Chapter 6

OSCAR-4 code integration

6.1 Introduction

The development of an improved transverse leakage and flux reconstruction method has been theoretically described (Chapter 3), implemented in a standalone Fortran code module (Chapter 4) and evaluated with a set of numerical fixed cross-section benchmarks (Chapter 5). The method represents a natural extension of the transversely integrated nodal methods to higher-orders and the so-called CQLA approach, coupled to appropriate iteration schemes, is shown to improve significantly upon the accuracy of the standard quadratic transverse leakage approximation. For the set of fixed cross-section problems, the computational cost associated with this improvement is well in-line with the original requirement of around a 50% increase in total calculational time.

The question however remains whether the claimed generality in coupling HOTR to an industrial nodal solver is justified. Further, an indication of the performance and accuracy of the proposed scheme within such an independently developed code system should be evaluated. Specifically of interest is the claim that the method is relatively insensitive to the solution, iteration and acceleration schemes employed by the host code.

In order to investigate these issues, this chapter describes the coupling of the HOTR module to the OSCAR-4 code system. OSCAR-4 has been developed at Necsa over a period of 20 years as the primary reload and core-follow neutronic analysis tool for the SAFARI-1 research reactor operated by Necsa. The code system employs a variety of iteration and acceleration schemes and allows the accuracy and performance of HOTR to be evaluated on realistic problems, such as actual SAFARI-1 cycle assessments and and core-follow calculations. Such problems allow the evaluation of additional levels of complexity such as multi-step calculations, material bur-

nup, cross-section feedback and critical control rod position search to be included in the analysis of HOTR. It should be stated that no specific effort has been made to optimize the integration between HOTR and OSCAR-4, as is surely needed in a final coupling (as done for the SANS link) and it could be expected that the performance obtained would at this stage not directly compete with results from the SANS-HOTR link from Chapter 5 in terms of calculational time. The underlying reason for this is that OSCAR-4 is a highly optimized code package and as such the nodal diffusion solver in OSCAR-4, called MGRAC (discussed in more detail later on in this chapter), would out-perform SANS on similar problems. HOTR then, would require similar optimizations as those in MGRAC during a final implementation, with special attention given to issues such as calculational statement optimization, basic loop parallelization, streamlined memory and data management, topology integration between HOTR and MGRAC and detailed iteration scheme integration. This additional work is intended as part of a final implementation step following from the completion of this thesis. Nevertheless, conclusions regarding the accuracy and performance of HOTR coupled to MGRAC in its current form, in a relatively straight-forward manner, should still prove valuable information with regard to issues around the accuracy, complexity and performance of the proposed schemes.

In this chapter, we briefly describe the OSCAR-4 system, discuss the coupling strategy to HOTR and evaluate the accuracy and performance of both fixed cross-section and steady state calculations with feedback. The MGRAC code proves to be a useful driver platform in this regard, since a large number of so-called "advanced" extensions to the basic nodal model are implemented in MGRAC and allow these options to be independently evaluated in conjunction with HOTR.

6.2 The OSCAR-4 Code System

The OSCAR nuclear reactor calculational software package has been incrementally developed over 20 years. The package is the primary calculational tool for the SAFARI-1 (Ball, 1999; Stander et al., 2008) research reactor operated by Necsa in South Africa, which is widely accepted as one of the best commercialized research reactors in the world, given its refocus toward medical isotope production. OSCAR also became the primary calculational system for the R2 reactor at Studsvik, Sweden (Reitsma et al., 2004) while it was still operating and is currently used at the HFR reactor at Petten in the Netherlands (Hendriks, 2006) and the HOR reactor at TU-Delft (Leege and Reitsma, 2001) in the Netherlands.

OSCAR-4 is the latest version of the OSCAR code system. OSCAR-4 employs the traditional deterministic calculational path, hence utilizing transport solvers for spatial homogenization and spectral condensation and then full core multi-group nodal diffusion solutions for the global reactor problem. This process is schematically illustrated in Figure 6.1.

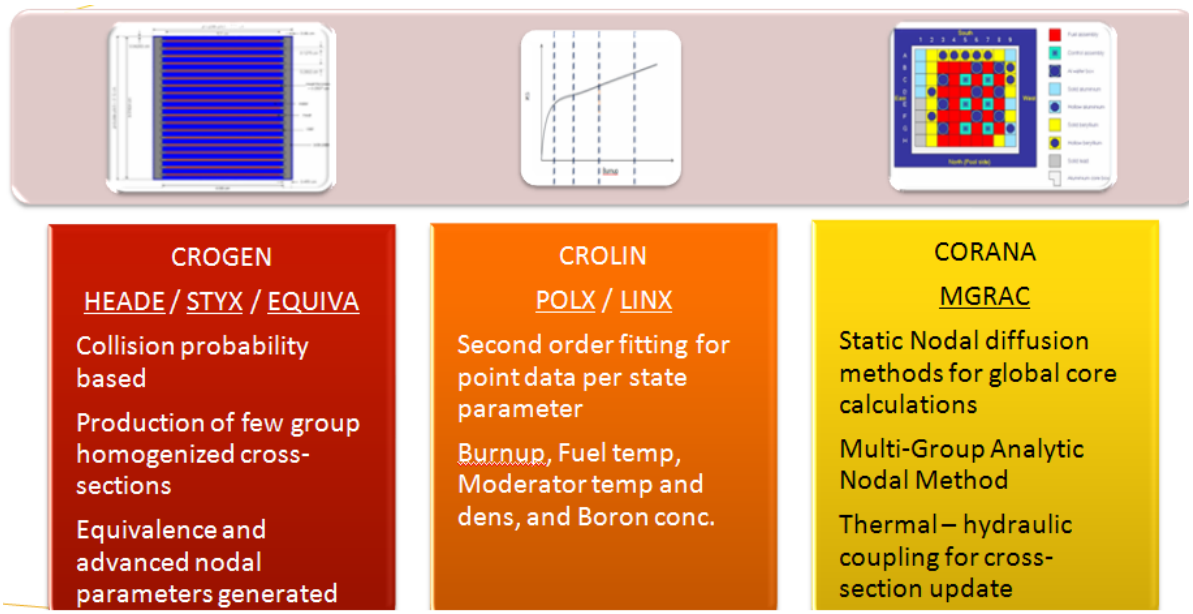

Figure 6.1: Schematic breakdown of the OSCAR-4 system.

