Chapter 5

Numerical results and discussion

5.1 Introduction

Chapters 3 and 4 laid the theoretical foundation for an improved transverse leakage approximation. In Chapter 3 a simplification to full higher-order methods was derived, specifically tailored towards the treatment of the transverse leakage terms as they appear in transversely-integrated nodal methods. The preferred approach is termed CQLA or Consistent Quadratic Leakage Approximation. Further performance enhancement was developed in Chapter 4, with three iterative schemes combining into a final solution strategy termed RLCS, or Reduced Leakage Correction Scheme. Although the combination of CQLA and RLCS is the primary outcome of this work, a number of options are available within the HOTR code module and thus the following are numerically investigated in this chapter:

- the accuracy of the proposed full higher-order solution scheme (as discussed in Section 3.4), of arbitrary order, as compared to published results. It is the intention to utilize this option to generate reference solutions for problems in this work, since higher-order methods are published as capable of producing results comparable in accuracy to extrapolated finite-difference results for sufficiently high orders;
- the efficiency and accuracy of combining various "base" higher-order solutions (e.g. CQLA or FHO$_n$), with the developed iteration schemes. In particular, the quality of the CQLA method and the associated improvement due to RLCS, are of primary interest; and
- the accuracy and efficiency of the homogeneous flux reconstruction capability (2D or 3D), as discussed in Section 4.5, with scalable accuracy determined by the order of the method.

These points are, in this chapter, assessed on a set of seven fixed cross-section benchmark problems, namely:

- the 3D version of the MOX C5 G2 problem, defined in Lefebvre et al. (1991) and Lewis (2001) as a difficult mini-core benchmark problem constructed for the purpose of testing transport codes, nodal solvers and flux reconstruction methods;
- the 2D version of the well known 2-group IAEA LWR benchmark problem, originally defined in Lee (1977) for the testing of nodal solvers;
- the 3D version of the IAEA LWR benchmark problem, also taken from Lee (1977);
- the 2D 6-group KOEBERG benchmark problem, as defined by Muller and Weiss (1991);
- the 2D 2-group BIBLIS benchmark problem as described in Smith (1979);
- the 2D 2-group ZION benchmark problem as described in Smith (1979); and
- a realistic 3D, 6-group SAFARI-1 research reactor benchmark problem, defined in Prinsloo et al. (2008).

For all the above problems, the driver code (SANS) utilizes a fourth order source expansion in the one-dimensional equations (with the exception of the FHO_6 calculations which require a sixth order in the driver code) and convergence criteria was set to 10^{-6} in a maximum relative node-averaged flux error for all the problems. Fission source extrapolation is used as the outer acceleration in all the problems and it should be stated that this approach could influence the timing results somewhat, since not all the strategies would converge with the same number of outer iterations. HOTR, the higher-order module coupled to SANS, utilizes a one-dimensional source expansion up to the method order - hence a full second order solution would only employ one-dimensional higher-order sources up to the second order. If not otherwise specified, the iteration structure (order and number of iterations) utilized in the calculations of this chapter is given in Table 5.1.

Table 5.1: Iteration structure for numerical problems.

Iteration level	**Maximum iterations**
Outer fission source	–
Leakage iteration	1
HOTR outers/leakage iterations	2
Up-scatter source	5
Energy group iterations	1
Spatial flux iterations	3

5.2 Reference Result Generation

In the numerical benchmark problems discussed in this chapter, a full sixth order nodal diffusion calculation, performed by the HOTR module, is utilized for the generation of reference solutions. The usage of higher-order methods for the generation of nodal reference solutions (thus on coarse, un-subdivided meshes) has been proposed by Rajić and Ougouag (1989) and Muller and Weiss (1991) as an alternative to fine-mesh finite-difference calculations. This view is strongly supported by Tomašević and Larsen (1993) who demonstrated clearly, on a set of well known 2D benchmark problems, that the superior convergence properties of higher-order methods in both eigenvalue and power distribution allow the generation of reference solutions with these methods. In actual fact, fine-mesh finite-difference methods face significant challenges in producing true reference results for 3D problems, given the limitations in memory, long calculational times and the need for multiple refinement levels used in solution extrapolations. In all these publications, a variety of 2D benchmark problems were utilized in these numerical arguments and there is no reason to expect that the argumentation does not extend to 3D.

Thus, these prior analyses, coupled with an investigation into reference solutions from other sources, such as Smith (1979), are used to support the claim that we can utilize the sixth order HOTR solution option to generate reference flux and power distributions for the set of problems considered in this chapter. Table 5.2 provides a comparison between the sixth order HOTR results and various published reference results for a subset of the numerical benchmarks considered in this work. All these

benchmarks are fully discussed in the rest of this chapter and here we are only interested in the issues surrounding the generation of reference solutions. In Table 5.2, CASTOR (2D higher-order nodal diffusion) results are taken from Tomašević and Larsen (1993), LABAN (2D higher-order response matrix) results are taken from Muller and Weiss (1991), ILLICO (2D higher-order nodal diffusion) are taken from Rajić and Ougouag (1989) and ANM results are taken from Smith (1979) (4 times sub-divided nodal mesh).

Table 5.2: A comparison of reference results from various sources for a selection of the problems considered in this work.

Code and approximation	**Benchmarks (k_{eff} (Maximum % nodal power error))**			
	ZION-1	**BIBLIS**	**IAEA2D**	**IAEA3D**
HOTR (6^{th} order)	1.27489 (0.00)	1.02511 (0.00)	1.029585 (0.00)	1.02907 (0.00)
CASTOR (4^{th} order)	1.27489 (0.00)	1.02511 (0.00)	1.029585 (0.01)	
[a]CASTOR (6^{th} order)	1.27489 (ref)	1.02511 (ref)	1.029585 (ref)	1.02907 (ref)
ILLICO (4^{th} order)	1.27489 (0.02)	—	1.029585 (0.02)	—
LABAN (4^{th} order)	1.27489 (0.00)	1.02511 (0.00)	1.029585 (0.00)	—
LABAN (6^{th} order)	1.27489 (0.00)	1.02511 (0.02)	1.029585 (0.00)	—
ANM (4×4 refined)	1.27489 (0.04)	1.02512, (0.04)	1.029585 (0.02)	
Finite-difference	—	—	1.029585	1.02903 (1.99)

[a]Considered as a reference based on the arguments in Tomašević and Larsen (1993).

Table 5.2 indicates that the higher-order implementation in HOTR, when coupled to a driver nodal code, reproduce the correct full sixth order higher-order result for all four listed problems. It can further be seen that the various higher-order nodal solvers all agree to within 0.02% for both fourth and sixth order solutions, independent of the solution method employed. Slightly larger differences (0.04%) are observed when compared to alternative references produced via 16 node per assembly ANM calculations, as used by Smith (1979). No well described 3D higher-order published results are available, but unpublished results provided by the author of the CASTOR code confirms the HOTR results for the IAEA 3D benchmark, while significant differences exist as compared to the originally published reference solution in Lee (1977). A detailed analysis of this issue is provided in Appendix B.3.

5.3 Results and Discussion

This section aims to place the performance and accuracy of the developed approach when resolving the transverse leakage terms, into perspective. The additional computational burden of the new scheme may of course be numerically quantified in a number of ways, depending on the scope, complexity and size of the benchmark problems considered. We may, for example, be interested in the relative cost of a single transverse leakage update, the total cost of transverse leakage updates during the iteration process, or the total calculational time spent on solving the steady state neutronic problem. In this work the emphasis is on the development of a **practical** scheme for the purpose of realistic, full core, 3D, multi-group diffusion problems and as such the latter measure of total calculational time is of the greatest concern. The fraction of calculational time spent on leakage updates is however also reported in order to indicate how sensitive a given problem is to the cost of the transverse leakage solution.

With this in mind, the range of selected problems gradually increase in scope and complexity, starting with the simple mini-core MOX benchmark problem and concluding with a multi-group 3D full core SAFARI-1 reactor model containing 2520 different materials. Each of the problems are discussed and summative conclusions are drawn at the end of this chapter. Detailed results for each problem, in tabular format, are provided in Appendix B, whereas summary accuracy and performance indicators are highlighted and discussed here.

5.3.1 OECD/NEACRP two-group MOX C5 benchmark problem

In this section the results for the 3D version (Lewis, 2001) of the well known MOX C5 benchmark problem are presented. Two-group homogeneous cell cross-sections were specified in Lefebvre et al. (1991). Assembly homogenized cross-sections are obtained from pin cell data using a discrete ordinate calculation (TWODANT (O Dell et al., 1982)) on the spatial mesh 132 $\times$ 132 and S_{16} quadrature set for each assembly with reflecting boundary conditions. Figure 5.2 depicts the radial thermal flux profile through the axial core centreline, produced by HOTR in the fourth order reconstruction mode (FHO_4). From this figure the sharp flux gradients at the UO_2/MOX interfaces are clearly visible.

The 2D version of the benchmark problem consists of a 3 $\times$ 3 core lattice with a 2 $\times$ 2 array of Uranium and Mixed-oxide fuel elements surrounded by a water reflector

as depicted in Figure 5.1. In 3D, the geometry is extruded to include top and bottom water reflectors.

