FULL-CORE, 2-D, LWR CORE CALCULATIONS WITH CASMO-4E

Kord S. Smith and Joel D. Rhodes, III
Studsvik Scandpower, Inc.
504 Shoup Ave., Suite 201
Idaho Falls, ID 83402

ABSTRACT

Modern large-memory computers have made possible direct application of lattice transport methods to large-scale computational models of LWR cores. However, many acceleration techniques common to lattice transport codes are not applicable to heterogeneous-geometry, high-dominance-ratio, eigenvalue problems. Consequently, large heterogeneous LWR problems have remained computationally intensive and impractical for routine applications.

This paper presents new synthetic acceleration methods that achieve speedups of more than a factor of 100 on large LWR problems. These acceleration techniques are non-linear, and they permit dramatic acceleration without geometrical limitations normally associated with traditional acceleration techniques, such as DSA. These new acceleration methods have been incorporated into the CASMO-4E lattice physics code, and they make practical the application of lattice transport methods to 2-D, whole-core, LWR computations - complete with fuel depletion and shuffling.

1. INTRODUCTION

Various versions of the CASMO code [1-2] have been used for many years to perform neutron and gamma transport computations for LWR lattices, which include thousands of computations for each type of LWR lattice introduced into an operating reactor. These computations include fuel assembly depletions and branches to all anticipated conditions (e.g., fuel temperatures, coolant temperatures, void conditions, boron concentrations, rod insertions, etc.).

The 2-D transport models in CASMO-3 [1] used rectangular cells with homogenenized pin-cell data. Those models have been replaced in CASMO-4 [2] with completely heterogeneous models solved using the Method of Characteristics (MOC) [3,4]. Accurate analysis of these lattices requires the use of very detailed spatial mesh. The assembly physical mesh (Figure 1) is divided into ~5000 regions, with each "pin-cell" being split into radial and azimuthal zones (Figure 2). Neutron sources in each zone are approximated as spatially flat and isotropic, and the heterogeneous 2-D geometry is preserved.

Fig. 1 Physical Geometry

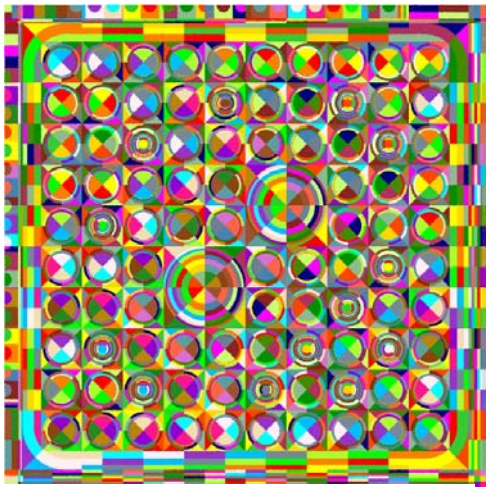

Fig. 2 Flat-Source Geometry

Discretization of the azimuthal angular variable in the Boltzman transport equation is achieved by laying down a set of finely spaced parallel rays across the entire problem for each of the desired azimuthal angles (which are approximately uniformly spaced), as depicted by the solid parallel lines in Figure 3. The precise ray spacing and azimuthal angles are chosen such that each azimuthal ray has a perfectly reflecting (or translating) counterpart in another of the azimuthal angles, as depicted by the dashed lines of Figure 3. This choice of cyclic azimuthal tracking allows all boundary conditions to be treated without approximation. Discretization of the polar angle variable is achieved by replicating the azimuthal rays for each desired discrete polar angle (i.e., a product quadrature set). This "long characteristic" discretization is common in MOC applications.

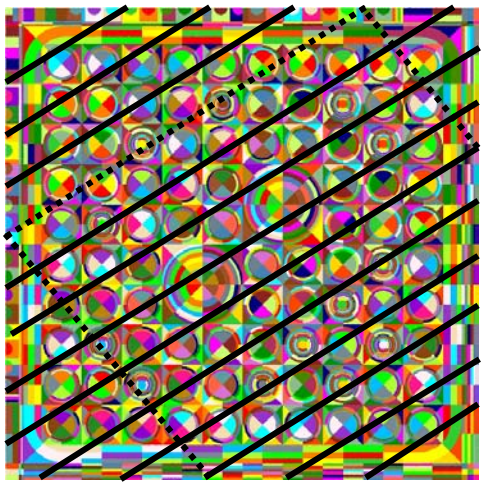

Fig.3 Characteristics Ray Tracing

For the case of isotropic scattering, the relationship between the outgoing and incoming angular flux for group g, azimuthal angle i, polar angle j, for track k, across flat-source zone m is given by:

$$\varphi_{g,i,j,k}^{m+} = \varphi_{g,i,j,k}^{m-} e^{-\Sigma_g^m \tau_k / \cos_j} + \frac{Q_g^m}{4\pi\Sigma_g^m}(1 - e^{-\Sigma_g^m \tau_k / \cos_j}), \qquad \text{(Eq. 1)}$$

where τ_k is the in-plane length of track k, and $\cos_j$ is the cosine of the polar angle (from the 2-D plane). The spatially averaged angular flux, in group g, azimuthal angle i, polar angle j, along track k is:

$$\bar{\varphi}_{g,i,j,k}^{m} = \frac{Q_g^m}{4\pi\Sigma_g^m} + \frac{(\varphi_{g,i,j,k}^{m-} - \varphi_{g,i,j,k}^{m+})}{\Sigma_g^m \tau_k / \cos_j}. \qquad \text{(Eq. 2)}$$

A full core 2-D BWR computation (~764 assemblies plus reflectors) in 16 energy groups has on the order of 100 million scalar flux unknowns. When the method of characteristics is used, the radial track density required to accurately predict the angular flux distribution is approximately 0.5 million tracks per assembly, resulting in approximately 500 million radial tracks (for each group and each of 3 to 5 polar angles). If careful attention is paid to compressing storage to a minimum, the computer storage required for this scale problem is a few Gigabytes. The simplicity of Eqs. 1 and 2 leads to a basic transport sweep which is extremely efficient, having only a few operations per track.

The principal problem that arises in going from single-assembly computations to full-core reactor computations is that the solution of the eigenvalue problem becomes far more difficult. This is easy to understand, since the lattice calculation with reflecting boundary conditions has a very low dominance ratio, while 2-D cores typically have dominance ratios of 0.99 or greater. Consequently, straightforward application of unaccelerated power iterations (as often used in normal lattice applications) requires thousands of iterations to achieve convergence in large 2-D LWR problems.

Conventional transport acceleration techniques, such as DSA [5] are difficult, if not impossible, to apply to the MOC equations. The heterogeneous nature (e.g., arbitrary collections of rectangles,

cylinders, and arcs of cylinders in CASMO-4) of the MOC representation makes it difficult for these authors to see how to define proper acceleration equations for DSA. Transport acceleration methods (TSA) have also been applied recently to acceleration of the MOC in-group scattering [6], but they still fall far short of the computational efficiency required for large-scale LWR problems. In order to overcome such limitations, two very different approaches to non-linear acceleration have been implemented into CASMO-4E. Both techniques permit acceleration of MOC equations and, perhaps, other transport formulations.

2. DIFFUSION ACCELERATION OF MOC SOLUTIONS

The non-linear iteration techniques commonly used in advanced nodal diffusion codes [7,8] (to accelerate higher-order diffusion models with low-order diffusion models) can also be used to accelerate directly the MOC transport solutions. This can be achieved by superimposing a rectangular acceleration mesh on the heterogeneous geometry, and tallying volume-averaged scalar fluxes and surface-averaged net currents during each transport sweep. Following each transport sweep, a set of finite-difference-like, coarse-mesh diffusion equations are constructed on the rectangular mesh. Coarse-mesh cross sections are computed by flux-volume weighting cross sections of each heterogeneous sub-region contained in each coarse mesh. Diffusion-like coupling coefficients for the coarse-mesh equations are constructed to match the average net current (from the transport sweep) on each face of the rectangular mesh (assuming that the coarse-mesh group fluxes are correct), as expressed (for group g and the plus side of coarse mesh M) in the following equation:

$$J_g^{M+} = -D_g^{M+}(\bar{\phi}_g^{M+1} - \bar{\phi}_g^{M}) - \tilde{D}_g^{M+}(\bar{\phi}_g^{M+1} + \bar{\phi}_g^{M}) . \qquad \text{(Eq. 3)}$$

D_g^{M+} is computed using the standard finite-difference expression for the weighting of diffusion coefficients, mesh spacing, and surface area from two neighboring nodes, and $\tilde{D}_g^{M+}$ (the non-linear coupling coefficient) is simply chosen such that Eq. 3 is exactly correct when using the scalar fluxes and net currents from the latest transport sweep. Substitution of Eq. 3 into the 2-D scalar neutron balance equation, results in a set of coarse-mesh, diffusion-like equations with five-point coupling. In the rare case that $\tilde{D}_g^{M+}$ is greater in magnitude than D_g^{M+}, new values of $\tilde{D}_g^{M+}$ and D_g^{M+} are selected such that the magnitudes of the two coefficients are identical. This flux-limiting condition [9] is required so that the coarse-mesh equations are guaranteed to be diagonally dominant.

