002	0.003	0.003	-0.003	0.000	0.000	0.003
0.006	0.000	0.002	0.002	0.000	0.000	0.003	0.001	0.000	0.006
0.000	-0.002	0.002	-0.001	-0.003	-0.004	-0.002	0.003	-0.001	0.001
0.001	0.003	0.000	-0.003	0.000	0.000	0.000	0.000	0.004	0.001
0.001	0.002	0.000	-0.004	0.000	0.000	0.001	0.000	0.003	0.002
0.000	-0.003	0.002	-0.002	0.000	0.001	-0.004	0.002	-0.001	0.000
0.005	0.000	-0.001	0.003	0.000	0.000	0.002	0.002	0.001	0.005
0.002	-0.002	0.000	-0.003	0.003	0.003	-0.002	0.000	-0.003	0.001
-0.014	0.002	0.006	0.000	0.000	0.001	0.000	0.005	0.001	-0.015

Figure A.3: Pin power relative errors, Figure 5.7 corrected for power in rods, after a step-2 optimization. All errors are below 2 %.

Bottom Slice, 0 GWd/tU										
-0.022	0.003	0.008	0.001	0.001	0.002	0.001	0.008	0.003	-0.022	
0.003	-0.001	0.000	-0.001	0.003	0.003	-0.001	0.000	0.000	0.004	
0.008	0.000	0.002	0.002	0.000	0.000	0.003	0.001	0.001	0.009	
0.000	-0.001	0.002	0.000	-0.003	-0.004	-0.001	0.003	0.000	0.001	
0.001	0.003	-0.001	-0.003	0.000	0.000	0.000	0.000	0.004	0.002	
0.001	0.002	0.000	-0.004	0.000	0.000	0.001	0.000	0.003	0.002	
0.000	-0.001	0.002	-0.001	0.000	0.001	-0.002	0.002	0.000	0.000	
0.007	0.000	0.000	0.003	0.000	0.000	0.002	-0.001	0.001	0.007	
0.002	-0.001	0.000	-0.001	0.003	0.003	-0.001	0.000	-0.001	0.002	
-0.023	0.003	0.008	0.000	0.000	0.001	0.000	0.007	0.002	-0.023	
Middle Slice, 0 GWd/tU										
0.000	-0.002	0.004	-0.001	0.000	0.001	0.000	0.006	0.001	0.000	
-0.002	0.000	-0.002	0.001	0.002	0.002	0.003	-0.001	0.002	0.000	
0.004	-0.002	0.000	0.002	0.001	0.000	0.001	0.002	-0.001	0.004	
-0.002	0.001	0.002	0.000	-0.004	-0.004	0.000	0.001	0.000	-0.001	
0.000	0.002	0.001	-0.004	0.000	0.000	-0.001	0.000	0.002	0.001	
0.000	0.002	-0.001	-0.004	0.000	0.000	-0.001	-0.001	0.002	0.001	
-0.001	0.002	0.000	0.000	-0.001	-0.001	0.000	0.001	-0.001	-0.002	
0.005	-0.002	0.002	1.000	-0.001	-0.001	0.001	0.002	-0.001	0.004	
0.000	0.001	-0.001	0.000	0.001	0.002	-0.001	-0.001	-0.001	0.000	
0.000	-0.002	0.003	-0.002	0.000	0.000	-0.002	0.003	0.000	0.000	
Bottom Slice, 18 GWd/tU										
0.000	0.001	0.000	0.000	0.000	0.000	0.000	0.001	-0.001	0.001	
0.001	-0.002	0.000	-0.001	0.000	0.001	0.003	0.000	-0.001	-0.001	
0.000	0.000	0.001	-0.001	-0.002	0.000	0.001	-0.002	0.000	0.000	
0.000	-0.001	-0.001	-0.001	-0.001	-0.001	0.000	0.001	-0.001	-0.001	
0.000	0.000	-0.002	-0.001	0.000	0.000	0.000	0.000	0.000	0.000	
0.000	0.001	0.000	-0.001	0.000	0.000	0.000	-0.001	0.001	0.000	
0.000	0.003	0.001	0.000	0.000	0.000	0.000	0.002	-0.001	-0.001	
0.001	0.000	-0.002	0.000	0.000	-0.001	0.002	0.009	0.002	0.001	
-0.001	-0.001	0.000	-0.001	0.000	0.001	-0.001	0.002	-0.001	0.001	
0.001	-0.001	0.000	-0.001	-0.001	0.000	-0.001	0.001	0.001	-0.001	

Figure A.4: Errors in pin powers from CPLEX approximation in step 2, by far the largest errors are found in the corner rods early in the cycle. For fuel bundle optimization accuracy this is good, since the most critical part of the cycle is when peak k_{∞} is reached, for this bundle at a burnup of 18 GWd/tU.