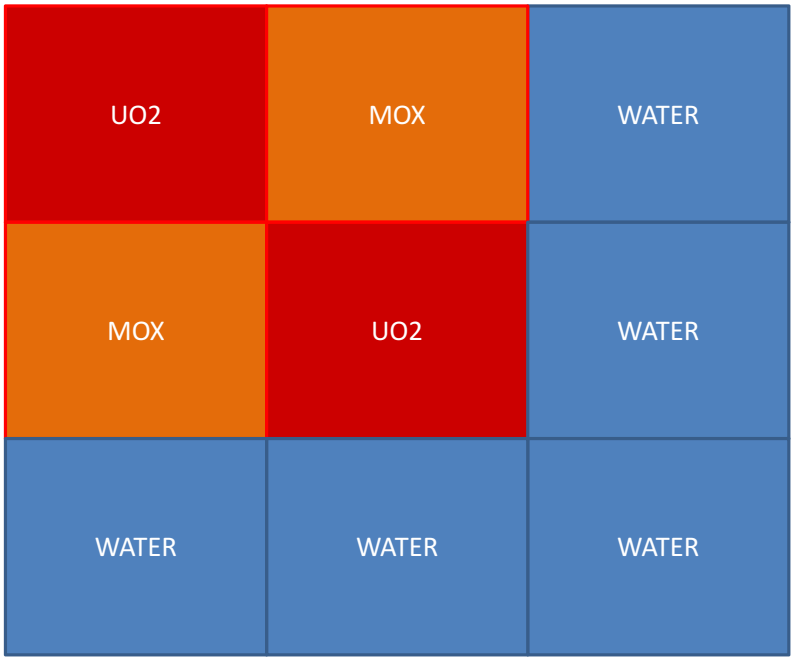

Figure 5.1: 2D geometric layout of the MOX C5 benchmark problem.

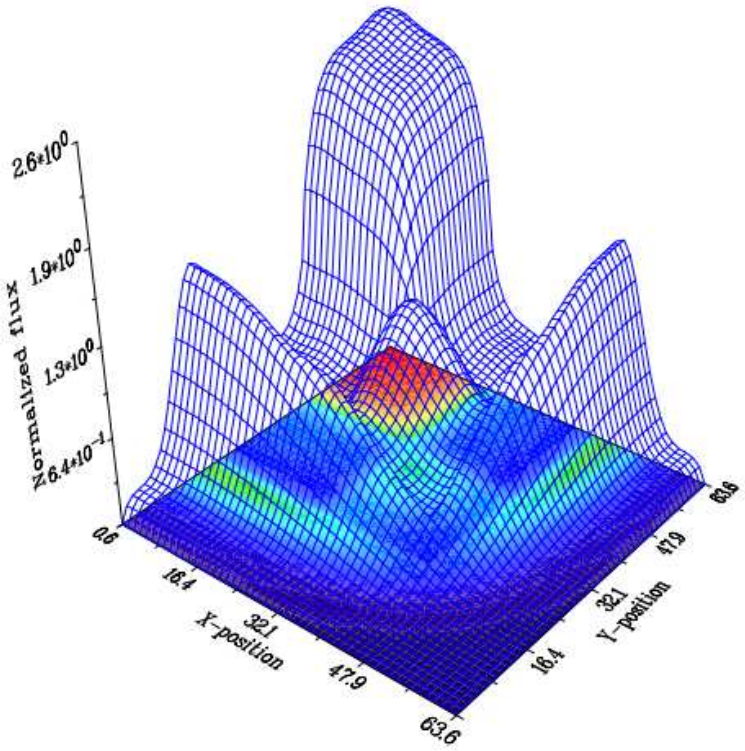

Figure 5.2: Reconstructed 2D thermal flux profile of the MOX C5 benchmark problem.

The MOX C5 benchmark problem is thus a mini-core configuration which lends itself to the parametric analysis of various solution approaches, relating to both the nodal and the flux reconstruction solvers. In this section we utilize this benchmark to quantify the efficiency of the various iteration schemes proposed in Chapter 4. Although the benchmark is a mini-core problem, it provides us with a useful platform to determine how the various iteration schemes scale with the method order. Results obtained for this problem are presented in Table 5.3 and all iteration schemes assumed their standard parameter values as discussed in Section 4.4. Detailed results of this problem are tabulated in Section B.1.

Table 5.3: Results for the 3D, two-group MOX C5 benchmark problem.

Iteration Scheme	k_{eff} **# (pcm)**	**Cost factor**	**Ass. Pow. Err Ave (Max) %**	**Nod. Pow. Err Ave (Max) %**
FHO_6	0.94042	86.6	—	—
$SQLA_{hotr}$	0.94060 (19)	1.0	0.22 (0.62)	0.22 (0.67)
FHO_6-MR	0.94042 (1)	68.5	0.02 (0.04)	0.04 (0.08)
FHO_6-PLC	0.94042 (1)	43.2	0.02 (0.04)	0.02 (0.06)
FHO_4	0.94040 (2)	20.4	0.01 (0.03)	0.01 (0.04)
FHO_4-MR	0.94040 (2)	13.0	0.005 (0.01)	0.02 (0.08)
FHO_4-PLC	0.94040 (2)	10.1	0.01 (0.02)	0.02 (0.06)
FHO_2	0.94036 (7)	3.4	0.04 (0.06)	0.06 (0.10)
FHO_2-MR	0.94035 (7)	2.7	0.04 (0.06)	0.06 (0.19)
FHO_2-PLC	0.94035 (7)	2.4	0.04 (0.06)	0.06 (0.10)
FHO_2 -QLAC	0.94035 (7)	4.1	0.04 (0.05)	0.11 (0.29)
CQLA	0.94026 (16)	2.6	0.10 (0.17)	0.15 (0.39)
CQLA-MR	0.94024 (18)	2.3	0.09 (0.16)	0.12 (0.24)
CQLA-PLC	0.94026 (16)	1.9	0.10 (0.17)	0.15 (0.39)
CQLA-MR-PLC	0.94026 (16)	1.8	0.10 (0.17)	0.14 (0.30)
CQLA-QLAC	0.94025 (17)	1.47	0.10 (0.17)	0.19 (0.39)
CQLA-RLCS	0.94025 (17)	1.41	0.08 (0.15)	0.15 (0.31)

Table 5.3 presents the accuracy and timing comparisons of various iteration strategies for the 3D MOX problem, which studies the efficiency and accuracy of the iteration schemes, parametrically against the method order. For the purpose of comparison, the SQLA solution, with a 19 pcm k_{eff} error, is scaled to a computation cost factor 1.0 and thus all timing comparisons are relative to it. The Model Reduction (MR) and Partial Leakage Convergence (PLC) iteration strategies are applied to the four different base solution strategies, namely full sixth order (FHO_6), full fourth order (FHO_4), full second order (FHO_2) and the CQLA simplification. The Quadratic Leakage Approximation Correction (QLAC) is applied only to the second schemes, since utilizing higher-order solutions to provide correction factors to the second order SQLA provides very little benefits. Figure 5.3 further illustrates the results in Table 5.3 by analyzing the performance of the various base schemes (SQLA, CQLA and

FHO_n).

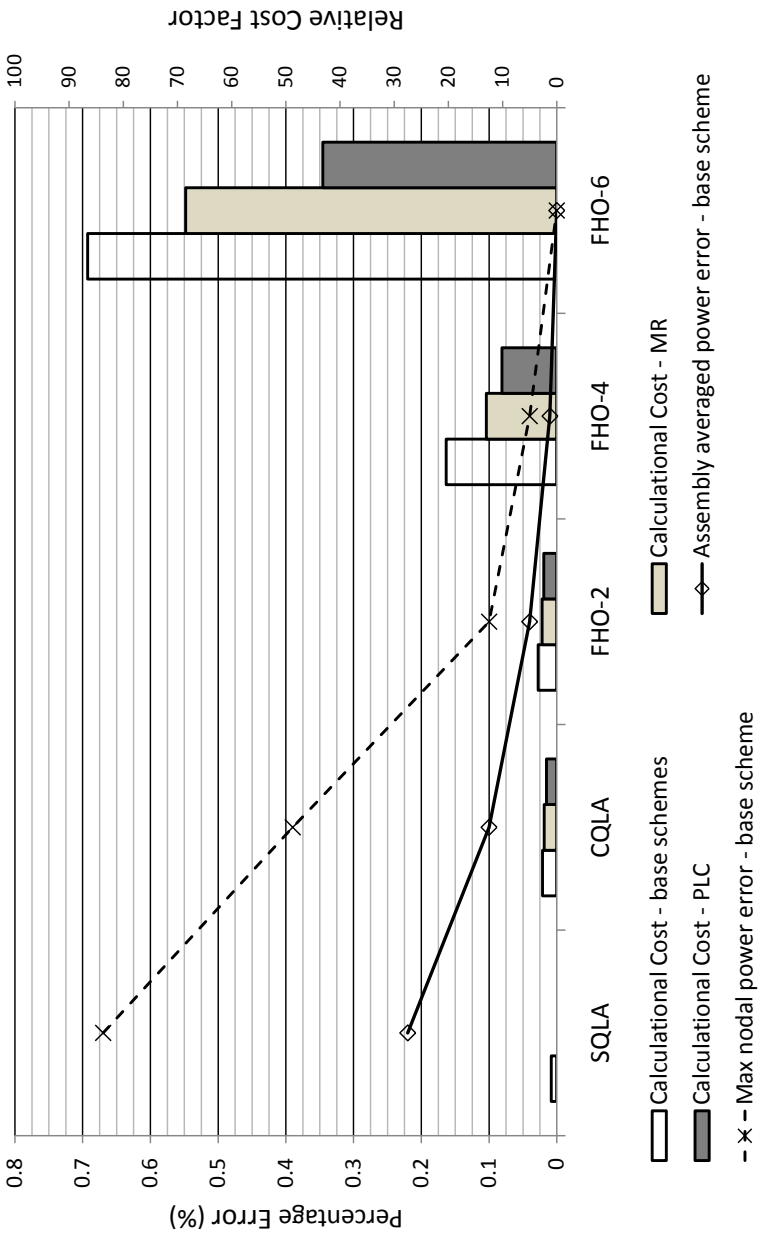

Figure 5.3: Analysis of the performance and accuracy of various solution schemes as applied to the MOX C5 benchmark problem.