The coarse-mesh diffusion equations are constructed in CASMO-4E on the "pin-cell" mesh, and these equations are solved using standard fission source and Gauss-Seidel inner iterations. There are no restrictions on the structure or complexity of flat-source regions contained in each coarse mesh regions, and each acceleration mesh can contain an arbitrary collection of objects or partial objects that are further subdivided into flat-source regions. This is quite different from standard DSA techniques for which the spatial mesh is identical in the transport and diffusion operators.

The ratio of converged-to-initial coarse-mesh diffusion theory scalar fluxes are used to update flat-source region fluxes prior to the n+1 transport sweep by using the n-th transport sweep fluxes (for flat-source region m contained in coarse mesh M):

$$\bar{\phi}_g^{m,n+1/2} = \bar{\phi}_g^{m,n}\left(\frac{\bar{\phi}_g^{M,\infty}}{\bar{\phi}_g^{M,0}}\right) . \qquad \text{(Eq. 4)}$$

The fine mesh "correction factors" are multiplicative (unlike the additive corrective factors normally associated with DSA), and easier to apply in eigenvalue problems in which flux levels may change by orders of magnitude in early iterations. Note from Eqs. 3 and 4 that when the diffusion acceleration is converged, the fine-mesh flux correction factors are unity. Boundary angular fluxes are scaled similarly, by using the scalar flux ratios for the boundary coarse mesh.

In large BWR core configurations, it is impractical to overlay a contiguous rectangular acceleration mesh because the core may consist of many different types of fuel assemblies (e.g., 8x8 fuel pins, 9x9 fuel pins, and 10x10 fuel with water crosses, etc.). Consequently, a contiguous mesh would consist of many very narrow meshes that make the coarse-mesh diffusion equations difficult to converge. In CASMO-4E, the acceleration mesh is defined to be contiguous only within each fuel assembly. Figure 4 displays an acceleration mesh for two hypothetical assemblies (one with 5x5 fuel pins and the other with 4x4 fuel pins). In such a case, the current coupling equations for an interface are constructed such that the net currents in Eq. 3 are satisfied when using the scalar fluxes for the mesh under consideration and the neighboring coarse mesh having the closest geometrical center. Consequently, the diffusion equations retain their five-point coupling (four neighbors for each node), even though a node may have a larger number of physical neighbors. Note that the coupling matrix will be nonsymmetric in such a case. (Actually the non-linear coupling coefficients result in a nonsymmetric matrix even for contiguous acceleration meshes.)

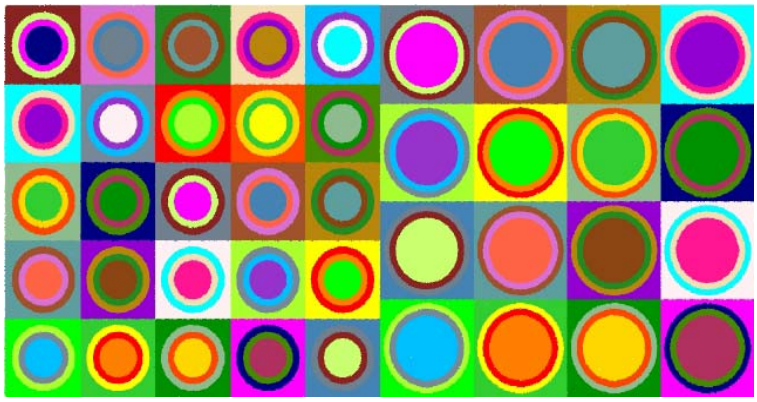

Fig. 4 Non-Uniform Rectangular Diffusion Acceleration

3. MACRO-TRACK TRANSPORT ACCELERATION OF MOC SOLUTIONS

The second acceleration method that has been implemented into CASMO-4E is called the "macro-track" method. In this method, neither the rectangular acceleration mesh nor the diffusion approximation is used. Rather, a set of coarse "macro-tracks" are defined to lie directly on the MOC tracks, as depicted schematically for two macro-tracks of one azimuthal angle in Figure 5. Each macro-track (composite colored line segments) consists of an integral number of "micro-tracks" from the basic MOC tracks (individual colored line segments). Macro-track lengths and optical path lengths are known from the basic MOC tracking information. Consequently, as the basic MOC sweeping is performed, it is possible to define spatially flat "equivalent" sources (using the macro-track analogue of Eq. 1) which exactly preserve the outgoing angular fluxes of the macro-track, as presented in Eq. 5. Thus, one constructs a set of "coarse-mesh" MOC equations that (when converged) exactly preserve the outgoing angular flux of the "fine-mesh" MOC solution along each macro-track. The macro-tracks can be chosen in a number of ways (e.g., a maximum optical path length, a fixed number of fine-tracks, etc.), but the method used in