In Figure 6.1, the CROGEN subsystem refers to the subset of transport solvers (HEADE, STYX, EQUIVA) utilized for cross-section generation, the CROLIN subsystem to the cross-section parametrization tools (POLX, LINX) and the CORANA subsystem to the global nodal diffusion solver MGRAC. MGRAC is of specific interest in this work, since it is to this code within the OSCAR-4 system that HOTR is coupled. The capabilities of MGRAC is described in somewhat greater detail in the following section, as summarized from Stander et al. (2008).

6.2.1 The MGRAC nodal diffusion solver

In MGRAC, the calculation of the steady-state neutron flux distribution throughout a reactor core and part of its non-fuel peripheral regions is based on the solution of the three-dimensional multi-group time-independent diffusion equation by means of a modern transverse-integration nodal method for Cartesian geometry. This nodal method, which is known as the Multi-group Analytic Nodal Method (MANM) (Vogel

and Weiss, 1992), engages an analytic solution to the one-dimensional transverse-integrated multi-group diffusion equation in order to determine a relationship between node side-average net currents and node-average fluxes. It is subject to only one approximation, namely that of a finite-order polynomial approximation for the transverse leakage inhomogeneous source term in the one-dimensional equation.

By reordering the terms of the equation system, different iteration strategies that exhibit different iteration behaviour are obtained (Muller and Wiederhold, 1995). Iteration acceleration methods include the Wielandt eigenvalue shift (Sutton, 1988) combined with the asymptotic fission source extrapolation for the outer iterations and the Chebyshev Cyclic Semi-Iterative (CCSI) method for the inner iterations. The default scheme in MGRAC is the so-called transverse leakage source iterative method (TLSIM) which utilizes the leakage source as a driver for the eigenvalue problem, instead of the fission source.

OSCAR-4 utilizes a microscopic depletion model. Depletion history tracking in MGRAC involves both fuel exposure and nuclide (an arbitrary number of actinides, fission-products and burnable absorbers) inventory tracking. In MGRAC the depletion tracking mesh (the exposure mesh) is quite independent of the neutronic mesh. An exposure mesh is assigned to each component (fuel assembly, control rod, detector string, reflector assembly, irradiation rig, etc.) individually at the beginning of life of the component (at the time the component is first introduced into the calculation system). Since the exposure mesh is now coupled to a component, this mesh can be moved in unison with a component that is subjected to axial motion. This feature, often termed axial homogenization (Reitsma and Muller, 2002), has proven to be particularly useful in capturing the depletion history of the follower type control rods (consisting of a fuel assembly section connected to a control rod section by a non-burnable coupling piece) that are employed in many research reactors such as SAFARI-1. The generation of averaged nodal cross-sections on the global calculational mesh is achieved via a flux-volume weighting procedure to average cross-sections from the material mesh. This axial homogenization procedure utilizes single channel diffusion fluxes from an auxiliary refined sub-node mesh calculation as a cross-section weighting function.

A predictor-corrector method is used for the depletion calculations, thus involving two converged nodal flux solutions per burn-up step. A constant-flux explicit time integration method is optionally available for faster calculations with reduced accuracy requirements. One of the most significant improvements that have been incorporated in the new version of MGRAC is the capability to deplete burnable

absorbers (rods, wires, plates) as heterogeneous objects using local reconstructed heterogeneous fluxes and heterogeneous (not smeared over assembly radial domains) microscopic cross-sections. These models allow for a much more accurate treatment of both the reactivity and depletion effects of burnable absorbers.

6.3 Coupling of HOTR to OSCAR-4

In principle, HOTR is coupled to MGRAC in the same way as described for the coupling between SANS and HOTR in Section 4.3. A number of additional levels of complexity however exist, due to additional features of the MGRAC code. These include:

1. Various iteration schemes implemented in MGRAC, such as both the standard fission source iterative scheme (SFSIM) and the transverse leakage source iterative method (TLSIM) as discussed by Muller and Wiederhold (1995). In TLSIM the transverse leakage source is utilized as the primary driving source for the eigenvalue problem and this method has been shown to converge significantly faster that SFSIM;

2. The MGRAC strategy for solving the multi-group steady state nodal diffusion equations simultaneously via the Multi-group Analytic Nodal Method (MANM) described by Vogel and Weiss (1992), as compared to the group-by-group solution implemented in SANS and analyzed in Chapter 5;

3. The selective use of flux moments from the driver code vs the requirement that HOTR recalculates the zero-order one-dimensional flux moments independently. This is indeed the case for the MGRAC coupling, since the one-dimensional flux moments in MGRAC are only updated for use with the intra-nodal cross-section shape source and thus are not necessarily available when needed by HOTR. In this implementation, one-dimensional flux moments of the zero-order transversely-integrated equations are generated by HOTR, as compared to receiving them as additional data items from MGRAC, as in the case of the coupling with SANS. This functionality of HOTR is useful for coupling to driver nodal codes which do not necessarily have the capability to calculate one-dimensional flux moments;

4. The possibility of multiple calculational cases in a single MGRAC run (including time stepping and depletion);

5. The possibility to update the state conditions of fuel assemblies (or other core components) between calculational cases, thus requiring cross-section updates;

6. The possibility of component movement and thus adjustment of material cross-sections. This includes the use of the axial homogenization procedure as described by Reitsma and Muller (2002) to counteract the so-called "cusping effect" which appears due to the placement of the control rod tip between axial mesh boundaries;

7. The possibility of control rod search calculations which require material movement within a single calculational case; and

8. Advanced non-linear extensions to the standard nodal solution, such as the use of intra-nodal cross-section shapes to compensate for the shortcomings of "block depletion" (or depletion on the coarse nodal mesh).