In Figure 5.3 a number of schemes are graphically compared in terms of the

average assembly power error and the maximum nodal power error. The bar chart, on the secondary $y-$axis, measures the relative computational cost factor of each of the schemes. In order to clearly distinguish the relative efficiencies of the lower order schemes, Figure 5.4 highlights the data of Figure 5.3 for the SQLA, CQLA and FHO-2 cases and includes the results of the QLAC scheme.

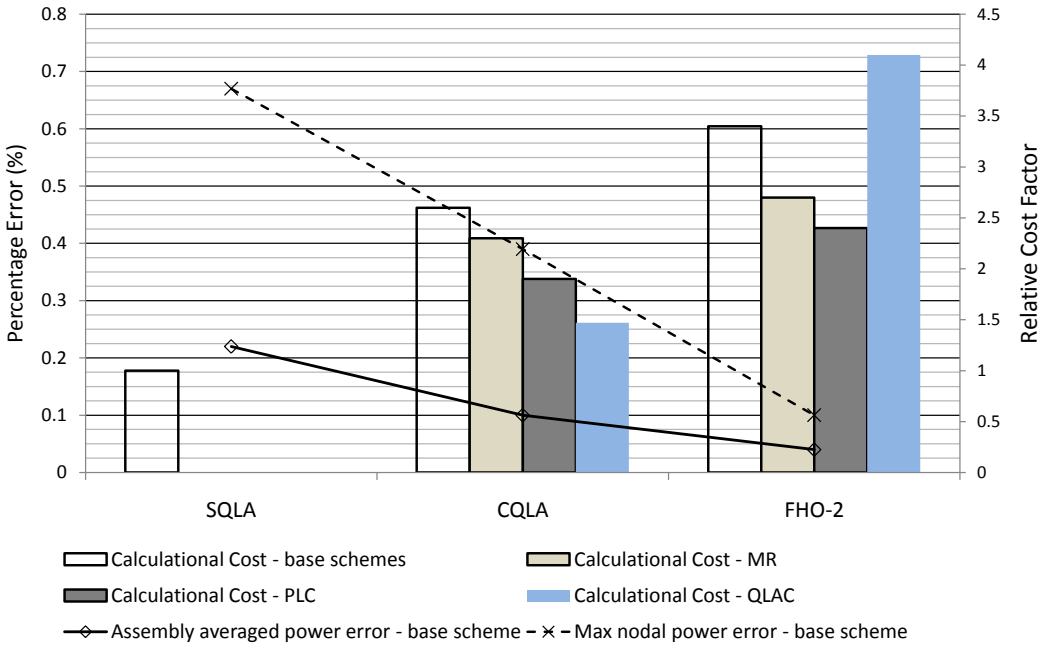

Figure 5.4: Specific focus on the efficiency of lower order schemes as applied to the MOX C5 benchmark problem.

A number of observations are made from Table 5.3, Figure 5.3 and Figure 5.4:

- All approaches are generally accurate for this problem. The accuracy clearly improves with the order and importantly the CQLA solution is more accurate than SQLA by about a factor of two. The scaling of accuracy versus method order is typical for higher-order methods and follows the trends as published in Ougouag and Rajić (1988). For the case of SQLA, about 5% of the total calculational time is spent on resolving the transverse leakage source, while for the CQLA case this fraction increases to just over 30% for this problem;

- Partial Leakage Convergence (PLC), Model Reduction (MR) and QLA Correction (QLAC) generally maintain the accuracy of the base scheme they are "attached" to, be it FHO_n or CQLA. This is an important feature, since the reduction of calculational time associated with these schemes come with almost no additional accuracy penalty;

- The PLC approach improves upon the efficiency of its base scheme by about a factor of two. As input, this scheme requires the relative convergence criteria at which the leakage shape should be frozen. For all these calculations performed, the set point was placed at 10^{-4}. A lower value for this set point would improve the accuracy at the cost of the performance;

- The MR approach shows a larger cost benefit for higher-order solutions, as expected, than for lower order solutions. In the case of FHO_4-MR, 63% of all the moments in the system were rejected. For the CQLA case, only 13% of the moments were rejected, leading to an effective (although small) increase in calculational time due to the overhead of the ANOVA implementation. For these calculations, the set point for the inclusion of moments, measured as a percentage contribution to the total function variance, was set to around $10^{-2}(1\%)$, which places a limit on the achievable accuracy. In general, the benefit for the higher-order solutions of this problem leads to about a 30% reduction in calculational time. The CQLA-MR-PLC case refers to the scheme which applies to both PLC and MR of the CQLA base scheme and shows that some additional benefits may be derived from the combination of the iteration schemes;

- The QLAC and RLCS schemes maintain the accuracy of the CQLA solution, with a reduction of the calculational cost of 1.47 and 1.41, respectively. Although this signifies a good reduction in cost, it is expected that, for this relatively small benchmark problem, the overhead involved in applying these schemes, may influence the timing results. The application of QLAC to the FHO_2 base solution does not yield an efficient solution. The primary reason for this is the fact that this is the slow rate of convergence of the SQLA correction factors, given the additional cross-terms present in the FHO_2 solution, as compared to CQLA which only solves for the uni-variate leakage coefficients; and

- For this problem, the most efficient scheme, namely CQLA-RLCS, provides a two-time improvement in accuracy at a calculational cost of 1.41 as compared to SQLA. In the concluding section of this chapter we shall compare the performance of various problems via a constructed Figure Of Merit (FOM), defined as

$$\text{FOM} = \frac{\text{CPUtime}_{\text{rel}}}{\text{Accuracy}_{\text{rel}}}$$

  which in this case, for the CQLA-RLCS scheme, would yield $1.41/(0.67/0.31) = 0.65$. This definition indicates better performance for lower values of the FOM and is relative to a value of 1.0 for SQLA.

5.3.2 IAEA LWR two-group benchmark problem

This benchmark problem is widely known and used to evaluate and compare diffusion solvers. The benchmark is a simplified representation of two-zone, 177 fuel element LWR, with an active height of 340 cm. Nine inserted control rods are represented as homogenized fuel/absorber cross-sections, with four partially inserted rods present in the 3D version. The 2D core layout is depicted in Figure 5.5.

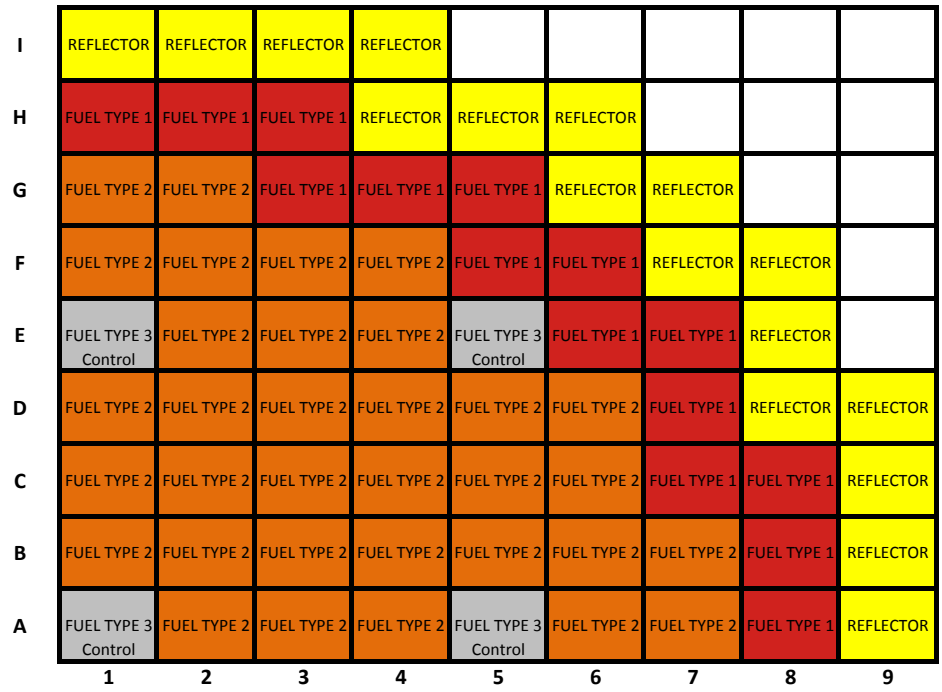

Figure 5.5: IAEA LWR 2D core layout.

The benchmark has a symmetry line through the centre of row A and column 1. The problem also exhibits octant symmetry, but all the results presented in this section solve the quarter symmetry problem, with the symmetry line through row A and column 1 modelled as a reflective boundary condition. Reference results of the problem, in terms of both eigenvalue and power distribution, are provided from the original benchmark publication (Lee, 1977) which utilized an extrapolated fine-mesh finite-difference technique. Some uncertainty has been expressed by various users of this benchmark, as mentioned in Lawrence (1986), that the extrapolated 3D result was not fully converged and an analysis of the quality of the published benchmark is performed in Appendix B.3. This investigation shows that indeed, specifically in the outer axial fuel zones, some nodal power errors in the order of 2% are found and as such the full sixth order higher-order solution utilized here can be regarded as more accurate than the published reference. The issue is discussed further in Section B.3.

5.3.2.1 2D version

Firstly we investigate the 2D version of the problem, given that it is most often cited and used to evaluate nodal solvers. Detailed results of this problem, in tabular format, is provided in Section B.2. Prior to analyzing the efficiency of the various numerical schemes, it is insightful to obtain an understanding of the nature of the problem. Figure 5.6 depicts the thermal flux profile in the system, produced by HOTR in fourth order reconstruction mode (FHO_4). The sharp flux gradients in the reflector and the severe flux depression in the control positions, are the primary reason why this problem is considered to be difficult.