These eight additional features of MGRAC, not available in SANS and thus not yet evaluated, are the focus of this chapter. It is of interest to evaluate the stability and performance of the proposed leakage algorithms, as implemented in HOTR, against these additional levels of complexity and in so doing identifying areas of further development to facilitate a final "industrial" implementation. Section 6.4 investigates these issues on both the 3D IAEA LWR benchmark and an example of the SAFARI-1 reload and core-follow analysis for an actual historic cycle.

Taking into account a simplified representation of the algorithmic layout of the MGRAC code, Algorithm 3 describes the implemented interaction between MGRAC and the HOTR module.

Algorithm 3: MGRAC solution algorithm

Input: Component placement, problem geometry and case structure
Result: Eigenvalue, nodal flux and power distribution per case

```
HOTR calc init  initialization
1               for m ← 1 to numcases do
HOTR case init      initialize case
2                   while Feedback level iterations not converged do
3                       for c ← 1 to num_core_channels do
                            Move control rods if needed
                            Perform axial homogenization if needed
HOTR xs update              Update nodal cross-sections if needed
                            Update coupling coefficients if needed
                        end
4                       for o ← 1 to num_outers_per_feedback do
                            Update driving outer level source according to chosen scheme
                            Update power distribution and eigenvalue
                            Update side-averaged fluxes and side-averaged currents for all
                            groups
HOTR leakage                Calculate transverse leakage source for all groups
                            Update intra-nodal cross-section shape source if needed
5                           for o ← 1 to num_inners_per_outer do
                                Update tensorial sources for balance equation
                                Perform multi-group inner sweep according to chosen scheme
                            end
                            Update one-dimensional flux moments if needed (for intra-nodal
                            cross-section shape source)
                        end
                    end
                end
```

Algorithm 3 describes the various iteration levels in MGRAC, which are limited to a feedback level, an outer iteration level and an inner iteration level. These iteration levels are applied within each calculational case, as specified by the user. We note that, in the coupling to HOTR proposed here, four interaction points are defined (in the left-hand margin).

- **HOTR calc init** refers to the passing of the geometric and material data of the problem as set out for the initial calculational case;
- The **HOTR case init** call is utilized to reset iteration parameters for the next case, while allowing the latest values of higher-order moments to be carried over to the next case as an initial guess;

- The **HOTR xs update** call updates the higher-order nodal coupling coefficients (when the material layout changes for whatever reason). The reason for the update is taken into account (e.g. update of control position vs update of intra-nodal source shape) in order to determine the appropriate response in HOTR; and
- **HOTR leakage** is the call to update the leakage coefficients. Thus the **HOTR case init** and **HOTR xs update** interfaces have been additionally introduced in order to facilitate the multi-case structure and material configuration changes allowed by MGRAC.

The usage of four interaction points increases the complexity of the coupling and additional information about the state of the calculation is passed from the driver code to HOTR. Further analysis of the MGRAC code structure could yield improved methods of interaction and the coupling proposed here is most likely non-optimal. The efficiency of this proposed scheme is evaluated in the next section.

6.4 MGRAC-HOTR Analysis

Two problems are selected to investigate the performance of HOTR against the set of capabilities identified in Section 6.3. In order to verify the implementation in MGRAC, we firstly select the 3D version of the IAEA LWR benchmark problem, which was discussed extensively in Chapter 5. For this problem, the reference solution is known from the previous chapter.

The second problem is actually a series of calculations needed to perform reload and core-follow analysis for a given operating cycle of the SAFARI-1 reactor. Although no explicit reference solution is available for the problems, the expected accuracy of the CQLA approach (and its associated iteration schemes) is already known from the fixed cross-section version of this problem analyzed in the previous chapter. We are thus primarily interested in the stability and performance of various calculational options when applying the coupled MGRAC-HOTR code to these cases.

6.4.1 The 3D IAEA LWR benchmark problem

For this problem, calculations are performed for various iteration schemes in MGRAC in order to test the correctness of the coupling. We are thus largely evaluating points 1-3 from the list proposed in Section 6.3 regarding this problem and leave the other items for the evaluation of the SAFARI-1 problem for Section 6.4.2.

We evaluate this problem then for three different iteration/acceleration schemes, namely SFSIM with fission source extrapolation (as was the case during the coupling with SANS), SFSIM with both fission source extrapolation and Wielandt eigenvalue shift acceleration and TLSIM. Results are summarized in Tables 6.1 to 6.3.

Note that the SQLA results differ somewhat from those presented in Chapter 5, since the MGRAC model is a full-core representation of the benchmark problem, where the model in the previous chapter utilized a quarter core symmetry segment with reflective boundary conditions applied to the centre of the subdivided fuel assemblies (see the discussion in Section 5.3.2). The use of the 10 cm subdivided fuel assembly nodes at the model edge produces a different three-node quadratic fit for the transverse leakage approximation as compared to the three full 20 cm nodes in the case of the full symmetry model utilized here. This does not influence the higher-order or reference results. Ideally, the implementation of symmetry conditions in SANS would remove this difference, but this was not done in the scope of the SANS test code development.

Table 6.1: Results for the 3D, two-group IAEA LWR benchmark problem for HOTR as coupled to MGRAC, for the SFSIM iteration scheme.

Code and scheme	**Iteration scheme**	**Acceleration scheme**	k_{eff} **# (pcm, outers)**	**Cost factor MGRAC (SANS)**
[a]Reference	SFSIM	Extrap.	1.02907	–
MGRAC (SQLA)	SFSIM	Extrap.	1.02916 (9, 128)	[b]1.00 (1.00)
MGRAC (CQLA-FHO_2)	SFSIM	Extrap.	1.02906 (1, 109)	6.27 (6.01)
MGRAC (CQLA-PLC)	SFSIM	Extrap.	1.02909 (2, 109)	2.75 (1.81)
MGRAC ($CQLA_{rlcs}$)	SFSIM	Extrap.	1.02909 (2, 103)	1.55 (1.04)

[a]Reference as produced by SANS-HOTR in Chapter 5

[b]Both MGRAC and SANS SQLA were scaledd to a relative performance of 1 for comparison. In actual fact, MGRAC is about 50% faster than SANS for the same problems.

Table 6.1 gives k_{eff} accuracy and solution performance for HOTR coupled to MGRAC, which is the same iteration scheme utilized in the previous chapter in the SANS host code - namely the standard fission source iterative method (SFSIM) with fission source extrapolation. The final column in the table shows the performance factors obtained with HOTR coupled to MGRAC, followed by the corresponding result

obtained for the SANS coupling (as given in Table 5.6) in brackets. Here the CQLA solution in HOTR is combined with the partial leakage convergence (PLC) iteration scheme, since this approach was shown to provide a useful performance improvement without any notable accuracy penalties.

The table shows a reasonable comparison between the efficiency of CQLA-PLC in the coupling of HOTR with SANS (value of 1.81) as compared to the coupling with MGRAC (value of 2.75). As expected, HOTR proves somewhat less efficient when coupled to MGRAC than that of SANS, but the relative improvement when utilizing the CQLA$_{rlcs}$ scheme compares quite well (2.75 to 1.55 vs 1.81 to 1.04). The performance of the FHO$_2$−PLC solution (full higher-order with partial leakage convergence) compares very well with the previously obtained results. All the k_{eff} results are identical to the results obtained in Section 5.3.2. In general, we may conclude that this full symmetry model of the IAEA 3D benchmark problem performs similarly in MGRAC for the SFSIM scheme with the fission source extrapolation as in the case for the SANS host code.

Table 6.2 continues the iteration scheme analysis by utilizing the Wielandt eigenvalue shift acceleration option in MGRAC, followed by Table 6.3 which employs the transverse leakage source iterative method (TLSIM) option.

Table 6.2: Results for the 3D, two-group IAEA LWR benchmark problem for HOTR as coupled to MGRAC for the SFSIM iteration scheme with Wielandt acceleration.

Code and scheme	**Iteration scheme**	**Acceleration scheme**	k_{eff} **# (pcm, outers)**	**Cost factor** #
Reference	SFSIM	Extrap.	1.02907	–
MGRAC (SQLA)	SFSIM	Wielandt+Extrap.	1.02916 (9, 19)	1.00
MGRAC (CQLA-PLC)	SFSIM	Wielandt+Extrap.	1.02909 (2, 25)	3.91
MGRAC (CQLA$_{rlcs}$)	SFSIM	Wielandt+Extrap.	1.02909 (2, 25)	4.40
MGRAC (CQLA$_{rlcs}$)[a]	SFSIM	Wielandt+Extrap.	1.02909 (2, 23)	2.91

[a]Relaxed criteria for correction factor convergence

Table 6.3: Results for the 3D, two-group IAEA LWR benchmark problem for HOTR as coupled to MGRAC for the TLSIM iteration scheme.

Code and scheme	**Iteration scheme**	**Acceleration scheme**	k_{eff} **# (pcm, outers)**	**Cost factor** #
Reference	SFSIM	Extrapolation	1.02907	–
MGRAC (SQLA)	TLSIM	—	1.02916 (9, 29)	1.00
MGRAC (CQLA-PLC)	TLSIM	—	1.02909 (2, 32)	3.82
MGRAC (CQLA$_{rlcs}$)	TLSIM	—	1.02909 (2, 29)	3.84
MGRAC (CQLA$_{rlcs}$)[a]	TLSIM	—	1.02909 (2, 27)	2.31

[a]Relaxed criteria for correction factor convergence

Table 6.2 shows results for the case of the highly efficient Wielandt acceleration, combined with fission source extrapolation and Table 6.3 follows with the same performance quantification for the TLSIM option in MGRAC. It should be stated that both these schemes provide a performance improvement of around a factor of 5 as compared to SFSIM with extrapolation. We note that in these two tables the picture is somewhat different. The performance of CQLA-PLC is still generally comparable to previously obtained estimates, but the CQLA$_{rlcs}$ scheme shows an effective degradation in efficiency. The underlying reason for this deterioration is related to the dramatic decrease in the outer iterations required for convergence by both the Wielandt acceleration scheme (decrease of around a factor of 6) and the TLSIM iterations scheme (decrease by a factor of 4).