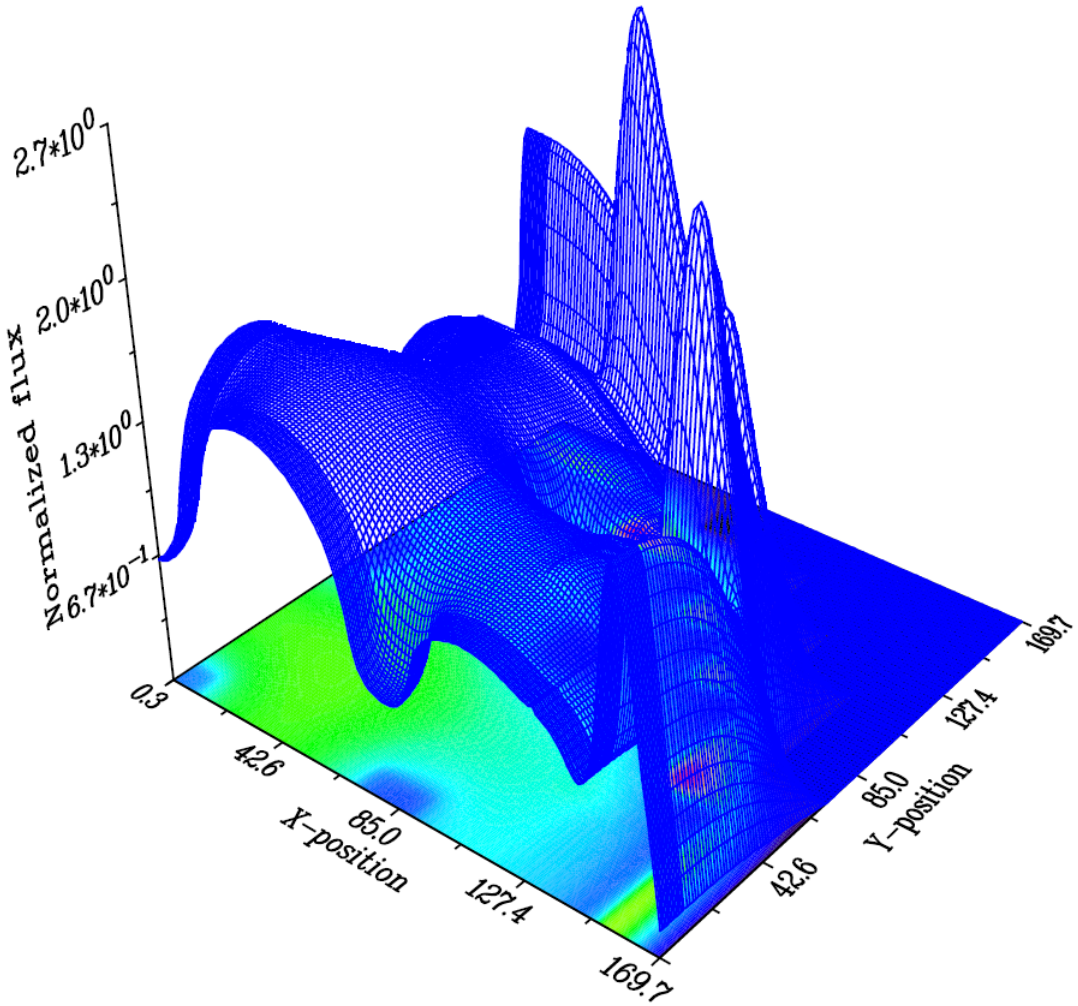

Figure 5.6: IAEA LWR 2D reconstructed thermal flux from HOTR.

Table 5.4 presents the results of the numerical analysis and comparison of various solution schemes as applied to the 2D version of the IAEA LWR benchmark problem.

Table 5.4: Results for the 2D, two-group IAEA LWR benchmark problem.

Solution Scheme	k_{eff} **# (pcm error, outers)**	**Cost factor**	**Power Error Ave (Max) %**	**Flux error Ave (Max) %**
FHO_6	1.029585 (–, 96)	125 (106.00)	—	—
SQLA_{sans}	1.029615 (3.0,82)	0.85 (0.85)	0.25 (0.91)	0.78 (9.89)
SQLA_{hotr}	1.029615 (3.0,82)	1.00 (1.00)	0.25 (0.91)	0.78 (9.89)
FHO_4	1.029584 (0.1, 172)	61.5 (29.32)	0.01 (0.04)	0.01 (0.10)
FHO_2	1.029574 (1.1, 118)	9.10 (6.32)	0.07 (0.35)	0.11 (0.45)
CQLA	1.029601 (1.3, 91)	3.14 (2.82)	0.04 (0.16)	0.11 (1.14)
CQLA_{rlcs}	1.029590 (0.8, 112)	2.61 (1.91)	0.04 (0.14)	0.15 (1.26)
ILLICO_{FHO2}	1.029575 (1.1, 45)	3.32 (6.05)	— (0.32)	—
NEM	1.02965	—	— (1.4)	—

Table 5.4 provides an overview of the accuracy and performance of the various solution schemes in HOTR, as connected to the SANS nodal solver. The first column labels the solution scheme, the second the achieved k_{eff} value followed, in brackets, by the associated error in pcm and the number of outers needed to solve the problem. Column three rates the computational efficiency of the scheme, as compared to the reference performance of 1.00 assigned to the SQLA solution from HOTR (denoted by SQLA_{hotr}). The computational cost column contains an additional bracketed result which gives the average cost per outer (as compared to a cost/outer of 1.0 for SQLA). This additional measure, although somewhat unrelated to the actual running time, provides an alternative approach to comparing calculations which converge in a significantly different number of outers, as in the case of these results given the differing number of fission source extrapolations performed. Columns four and five provide the average (and maximum) power and flux errors, respectively. The flux errors in column six are given for the energy group with the largest error, which in all the cases was the fast group (group one).

The following observations are made based on Table 5.4:

- It should be noted that the results presented here are all calculated with a 3D code and hence the 2D problem is set up with a 3D nodal model utilizing a single axial mesh (with reflective boundary conditions). Hence, it is expected

that the calculational time scales with a number of unknowns is roughly shown as $(M+1)^2$ as would be expected for 3D;

- A relatively small timing difference exists between the SQLA solution in SANS (SQLA_{sans}) and in HOTR (SQLA_{hotr}) - about 15% slower via HOTR, which can be attributed to the additional overhead in coupling to HOTR. The primary concern in this regard is, currently the geometry and cross-section model of the driver code is not directly utilized by HOTR for the sake of independence and generality. In an industrial, full code integration between a driver nodal code and HOTR, it would be expected that HOTR makes use of the driver code topology modules and thus largely eradicate this small overhead burden. As a result of this expectation, all the cost comparisons in this chapter are performed relative to SQLA_{hotr}; and

- The average and maximum nodal power errors, due to SQLA, are 0.25% and 0.91% respectively, with a maximum flux error in the reflector of about 10%. Although, in this problem, these errors are not that significant, this simplified 2D version of the benchmark remains a good platform on which to evaluate potential accuracy improvements. Table 5.5 shows a reference relative power density distribution for this problem, as well as the distribution of a relative power density error associated with the SQLA and CQLA solutions in databar format;

Table 5.5: Reference relative power density results for the 2D IAEA LWR benchmark, with SQLA and CQLA percentage errors indicated in databar format in each cell.

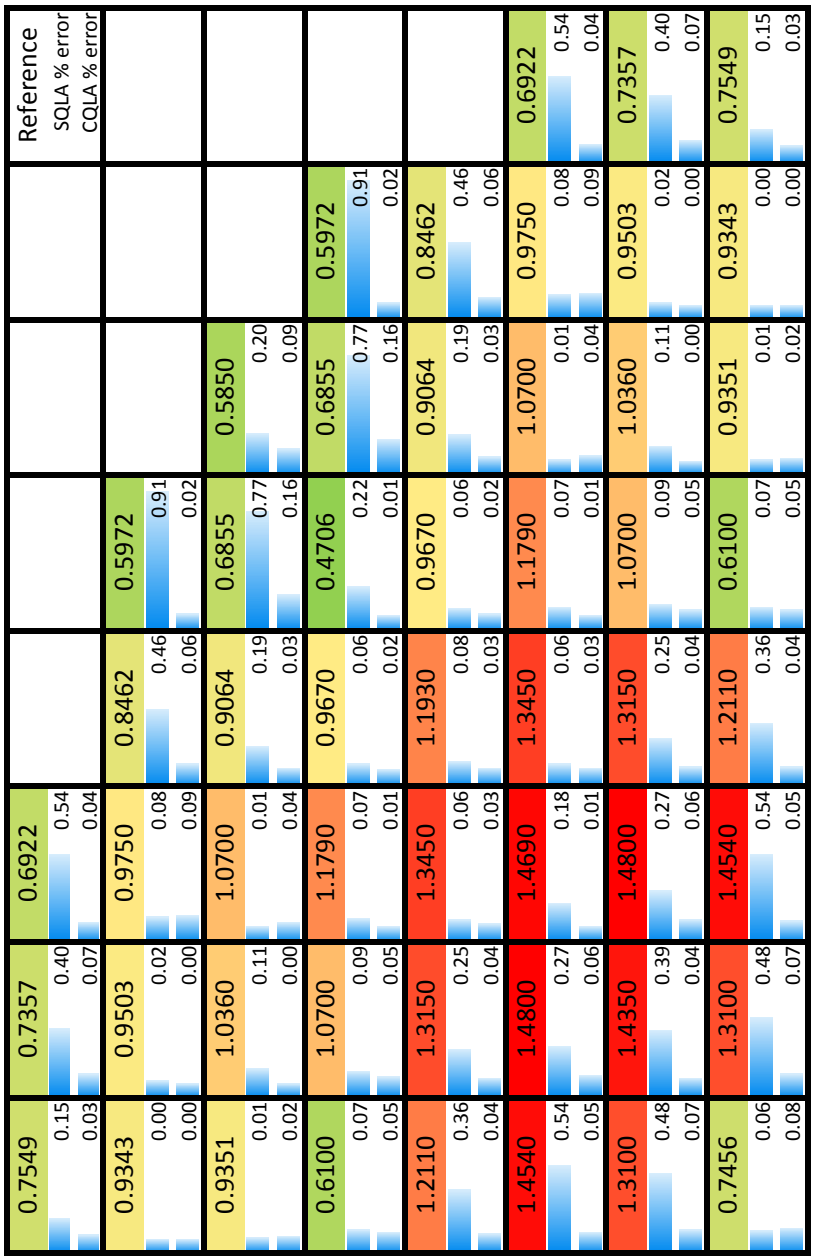

- We notice from Table 5.5 that the primary source of the errors in SQLA occurs at the fuel/reflector interfaces, at the periphery of the model. The large errors in fast flux in the reflector regions, as can be seen in Appendix B Table B.3, are the major contributing factors to these errors;