For the case of Wielandt acceleration, the number of higher-order outers needed to converge the quadratic leakage corrections (as needed by CQLA$_{rlcs}$) also decreases (from 14 without Wielandt to 9). However, it is clear that the decrease in higher-order iterations does not occur at the same rate as for the overall number of outers and as such the fraction of higher-order outers during Wielandt acceleration effectively increases. This is to some extent understandable, since Wielandt acceleration significantly improves the rate of convergence of the fission source distribution, but does not directly improve the estimation of the leakage source. Further, large "jumps" in the leakage coefficients due to the accelerated convergence of the fission source could influence the convergence estimation of the correction factors.

With this in mind, Tables 6.2 and 6.3 both contain a second CQLA$_{rlcs}$ result, in which a possible remedy to this problem is illustrated. For both these additional results, the convergence criteria for the SQLA correction factors are relaxed by about

one order of magnitude, which decreases the number of higher-order outers needed for convergence, whilst in both cases (Wielandt and TLSIM) maintaining the accuracy of the solution. The associated calculational cost improves noticeably. A possible explanation for this behaviour may be found in the fact that the fast rate of convergence of the driving source in the system could fictitiously delay the numerical estimates for the convergence of the correction factors and as such lead to eventual over-convergence. An investigation into the rate of convergence for these factors indeed shows that, in the case of Wielandt, the maximum obtained convergence is between one and two orders of magnitude greater than for the standard SFSIM at the termination of the higher-order iterations. This interplay between convergence estimates and iteration schemes is an area which requires additional future work, which would be needed in a final "industrial" implementation of HOTR into any given host code.

Nevertheless, from Tables 6.2 and 6.3 we conclude that for the current implementation, the CQLA_{rlcs} scheme does not retain its efficiency when applied to a solution or a problem with a very small dominance ratio (as is exhibited by both Wielandt and TLSIM). The CQLA-PLC scheme however, still approximately scales as expected. The SFSIM scheme, with fission source extrapolation results, generally confirm the conclusions of Chapter 5.

In order to investigate these and the rest of the eight issues highlighted for evaluation in MGRAC, we now proceed to the realistic case of performing a reload and core-follow analysis for an actual SAFARI-1 operating cycle.

6.4.2 SAFARI-1 reload and core-follow analysis

In this section we revisit the SAFARI-1 core model, but in this case for the realistic calculational scenario of performing the reload and core-follow calculation for operating cycle C1001-1 (first cycle of 2010). This model is selected primarily since detailed model and plant data is available for this problem, given that OSCAR-4 is the primary calculational tool of SAFARI-1. As could be seen from the fixed cross-section benchmark of SAFARI-1 discussed in the previous chapter, this model does not highlight the shortcomings of SQLA as clearly as larger power reactor problems do, but nevertheless provides a realistic calculational environment to confirm that the HOTR code is correctly integrated into OSCAR-4. This definition of "realism" originates from the additional set of requirements posed by the reload and core-follow calculational procedure, listed below (which were not evaluated in the case of the fixed cross-section SAFARI-1 benchmark in Section 5.3.5). A brief explanation of

how the new issue is addressed in HOTR is included with each item. During reload and core-follow analysis:

- **Cross-sections** are obtained from a parametrized cross-section library for each node, depending on the state conditions of that node. In this polynomial cross-section library, cross-sections are tabulated against assembly burnup, moderator temperature, moderator density and fuel temperature. HOTR was not particularly adapted for this purpose, since cross-sections are passed to HOTR only after they were reconstructed within each node. What is of importance, is that cross-sections are, for various reasons, sporadically updated during the MGRAC calculation and an additional interface was defined to allow this update in HOTR;

- **Multi-case calculations** occur within a single MGRAC calculation, which requires that the stepping from case to case is correctly implemented. Cases in MGRAC refer to various nodal solutions ordered within a single MGRAC calculation. Iteration parameters should be correctly reset from case-to-case, while the solution obtained in a given case should be correctly utilized as an initial guess in the subsequent case. This is particularly important for the values of the higher-order moments in order to speed up convergence. In the case of RLCS, iteration parameters have to be carefully controlled in order to allow an update of the SQLA correction factors for the new case, without reverting to a pure SQLA solution for the first couple of outers. If material placement changes from case to case (control rod movement, rig movement or material depletion), parameters such as nodal coupling coefficients and (in the case of RLCS) quadratic leakage correction factors should be updated within HOTR;

- **Control rod search calculations** are utilized for the determination of critical rod positions, which require multiple nodal solutions within a single case and typically induce only small perturbations due to the movement of control rods by the control rod search algorithm. There is no need to activate the higher-order solution prematurely, given that SQLA is sufficiently accurate to identify the critical control rod position to within some reasonable convergence. In this implementation the higher-order solution is activated when the k_{eff} is within 250 pcm from the target k_{eff} (which is user input). Furthermore, movement in control rod placement implies that a change in nodal cross-sections and parameters such as nodal coupling coefficients should again be updated. In the

case of RLCS, correction factors generally become invalid once the material and geometric layout of the problem is adjusted and as such the tabulation of correction factors is repeated after each control rod movement;

- **Depletion calculations** are needed, which requires a multi-case calculation, spanning the entire cycle. Between calculational cases, material depletion is performed and depending on the depletion method (explicit time integration versus a predictor-corrector type approach) different numbers of nodal solutions are needed to perform the depletion. For example, in a predictor-corrector solution (with only a single corrector step), depletion is performed with both the beginning of step flux as well as the end of step flux, after which the number densities are averaged;

- Control rods, amongst others, are modelled as **moveable components** and as such are placed in various locations depending on the scenarios defined for the modelling. When activated, the axial homogenization procedure implemented in MGRAC plays a role in these cases. This axial homogenization solution implemented in MGRAC utilizes a one-dimensional solution in the channel of interest, on a finer submesh, to determine an axial flux profile used for the flux-volume weighting of the nodal cross-section on the calculational mesh. Since cross-sections on the calculational mesh are passed to HOTR, it is expected that no additional complexity is needed in HOTR to handle such a capability in the driver code. However, the result of axial homogenization is an updated set of cross-sections in control rod channels, which requires an update to HOTR;

- **Various core states** are considered, depending on the parameters of interest. These states include hot or cold conditions, zero xenon or xenon equilibrium conditions and configurations with different irradiation rig loadings. The latter primarily refers to the state of the Mo-99 production rigs (positions C3, E3, G3, B6, B8, D8 and F8 in Figure 5.9), which are essentially mini-fuel elements for the production of Mo-99 from fission. The core states are mostly reflected in the cross-sections passed to HOTR and no specific intervention is needed; and

- Data regarding the **intra-nodal cross-section shape**, as accumulated from the historical burnup of the fuel elements in past cycles, is available for use and thus, when activated, the intra-nodal cross-section shape source would contribute to the solution. The issue here is that this source is not explicitly

represented in the higher-order equations (as is surely possible with further development) and the question is whether treating the source only in the zero-order equations leads to stability problems in the coupled solution.

Considering the last point above, the issue is that in standard zero-order transversely-integrated methods we assume that both an intra-nodal flux shape and an intra-nodal reaction rate are represented with separable expansions as:

$$\phi(u,v,w) = f(u) + g(v) + h(w) - 2\overline{\phi}$$

and

$$\Sigma(u,v,w)\phi(u,v,w) = \Sigma(u)f(u) + \Sigma(v)g(v) + \Sigma(w)f(w) - 2\overline{\phi\Sigma}.$$

In the higher-order approach we assume the full tri-variate expansion for an intra-nodal flux (as discussed in previous chapters), but we do not have to assume that the intra-nodal reaction rate is represented via this same intra-nodal tri-variate expansion. We are thus free to choose any representation for a reaction rate and in particular one that would lead to the set of quasi-one-dimensional equations with an intra-nodal cross-section shape that is zero in the higher-order equations (and non-zero in the zero-order equations). This approach would give the expected results (a set of zero- and higher-order equations) with a weighted transverse integration.

However, if we use the variational approach, we would not be able to obtain a variation of the functional with respect to an independent variation of the unknown functions ($f_{kl}(u)$), because we have a different flux representation in the reaction rate and in the leakage terms. This issue should, in future work, be carefully examined, since it may be problematic to obtain a tri-variate expansion of the reaction rate from assembly calculations (and from codes that solve depletion equations and thermal-hydraulic states).

All of these listed issues are involved in performing SAFARI-1 reload and core-follow calculations and allow for the evaluation of the additional MGRAC features (see Section 6.3) defined earlier.

6.4.2.1 Description of the SAFARI-1 reload and core-follow procedure

The SAFARI-1 reload and core-follow calculational process involves a number of MGRAC calculations, simulating various reactor states in an effort to calculate a defined set of safety and utilization parameters. The so-called "reload" calculations are performed in preparation of every operating cycle, to confirm that the characteristics of the upcoming cycle and the "core-follow" calculations are performed in order

to update the depletion state of the fuel elements once the cycle is complete. The process is graphically illustrated in Figure 6.2.

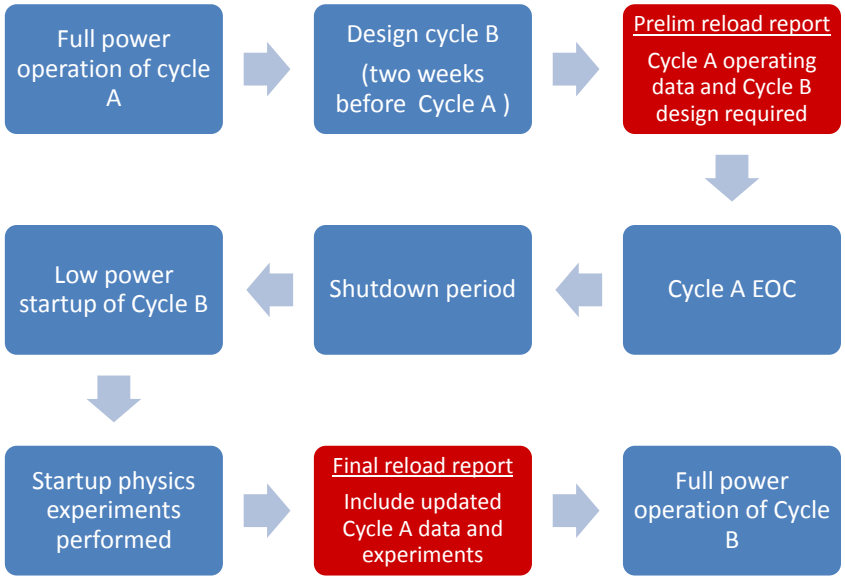

Figure 6.2: Reload and core-follow calculational procedure (as a flow chart) as applied to the SAFARI-1 research reactor.