- As expected, the various higher-order schemes significantly improve upon the accuracy of SQLA. Of particular interest to this work is the accuracy of the CQLA solution, which improves upon SQLA in all measures by a substantial margin. The maximum power and flux errors have decreased by factors of 6 and 9, respectively. For this problem, FHO_2 and CQLA performs similarly in accuracy, which is one of the original aims in developing CQLA. For the case of SQLA, about 7% of the total calculational time is spent on resolving the transverse leakage source, while for the CQLA case this fraction increases to just over 40%, resulting in an associated increase in the overall calculational cost of CQLA;

- The computational cost of the FHO_2 solution is about 9 (6 if the alternate measure is used) times more expensive than SQLA. CQLA reduces this cost to a factor of 3, which is generally in-line with the expectations in Chapter 3 regarding the reduction in the number of unknowns;

- The RLCS scheme, as applied to CQLA (termed $CQLA_{rlcs}$) further reduces the cost factor only slightly to 2.6 (1.9 if an alternate measure is used), which is less than anticipated, but may be understood, given the additional outer iterations performed. $CQLA_{rlcs}$ retains the accuracy of CQLA, with very little deterioration due to the selected iteration parameters, such as $p_{qlac} = 95\%$ (see Table 4.5), which implies that only 95% of the system requires converged SQLA corrections. It is foreseen that $CQLA_{rlcs}$ would show better performance for larger 3D problems, since for this 2D problem the incurred overhead cost is not negligible, as compared to the short total running time. The obtained FOM for this problem is 0.4; and

- The final two entries in the table are sourced from published literature. The first is for the ILLICO higher-order code as published in Ougouag and Rajić (1988) and the second is for the NEM code as published in Finnemann et al. (1977). The ILLICO timing is relative to its own "flat" leakage solution and closely matches the FHO_2 result from HOTR. It confirms that the hierarchical approach for the construction of the higher-order moments, as discussed in Section 3.4,

is consistent with the originally published method of sweeping the full system for the determination of the node-averaged higher-order flux moments. The computational cost per outer associated with these approaches is nearly the same (around 6 in both cases), since the HOTR approach places a greater burden on the outer iterations, as compared to ILLICO which resolves the higher-order moments via a dedicated additional iteration level. However, the actual calculational time as compared to SQLA differs significantly due to the different number of outers, different convergence criteria (10^{-5} for ILLICO) and due to the fact that the ILLICO is a 2D code, whereas the HOTR solution still solves a 3D problem with reflective boundary conditions on the top and bottom surfaces of a single axial node for this problem.

5.3.2.2 3D version

The 2D version of the problem indicates that CQLA provides an accurate solution and a good general performance improvement as compared to full higher-order solutions. This 2D problem, however, might still not provide a good platform for determining the method efficiency of the RLCS scheme as applied to realistic 3D problems, since the short run-times cause timing results to be compromised by method overhead costs. A more appropriate measure and a more realistic core benchmark, is found in the full 3D version of this problem. The radial layout of the 3D version is the same as for the 2D shown in Figure 5.5, with an additional partially inserted control rod in position C3. Results for the various solution schemes applied to the 3D version of the benchmark are presented in Table 5.6.

Table 5.6: Results for the 3D, two-group IAEA LWR benchmark problem.

Solution scheme	k_{eff} **#, (pcm,outers)**	**Cost factor # (#/outer)**	**Ass. Pow. Err Ave (Max) %**	**Nod. Pow. Err Ave (Max) %**
FHO_6	1.02907 (229)	260.00	—-	—-
$SQLA_{hotr}$	1.02911 (4.2, 210)	1.00 (1.00)	0.26 (0.92)	0.28 (1.23)
FHO_2	1.02906 (0.7, 183)	10.12 (11.61)	0.08 (0.35)	0.08 (0.46)
FHO_2-MR	1.02906 (0.7, 192)	7.04 (7.71)	0.08 (0.35)	0.08 (0.46)
FHO_2-PLC	1.02906 (0.7, 187)	5.36 (6.01)	0.08 (0.35)	0.08 (0.46)
FHO_2-QLAC	1.02906 (0.7, 207)	2.71 (2.74)	0.08 (0.35)	0.15 (0.69)
$FHO_{2,rlcs}$	1.02906 (0.7, 162)	2.26 (2.92)	0.08 (0.35)	0.09 (0.55)
CQLA	1.02909 (1.6, 159)	2.82 (3.72)	0.04 (0.16)	0.05 (0.22)
CQLA-MR	1.02909 (1.6, 161)	2.42 (3.15)	0.04 (0.16)	0.07 (0.33)
CQLA-PLC	1.02909 (1.6, 159)	1.81 (2.39)	0.04 (0.16)	0.05 (0.22)
CQLA-QLAC	1.02909 (1.6, 172)	1.23 (1.50)	0.04 (0.16)	0.06 (0.30)
$CQLA_{rlcs}$	1.02909 (1.8, 160)	1.04 (1.37)	0.03 (0.16)	0.08 (0.34)

Table 5.6 follows the structure of Table 5.4. The k_{eff} in column 2 is followed by the error in pcm in brackets and the number of outer iterations required for the convergence of the given scheme. The cost factor column contains an additional bracketed result which gives the average cost per outer (as compared to a cost/outer of 1.0 for SQLA). Columns four and five give the average (and maximum) assembly and nodal power errors, respectively. For more detail, a full, axial layer-by-layer power density comparison between the published result in Lee (1977), the HOTR FHO_6 reference and CQLA are presented in Tables B.5 - B.21. In support of the results in Table 5.6, Figure 5.7 provides a graphical representation of the performance of the various schemes.

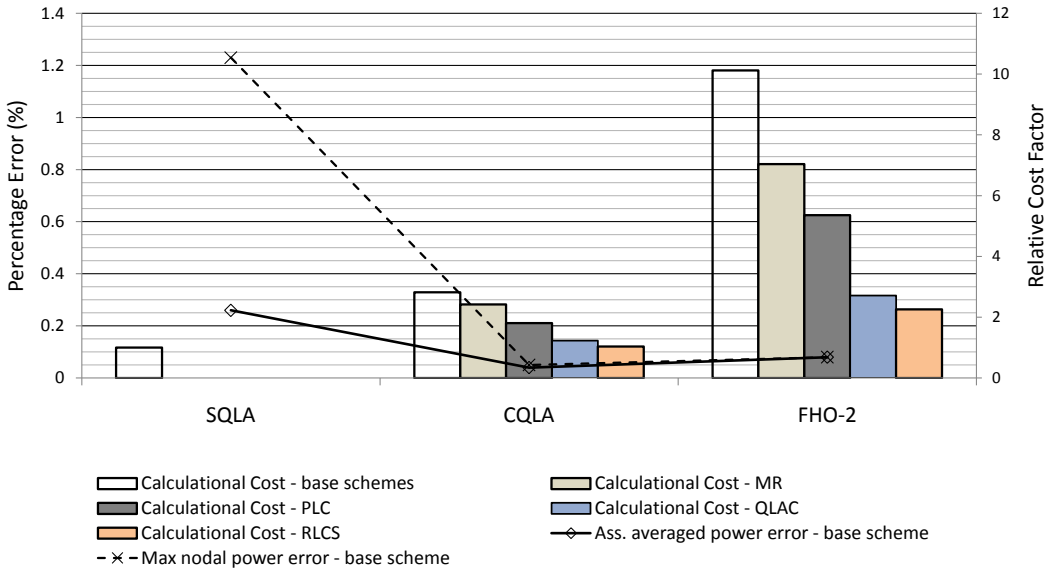

Figure 5.7: Analysis of the performance and accuracy of various solution schemes as applied to the IAEA 3D LWR benchmark problem.