Figure 6.2 describes the calculational requirements, fulfilled by the OSCAR-4 system, during the transition from any given cycle (say Cycle A) to the following cycle (Cycle B). In the figure EOC refers to End of Cycle. The production of both the preliminary reload report and the final reload report (highlighted in the figure) requires the following minimum set of MGRAC calculations, in order to determine the necessary parameters:

1. **A BOC operating critical calculation** is a MGRAC calculation containing a single calculational case at hot, zero xenon conditions with all the irradiation rigs in their typical operating states. The control rods are placed in their critical positions. The critical position is of course not known and a critical control rod search is required within the code;

2. **A BOC equilibrium xenon calculation** is a MGRAC calculation containing a single calculational case at hot, equilibrium xenon conditions with all the irradiation rigs in their typical operating states. The control rods are placed

in their critical positions via a critical rod search. This calculation is primarily utilized to extract typical power and flux distributions for the cycle, as needed for core utilization and irradiation rig load planning. Quantities such as the expected Mo-99 production yield for the cycle is of primary interest;

3. **A rod worth calculation** is a MGRAC calculation containing three subsequent calculational cases. The first with the control rods all out, the second with the control rods all in and the third with the control rods at their critical position. This calculation is performed at cold, zero xenon conditions with all the irradiation rigs in their most reactive state. Control rod worth, shutdown margin and excess reactivity of the core are estimated from this calculation;

4. **A power peaking calculation** is a MGRAC calculation containing two subsequent calculational cases, at BOC and after 1 day of full power operation, at hot conditions. In both cases, the primary parameters of interest are the maximum assembly power in the core, the maximum nodal power in the core and the maximum pin-power in the core (in this case referring to the highest power in a "piece of fuel plate"). The latter is determined from a power reconstruction calculation; and

5. **The cycle depletion** (or core-follow) calculation is a MGRAC calculation simulating the depletion of the reactor during a cycle, utilizing plant data to define the step sizes, power levels and control rod positions for every calculational case. This calculation is used to determine the EOC isotopic distribution in the fuel elements;

In this section, we repeat this set of calculations, as is needed for the production of the preliminary reload report in the third step of Figure 6.2. As stated, we perform these calculations specifically for the reload of reactor cycle C1001-1 and further perform the core-follow calculation for C1001-1 as would typically be done at the end of the cycle. These various calculations represent a thorough test of the coupling of MGRAC to HOTR.

After these calculations are performed, the set of safety and utilization parameters are extracted from the various result files and collated into a single reload and core-follow report. We perform the reload and core-follow calculations for various MGRAC option sets, utilizing the SQLA, CQLA-PLC and CQLA$_{rlcs}$ approaches in HOTR.

Prior to investigating the performance, a subset of the obtained parameters is compared between SQLA and CQLA (we obtain the same results via CQLA, CQLA-PLC and CQLA$_{rlcs}$) in Table 6.4. Note once more that we do not expect to find any larger discrepancies between SQLA and CQLA than reported in Section 5.3.5, but rather quantify the associated impact on parameters which the reactor operator or nuclear engineer would find more interesting.

Table 6.4: Summary results for the SAFARI-1 reload parameters as compared between SQLA and CQLA solution methods in MGRAC.

Reload Parameter	**SQLA result**	**CQLA result**	**Difference** #
Control Bank Worth	28.00 \$	27.97 \$	22 pcm
Excess Reactivity	10.02 \$	10.03 \$	7 pcm
Shutdown Margin	15.34 \$	15.30 \$	29 pcm
Relative Power Peaking Factor	3.95	3.99	1.01%
Critical Bank without Targets	47.44 cm	47.11 cm	0.69%
Critical Bank with Targets	43.71 cm	43.68 cm	0.07%
Total Mo-99 yield	3896 Ci	3863 Ci	0.85%
Flux in Irradiation position D6	2.12×10^{14}	2.09×10^{14}	1.41%

As demonstrated in Section 5.3.5, the relatively small fuel elements (7.71 cm $\times$ 8.1 cm) in SAFARI-1 do not present a significant problem to SQLA and results in Table 6.4 largely supports this claim. Of potential interest is an under-estimation of the power peaking factor by SQLA of about 1% as well as the overestimation of both the Mo-99 yield and the D6 thermal flux by 0.85% and 1.41%, respectively. Reactivity parameters are generally very well predicted by SQLA. The important conclusion to be drawn from this table is that the HOTR module seems to have been correctly integrated within the MGRAC code for the mixture of scenarios required to obtain this subset of parameters (a set of five calculations described earlier in this section).

We move on to the investigation of the performance of the HOTR module for these calculations and compile Table 6.5. The table lists the performance of CQLA-PLC and CQLQ$_{rlcs}$ on both a single critical search calculation (such as the BOC critical bank calculation) and a core-follow calculation spanning a 30 day operating cycle of SAFARI-1 (eight cases, utilizing fixed control bank positions from the plant at each step).

Table 6.5: Performance matrix for CQLA-PLC and CQLA$_{rlcs}$ against various solution options in MGRAC.

Calculation and Scheme	**Acceleration**		**Non-linear extensions**[a]	
	SFSIM	**SFSIM + Wielandt**	**Axhom**	**Intra-nodal model**
Critical Bank				
CQLA-PLC	1.75	2.93	X[b]	1.85
CQLA$_{rlcs}$	1.93	3.12	X	2.11
Core-follow				
CQLA-PLC	2.13	3.11	2.60	2.20
CQLA$_{rlcs}$	2.22	4.16	2.92	2.43

[a]Based on SFSIM solution

[b]Axial homogenization, coupled with critical rod searches, did not converge for SQLA or any higher-order solution

Table 6.5 illustrates the performance factors of CQLA-PLC and CQLA$_{rlcs}$ for various input options in MGRAC. Columns two and three compare the efficiency of the schemes for the case of SFSIM with extrapolation, with that of SFSIM with both extrapolation and Wielandt acceleration. Columns four and five investigate the impact of the two non-linear extensions evaluated, namely the axial homogenization option and the intra-nodal cross-section shape feedback option. These two cases are activated whilst utilizing the SFSIM option (with extrapolation) as the base solution.

Two scenarios are analyzed in the table, namely a single case critical search and a core-follow calculation with 8 flux and depletion steps. It should be noted that, since the higher-order calculation is only activated once the k_{eff} is within 250 pcm of the target value, a large portion of the CQLA related schemes still proceed via SQLA and hence produce somewhat flattering performance figures. Although this illustrates that HOTR performs acceptably in conjunction with control rod searches, more realistic timing comparisons of the performance in MGRAC may be extracted from the timing comparisons of the core-follow calculation.

Performance figures for SFSIM compare reasonably well with those obtained for the fixed cross-section version of the SAFARI-1 benchmark in Section 5.3.5, where HOTR was coupled to the SANS test code. A slight degradation is however visible for both the critical search and core-follow scenarios for the CQLA$_{rlcs}$ scheme, but can be understood considering the "blind" nature of the coupling performed between

MGRAC and HOTR, with no specific focus on optimizing data structures or integrating iteration schemes. Furthermore, the cross-section updates which occur during control rod movement, introduces some inefficiencies in the CQLA_{rlcs} scheme, given that the tabulation of correction factors is restarted after the movement. A possible improvement in this regard could relate to the triggering the the re-tabulation based on the magnitude of the cross-section perturbation, as compared to an indiscriminate re-tabulation.

Once again, as in the case of the 3D IAEA LWR problem earlier in this chapter (Table 6.2), the relative inefficiency of the CQLA_{rlcs} scheme, when Wielandt acceleration is activated, is clearly visible (highlighted in red). For the SAFARI-1 core, with a naturally small dominance ratio (large reflector region and small core), the number of outers needed for convergence decreases by a factor of 2 due to Wielandt acceleration and, as in the case of the IAEA LWR benchmark in Table 6.2, degrades the efficiency of the scheme. The discussion in Section 6.4.1 remains relevant and this observation shall be a focus of future study. It is noted that the CQLA-PLC solution is generally more stable in its performance measure across the various iteration schemes.

The final two columns indicate the variation in the obtained performance when axial homogenization and intra-nodal cross-section shape models are switched on, respectively, during the MGRAC calculation. We simply note that these models do not significantly influence the operation of HOTR and introduce a small degradation in computational time due to additional feedback iterations which they induce.

6.5 Conclusions

This chapter aimed at evaluating the impact of additional "realistic" calculational requirements on the operation of the HOTR module, or more specifically on the solution and iteration schemes introduced in Chapters 3 and 4, respectively. A number of issues were identified for evaluation and the question was posed regarding how strongly the accuracy and performance of the developed schemes are influenced by the introduction of these additional levels of complexity. Table 6.6 presents a summary of these findings, based on the 3D IAEA LWR benchmark and the more comprehensive SAFARI-1 reload and core-follow processes investigated in this chapter (including some experience from the previous chapter).

Table 6.6: Summary of HOTR code performance as coupled to the MGRAC nodal solver in OSCAR-4.

Host code capability	**Stability and accuracy**		**Performance**	
	CQLA	**CQLA$_{rlcs}$**	**CQLA**	**CQLA$_{rlcs}$**
SFSIM with group-by-group solution (SANS)	√	√	√	√
Simultaneous multi-group solution (MGRAC)	√	√	√	√
Standard fission source driven scheme	√	√	√	√
Transverse leakage source driven scheme	√	○[a]	○	×
Wielandt acceleration	√	√	○	×
Multi-case environment	√	√	√	○
Control rod-search calculations	√	√	√	○
Depletion calculations	√	√	√	○
Axial homogenization	√	√	√	×
Intra-nodal cross-section shape feedback	√	√	√	○
One dimensional moments from driver code	√	√	√	√
One dimensional moments from HOTR	√	√	√	√

[a] "All out" case during control rod worth calculation did not converge

In Table 6.6 the various capabilities which optionally exist in host nodal codes (and particularly those which are available between SANS and MGRAC) are listed. The stability and performance of HOTR, specifically regarding the CQLA and CQLA$_{rlcs}$ schemes, are then evaluated utilizing the following criteria:

- √ denotes a successful integration with HOTR. In the stability column this implies consistent smooth convergence for the set of problems investigated and in the performance column it implies that CQLA converged with a performance factor of less than 3.5 and CQLA$_{rlcs}$ converged with a performance factor of less than 2.0 for all the cases considered;

- ○ denotes a successful integration with HOTR. In the stability column this implies consistent smooth convergence for most problems, but some unexplained divergence for some problems. In the performance column it denotes that CQLA converged with a performance factor of less than 4.0 and CQLA$_{rlcs}$ converged with a performance factor of less than 2.5 for all the cases considered;

- $\times$ denotes a failure during the integration with HOTR. In the stability column this implies consistent divergence of the solution and in the performance column it denotes that CQLA converged with a performance factor of more than 4.5 and CQLA$_{rlcs}$ converged with a performance factor greater than 2.5 for any of the cases considered.

Table 6.6 summarizes the overall capability of HOTR and takes into account many of the issues which are often lumped into the concept of "practical" or "realistic" steady state reactor core calculations. Both methods (CQLA and CQLA$_{rlcs}$) generally converge smoothly and it is clear that CQLA performs consistently (at an average cost factor of about 3 times that of SQLA when combined with PLC). CQLA$_{rlcs}$ on the other hand has a number of occurrences of low and even weak performances, which most probably indicate a need for closer integration between the driving code iteration scheme and HOTR. In particular, these occur for TLSIM and SFSIM with Wielandt acceleration, in which case the number of outer iterations is dramatically reduced. The investigation and resolution of this issue is left as future work.

We may conclude that the implementation of HOTR inside MGRAC shows good agreement with set requirements for most of the cases considered in Table 6.6. Outliers are however identified for future analysis and shall be investigated once a final implementation of HOTR within MGRAC is performed.