Figure 5.7 indicates the average assembly power error and maximum nodal power error for a selection of the schemes given in Table 5.6. The bar charts, on the secondary y-axis indicate the relative computation cost of each of the schemes. From Table 5.6 we primarily conclude that:

- The maximum error for SQLA is larger for the 3D version of the problem, with a maximum nodal power error of 1.23% after 210 outer iterations. The maximum assembly power error is 0.92%;
- A relatively large variation in the total number of outer iterations exist between the various base schemes (SQLA, CQLA, FHO_2), which indicate different dominance ratios during convergence and a potentially smoother iteration behaviour. The alternative error measure of effective cost/outer is thus introduced to avoid these differences. It is noted though, that the higher-order based CQLA scheme consistently converges in fewer outers, which may indicate a more stable iteration scheme;

- Once again, the higher-order methods achieve a significant reduction in error, both in terms of the maximum and the average error measures. The assembly averaged power error measures closely follow those of the 2D benchmark, but the 3D nodal power distribution shows that CQLA improves upon both the average and the maximum errors in SQLA by a factor of 5. Unexpectedly, CQLA outperforms (in terms of accuracy) FHO_2, but the associated difference in accuracy may be regarded as negligible. The FHO_6 result is exceptionally slow in converging, indicating a worsening of the dominance ratio for the sixth order solution, primarily due to the slow convergence of the higher-order spatial moments;

- For the case of SQLA, about 15% of the total calculational time is spent on resolving the transverse leakage source, while for the CQLA case this fraction increases to just over 80%. For the $CQLA_{rlcs}$ case, this fraction is reduced to around 35%. These differences are expected, given that SQLA simply performs three local node quadratic fits in order to express the transverse leakage terms, while CQLA (or derivatives thereof) requires updates to higher-order fission and scattering sources during each outer to solve 2 two-node problems on each nodal surface interface;

- $CQLA_{rlcs}$ performs significantly better in the 3D case than in the 2D case. In total the $CQLA_{rlcs}$ iteration scheme achieves a cost factor of 1.04 when compared to SQLA. It is however clear that the low number of outers (160 in this case) contributed to this cost factor and thus the effective (cost/outer) is probably a more realistic reflection at a factor of 1.37. It should however be observed that the CQLA scheme consistently converges faster than SQLA, as could be expected considering the fact that SQLA is an ad-hoc extension to the ANM, while CQLA utilizes information which is naturally present in the transversely-integrated equations. In order to understand how this performance is obtained, it is insightful to analyze the iteration behaviour of the $CQLA_{rlcs}$ solution. Table 5.7 describes the workings of the RLCS scheme during the solution and shows that only 14 higher-order iterations were needed during the 160 total outers to achieve sufficient accuracy in the SQLA correction factors. The obtained FOM for this problem is 0.28;

Table 5.7: Iteration analysis for the 3D, two-group IAEA LWR benchmark problem in CQLA$_{rlcs}$ mode.

Iteration range	**Numerical scheme**	**Iteration behaviour**
1-44	SQLA	Standard iterations using SQLA up to preset convergence
33	SQLA	Fission source extrapolation performed
45-59	CQLA$_{rlcs}$	CQLA activated with SQLA correction factor tabulation
52	CQLA$_{rlcs}$	Model reduction due to parameter $f_{mr} = 0.6$: 12% of moments in the system rejected (no longer updated)
59	CQLA$_{rlcs}$	Higher-order moments frozen (parameter $f_{plc} = 0.8$)
61-160	SQLA	95% of moments converged and SQLA continues with corrections factors
95, 124, 147	SQLA	3 further fission source extrapolations take place

- The CQLA$_{rlcs}$ scheme shows some degradation in accuracy when compared to the CQLA solution, with the maximum error increased from 0.22% to 0.34%. Nevertheless, results for the CQLA$_{rlcs}$ case shows promise and provide an improvement in the assembly average and the maximum power error of almost one order of magnitude and an improvement in the maximum nodal power error of approximately a factor of 4, for a less than 40% increase in the calculational cost per outer (4% increase in total time).

5.3.3 KOEBERG benchmark problem

The 6-group KOEBERG benchmark problem (Muller and Weiss, 1991) is a 2D, 6-group diffusion benchmark and is derived from realistic cross-sections for the KOEBERG LWR operating in Cape Town, South Africa. The multigroup nature of the problem (with significant up-scattering), provides additional complexity. As illustration of the nature of the problem, the reconstructed flux from HOTR (in fourth order reconstruction mode) is presented in Figure 5.8, from which the checkerboard arrangement of fuel and control elements in this design is evident. The problem exhibits octant symmetry, but for the sake of testing the methods on a realistic full-core problem, all the calculations in this section are performed for the full core problem (no symmetry).

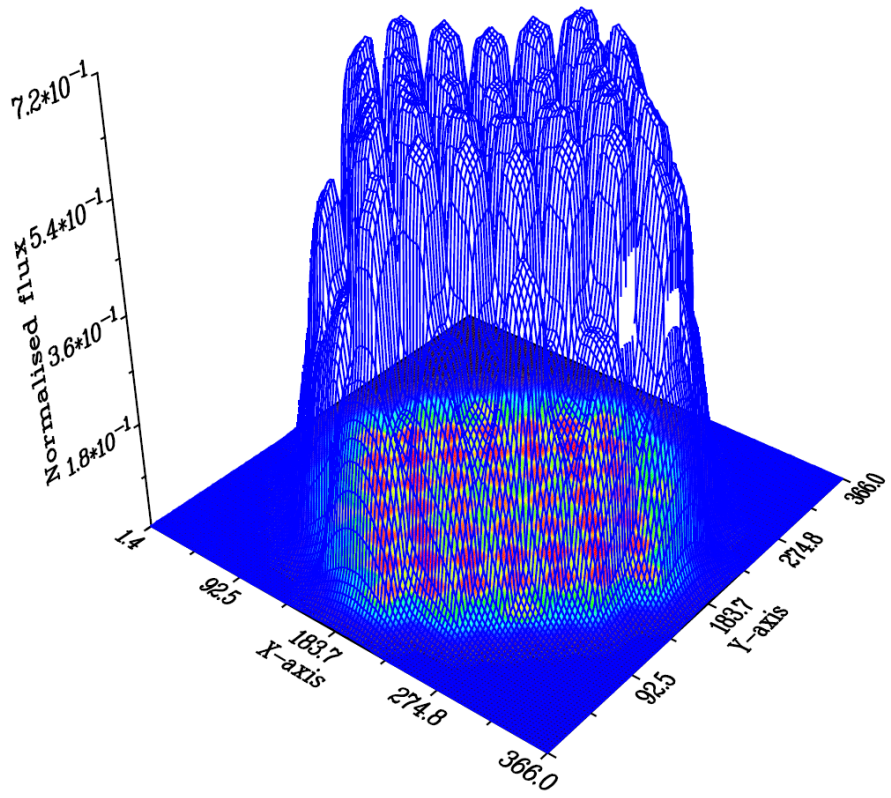

Figure 5.8: Reconstructed group 6 flux profile in the KOEBERG benchmark problem.

Summary results for the KOEBERG problem are presented in Table 5.8.

Table 5.8: Results for the 2D, six-group KOEBERG benchmark problem.

Iteration scheme	k_{eff}	k_{eff} **error (pcm)**	**Cost factor**	**Ass. Ave (Max) % Power Error**	**Ass. Ave (Max) % Flux Error**
FHO_6	1.00787	—	152	—	—
SQLA	1.00828	40 pcm	1.00	1.11 (2.24)	2.21 (15.2)
FHO_2	1.00786	1 pcm	1.53	0.02 (0.07)	0.04 (0.21)
CQLA	1.00786	1 pcm	1.20	0.14 (0.42)	0.32 (1.64)
$CQLA_{rlcs}$	1.00782	5 pcm	1.09	0.14 (0.49)	0.32 (1.72)

Table 5.8 tabulates the SQLA, CQLA and FHO_2 results for the KOEBERG problem. For this problem, the iteration structure, as presented in Table 5.1, is adapted in that the maximum up-scatter iterations is set to 100 to allow smooth error-reduction during convergence.

The most noteworthy result in this table is the large SQLA nodal power error, which is over 2% for this problem. CQLA (and CQLA$_{rlcs}$) improves upon this result by a factor of 4.5. The power error distribution for the SQLA calculation is presented in Appendix B in Figure B.1.

From Figure B.1 it is clear that the 2.2% error is not an isolated occurrence, but consistently shows up at the periphery of the model. More detail on the distribution is available in Table 5.9, which shows a consistent spread of SQLA errors and thus explains the large assembly averaged error.

Table 5.9: Reference relative power density results for the 2D KOEBERG benchmark (first quadrant), with SQLA and CQLA percentage errors indicated in databar format in each cell.

	1	2	3	4	5	6	7	8
H	8.293E-01 2.24 0.33	6.397E-01 2.06 0.32						Reference SQLA % error CQLA % error
G	9.566E-01 1.07 0.21	1.041E+00 1.23 0.16	9.647E-01 1.54 0.18	6.478E-01 0.86 0.19				
F	1.215E+00 0.05 0.00	1.060E+00 0.19 0.06	1.039E+00 0.51 0.11	9.798E-01 1.67 0.18	6.655E-01 1.51 0.17			
E	1.132E+00 0.56 0.07	1.245E+00 0.58 0.08	1.058E+00 0.32 0.01	9.999E-01 0.16 0.06	7.849E-01 1.75 0.22	6.655E-01 1.51 0.18		
D	1.168E+00 1.34 0.18	1.107E+00 1.10 0.11	1.226E+00 0.86 0.17	1.037E+00 0.51 0.03	9.999E-01 0.16 0.06	9.798E-01 1.67 0.19	6.478E-01 0.86 0.19	
C	1.046E+00 1.51 0.29	1.135E+00 1.69 0.26	1.094E+00 1.22 0.21	1.226E+00 0.86 0.17	1.058E+00 0.32 0.01	1.039E+00 0.51 0.11	9.647E-01 1.54 0.18	
B	1.090E+00 1.99 0.42	1.028E+00 1.69 0.33	1.135E+00 1.69 0.26	1.107E+00 1.10 0.11	1.245E+00 0.58 0.08	1.060E+00 0.19 0.06	1.041E+00 1.23 0.16	6.397E-01 2.06 0.32
A	1.008E+00 1.91 0.35	1.090E+00 1.99 0.42	1.046E+00 1.51 0.29	1.168E+00 1.34 0.18	1.132E+00 0.56 0.07	1.215E+00 0.05 0.00	9.566E-01 1.07 0.21	8.293E-01 2.24 0.33

The result of particular interest, is the CQLA$_{rlcs}$ performance factor of 1.09, with an average and maximum power density error of 0.14% and 0.49%, respectively, when compared to SQLA with 1.11% and 2.23%. In actual fact, all the higher-order schemes perform well in this problem, with the lower computational cost estimates resulting

from the fact that the major burden in this calculation is in resolving the up-scatter source, while the leakage update is performed at the outer level. This can even be seen in the case of the SQLA solution, with less than 1% of the calculational time being devoted to resolving the leakage source. In the case of CQLA_{rlcs} this value climbs to only 5%. This is an important result, which indicates that the burden of resolving the up-scatter source in the lower order nodal solver is not directly experienced by HOTR, as it constructs the higher-order shapes based on the provided zero order quantities.

Additional detail and the results of this problem, such as multi-group flux comparisons between SQLA and CQLA, are given in Section B.4. This problem exhibits a FOM of 0.23.

5.3.4 The ZION and BIBLIS benchmark problems

The ZION and BIBLIS benchmark problems, both well known and fully described in Smith (1979), represent two similar 2D, 2-group LWR benchmarks. They differ in that the ZION problem contains an explicit baffle model, while the BIBLIS problem exhibits a smeared baffle/water region surrounding the core.

The BIBLIS benchmark problem utilizes a checkerboard pattern, formed by fuel and control elements. The fuel elements are about 23 cm in pitch and the system is surrounded by water. The problem is realistic, given that it is a 2D representation of an actual operating LWR. The ZION benchmark contains a high enriched outer zone and a lower enriched inner zone, with, as stated, an explicit outer baffle which provides particular difficulties to the SQLA solution method. These difficulties are primarily a result of the nodal aspect ratios (Smith, 1979) between the baffle and the neighbouring fuel elements, which approaches a ratio of about 7:1. Results of these two problems are presented in Table 5.10.

Table 5.10: Results for the 2D, two-group BIBLIS and ZION benchmark problems.

Solution scheme	k_{eff} **# (pcm error)**	**Cost factor**	**Power error Ave (Max) %**
BIBLIS			
Ref. in Smith (1979)	1.02512		
FHO_6	1.02511	110.81	
SQLA_{hotr}	1.02531 (20)	1.00	0.65 (2.09)
CQLA	1.02510 (1)	2.11	0.04 (0.10)
CQLA_{rlcs}	1.02507 (3)	2.92	0.16 (0.48)
ZION			
Ref. in Smith (1979)	1.27489		
FHO_6	1.27489	85.22	
SQLA_{hotr}	diverges	–	–
FLAT_{hotr}	1.27505	1.00	0.74 (2.16)
CQLA	1.27489 (0)	3.54	0.02 (0.06)

Results for these considered benchmarks coincide with those of the KOEBERG problem in that all the 2D LWR cases show nodal power errors in excess of 2% for the SQLA case. For the ZION problem, no SQLA result could be reported, due to a non-convergence of the solution and only a result with a flat, or constant, leakage approximation was obtained (denoted by FLAT_{hotr}). This observation was also made by Smith (1979), in remedy of which he applied a hybrid constant/quadratic leakage approximation, which still yielded a maximum nodal error of 1.96%.

The CQLA solution shows a good accuracy improvement, with a 20 and 35 time improvement in the maximum nodal power error for BIBLIS and ZION, respectively. The CQLA_{rlcs} solution is reported only for BIBLIS, with the performance showing some deterioration when compared to CQLA. This may again be attributed to the additional overhead cost of CQLA_{rlcs} and the associated impact thereof on these smaller 2-group, 2D problems. For the ZION problem, no converged correction factors to SQLA could be obtained, due to the non-convergence of SQLA. This is an important result concerning the capabilities of CQLA_{rlcs} and highlights the fact that

such correction factors are only feasible in cases where SQLA is of reasonable accuracy. If such failure in convergence is detected, the solution method should revert back to pure CQLA, whilst still employing both the model reduction (MR) and the leakage freezing (PLC) iteration options to accelerate the solution.

For the BIBLIS problem, the calculated FOM is 0.67, but it can unfortunately not be evaluated for the ZION problem since SQLA does not converge.

5.3.5 SAFARI-1 benchmark problem

The 3D, 6-group SAFARI-1 benchmark problem analyzed in this section is a realistic representation of an actual BOC operating state of the 20 MW SAFARI-1 tank-in-pool type research reactor operated by Necsa, at the Pelindaba site, in South Africa. The cross-sections utilized in this benchmark are extracted from the official SAFARI-1 core-follow parametrized few-group cross-section library, utilizing the state conditions of every node to reconstruct each of the 2520 different cross-sections in the model. A schematic view of the SAFARI-1 core is presented in Figure 5.9.

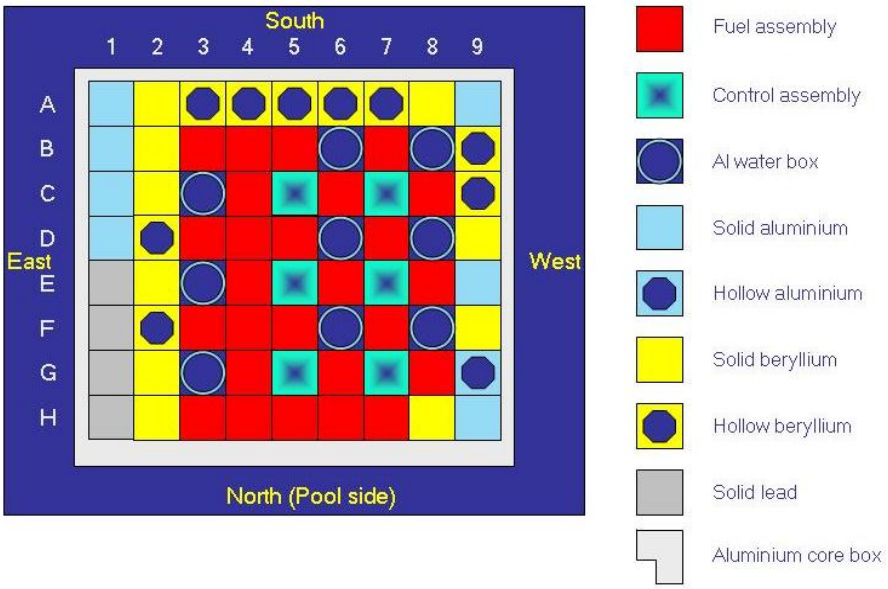

Figure 5.9: Schematic view of the SAFARI-1 benchmark core model.

The benchmark contains in-core elements such as fuel, follower type control ele-

ments, beryllium reflector elements and in-core irradiation devices which might contain a variety of samples for the purpose of isotope production or material testing. In this model, all irradiation positions are filled with aluminium water boxes. In the model, the core is surrounded with an explicitly modelled core box, which in turn is surrounded on all sides by two 10 cm water reflector nodes. Fuel assemblies have radial dimensions of 7.71 cm by 8.1 cm and as such this model exhibits smaller nodes than in the case of the LWR benchmark problems discussed before.

This model is of interest, since the use of nodal codes for research reactor modelling would not just involve the correct prediction of fluxes in fuel assemblies, but also in the other in-core components. However, prior to investigating these issues, Table 5.11 provides the standard accuracy and performance analyses we have performed for all the problems thus far. Table B.28 in Appendix B expands upon these results by providing additional detail on power density distribution errors for the SQLA and CQLA solutions.

Table 5.11: Results for the 3D, six-group SAFARI-1 benchmark problem.

Iteration scheme	k_{eff}	k_{eff} **error (pcm)**	**Cost factor**	**Ass. Ave(Max) % power error**	**Nodal Ave (Max) % power error**
FHO_6	1.021511	—	203		
SQLA	1.02288	134	1.00	0.33 (0.84)	0.40 (1.42)
FHO_2	1.02155	4	11.2	0.02 (0.14)	0.05 (0.18)
CQLA	1.02206	54	3.01	0.13 (0.49)	0.13 (0.59)
CQLA_{rlcs}	1.02206	54	1.71	0.13 (0.49)	0.14 (0.63)

Although the assembly pitch is relatively small, the SAFARI-1 benchmark problem shows large k_{eff} errors for all the approaches. Nevertheless, the results in Table B.28 are generally consistent with those obtained from previous problems, with CQLA_{rlcs} being slightly less efficient for this problem when compared to the larger LWR cores evaluated earlier in the chapter. This is primarily related to the relatively small core and overly large reflector region, causing the problem to exhibit a fairly small dominance ratio which in turn causes convergence in only 45 outer iterations. With the execution of the RLCS algorithm, 10 outers are needed to converge the SQLA correction factors in this case and thus approximately 20% of outers require higher-order solutions, when compared to the 9% of the IAEA 3D problem. This delay in

convergence relates to the slow convergence in the large reflector, which is not principally consistent with the way nodal methods should be applied, as the use of reflector equivalence parameters would be more appropriate for such a problem. Nevertheless, CQLA (and $\text{CQLA}_{\text{rlcs}}$) achieves an accuracy improvement of approximately a factor of 3 for averaged error measures and a factor of 2 for maximum error measures for this problem, at a computational cost of 1.71. The FOM for this problem assumes a value of 0.75.

As mentioned earlier in this section, the spatial and energy distribution of nodal fluxes in in-core irradiation positions are relevant in research reactor analyses. Figure 5.10 shows the SQLA errors in the unperturbed axial flux distribution for irradiation position B8, which in the SAFARI-1 reactor is typically utilized for the production of Mo-99 via the fission of specifically manufactured target plates.

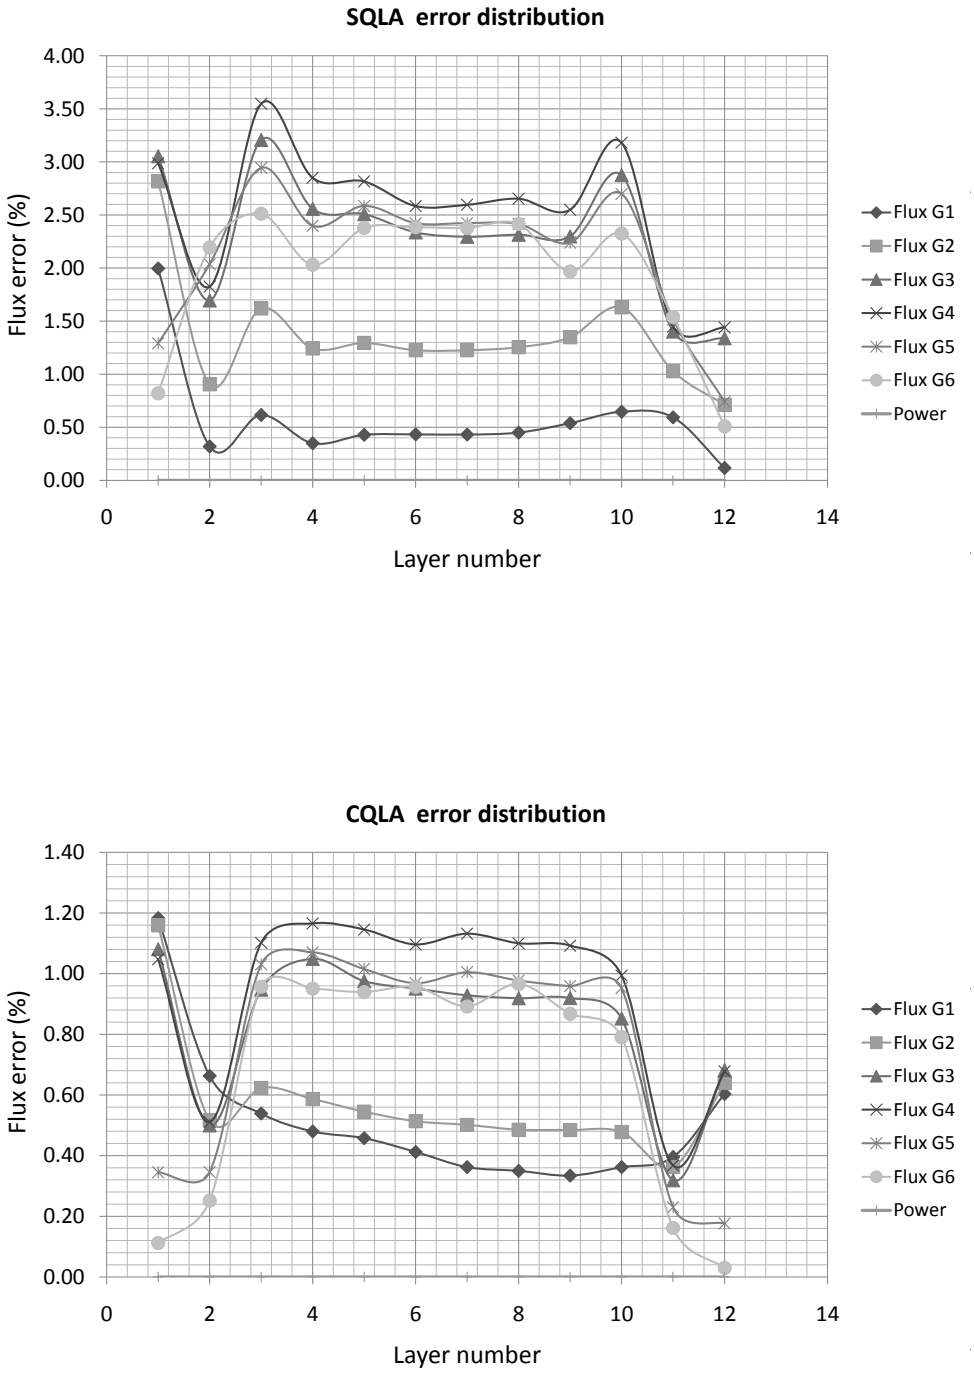

Figure 5.10: Axial flux distribution in SAFARI-1 core position B8 for SQLA and CQLA solutions, respectively.

These axial error profiles indicate that much larger flux errors occur in in-core irradiation positions when compared with fuel elements, with the maximum nodal flux errors in excess of 3% for SQLA. The CQLA solution improves these errors with around 1% and hence would provide a valuable contribution to the cycle planning and isotopic yield prediction. The SAFARI-1 reactor model is further investigated in Chapter 6, for cases where HOTR is directly coupled to the OSCAR-4 calculational system used to perform the day-to-day reload and core-follow calculations.

5.4 Summary Observations

Five fixed cross-section benchmark problems were investigated in this chapter. The solution schemes of primary interest were the SQLA solution, in order to summarize the error associated with it and the CQLA$_{rlcs}$ solution, in order to summarize the accuracy and efficiency of the proposed scheme. The overview results for these solutions are presented in Table 5.12.

Table 5.12: Summary results for the selected fixed cross-section benchmarks.

Benchmark	SQLA errors		CQLA$_{rlcs}$ errors		Calc cost
(node groups)	pcm (outers)	Ave (Max) power	pcm (outers)	Ave (Max) power	# (#/outer)
MOX 2D (18)	16.7 (32)	0.19 (0.55)	17.0 (44)	0.13 (0.16)	1.24 (0.9)
IAEA 2D (138)	3.0 (82)	0.25 (0.91)	0.8 (112)	0.04 (0.14)	2.61 (1.9)
MOX 3D (162)	19.0 (152)	0.22 (0.67)	17.0 (179)	0.15 (0.31)	1.41 (1.3)
ZION (338)	diverged	[a]0.74 (2.16)	diverged	[b]0.02 (0.06)	3.54 (–)
BIBLIS (578)	20.0 (121)	0.65 (2.09)	3.0 (170)	0.16 (0.48)	2.92 (2.5)
KOEBERG (1734)	40.0 (148)	1.11 (2.24)	5.0 (151)	0.14 (0.49)	1.09 (1.1)
IAEA 3D (2622)	4.2 (210)	0.28 (1.23)	1.8 (159)	0.08 (0.34)	1.04 (1.4)
SAFARI-1 (15120)	134.0 (55)	0.40 (1.42)	54 (50)	0.14 (0.63)	1.71 (1.8)

[a] FLAT solution, since SQLA did not converge
[b] CQLA since RLCS did not converge

Table 5.12 provides an overview of the performance and accuracy of the CQLA$_{\text{rlcs}}$ scheme. The benchmark names, followed by the number of node-groups in brackets, are given in column 1. The entries in the table are sorted according to the increasing number of node-groups, which is a measure of the problem size and are calculated as

$\text{nodes}_x \times \text{nodes}_y \times \text{nodes}_z \times$ groups. Error measures for SQLA and CQLA$_{rlcs}$, consistent with the definitions in earlier tables in this chapter, are provided in columns 2 - 5. The calculational cost of the CQLA$_{rlcs}$ scheme is given in the final column and refers to the relative performance in the total calculational time. The effective cost per outer iteration, which is a measure aiming to quantify the performance, independent of the acceleration scheme employed in the driver nodal code, is given in brackets in column 6. Note that the MOX 2D result is included for completeness.

The results in Table 5.12 indicate that the proposed scheme improves the average and maximum nodal power errors for all the problems. For the larger, more realistic core problems, such as the KOEBERG, IAEA 3D and SAFARI-1 benchmark problems (thus problems in excess of 1000 node-groups), the nodal power error measures generally improve by about a factor of four (in some cases much higher and in some slightly less), with the computational efficiency ranging from 1.04 to 1.7. The smaller problems show more erratic cost estimates, mainly due to the interaction of the overhead method cost and short total running times. Results of the ZION problem, in which SQLA and therefore also CQLA$_{rlcs}$ does not converge, indicate that the pure CQLA (not included in this table) solution remains an important solution option for cases where SQLA is severely inaccurate.

The obtained Figures Of Merits (FOMs) for the problems considered, range from 0.25 to 0.75. These numbers may be interpreted if we consider that a FOM of 0.5 indicates a method which is twice as good as SQLA in terms of both calculational cost and accuracy. In general, considering the cost/outer error measure as a better predictive performance indicator, it can be concluded that for realistic size problems, the CQLA$_{rlcs}$ scheme improves the solution accuracy by about a factor of 4, at a computation cost of between 1.1 and 1.8 times that of SQLA (which implies an average FOM of around 0.36 in terms of cost/outer). This is in line with the initial expectation of providing a significant improvement in the calculational accuracy of SQLA, at an acceptable cost penalty. To place this issue into context, consider that a 4 times reduction in error may be achieved by halving the calculational mesh in each direction when performing an SQLA solution. This would typically result in an increase in calculational cost by approximately a factor of 8 (or 800%) in 3D as compared to between 10% and 80% of the problems considered here.

It remains to be seen whether such figures are achievable when coupling HOTR to a wider spread of problems and a wider array of driver nodal codes. The upcoming chapter will touch upon this issue, by investigating the coupling of HOTR to the OSCAR-4 system as an alternative driver code.