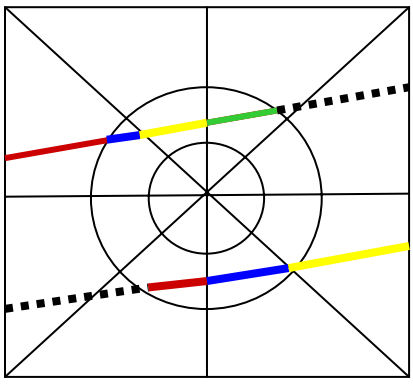

Fig. 5 Macro-Line Geometry

CASMO-4E is simply to make the length of each macro-track to be shorter than a user input length. When chosen this way, the macro-tracks conform to no regular boundaries, and the macro-track mesh is geometrically "unstructured". This feature is very convenient for geometries in which no regular structure exists.

$$\widetilde{\left(\frac{Q}{4\pi\Sigma}\right)}_g^{M,0} = \frac{\left(\varphi_{g,i,j,k}^{M+} - \varphi_{g,i,j,k}^{M-}\right)}{\left(1-e^{-(\Sigma\tau)_M/\cos_j}\right)} + \varphi_{g,i,j,k}^{M-} . \qquad \text{(Eq. 5)}$$

In order to use a macro-track flux iteration to accelerate the basic MOC solution, it is necessary to relate the flat-source region fluxes to the macro-track fluxes and the flat-source region fluxes to the macro-track "equivalent" sources. The flat-source region scalar fluxes at macro-track iteration n+1 are approximated by assuming that the flat-source region fluxes can be scaled with the change in macro-track flux integrals for all macro-tracks that cross flat source region m,

$$\bar{\phi}_g^{m,n+1} = \left(\frac{\bar{\phi}_g^{m,0}}{\sum_M \varphi_g^{M,0}\tau_k/\cos_j}\right)\sum_M \varphi_g^{M,n+1}\tau_k/\cos_j . \qquad \text{(Eq. 6)}$$

This approximation assumes that the fine-mesh angular flux distributions along all macro-tracks (which cross flat source region m) have fixed relative magnitudes (taken from the latest MOC transport sweep), as depicted schematically in Fig. 6. Given the flat-source region fluxes, the "equivalent" sources for macro-track M, at iteration n+1, are approximated by the change in the macro-track integrated sources for all flat-source regions crossed by macro-track M, as given in Eq. 7.

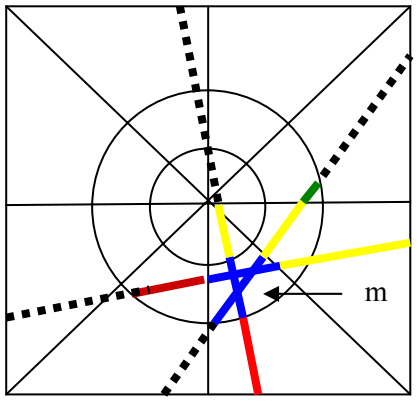

Fig. 6 Flux Approximation

The macro-track acceleration method is simply a non-linear acceleration technique in which the low-order operator is itself a coarse-mesh MOC method. The coarse-mesh operator is used to perform power iterations so that the fine-mesh MOC equations are swept only to update estimates of acceleration parameters.

$$\widetilde{\left(\frac{Q}{4\pi\Sigma}\right)}_g^{M,n+1} = \left[\frac{\widetilde{\left(\frac{Q}{4\pi\Sigma}\right)}_g^{M,0}}{\sum_m Q_g^{m,0}\tau_k}\right]\sum_m Q_g^{m,n+1}\tau_k \qquad \text{(Eq. 7)}$$

Note that unlike the diffusion method, the macro-track acceleration method requires computation of fine-mesh (flat-source region) fluxes at each iteration of the low-order operator. The macro-track acceleration method, as implemented in CASMO-4E, can be described by the following iterative procedure:

- Flat-source region scalar fluxes are set to unity
- Flat-source region sources are computed
- Fine-mesh MOC transport sweep is performed:
    - Macro-track equivalent sources are computed, using Eq. 5
    - Boundary angular fluxes are updated
    - Flat-source region scalar fluxes are accumulated
    - Flat-source region sources are updated
    - Coarse-mesh macro-track equivalent sources are approximated, using Eq. 7
    - Coarse-mesh macro-track equations are swept
    - Flat-source region scalar fluxes are approximated, using Eq. 6
    - k-eff is updated; scalar flux and k-eff convergence are tested
- Boundary angular fluxes are updated
- Scalar flux and k-eff convergence are tested

The macro-track method requires a small amount of initial overhead to establish the correspondence between macro-tracks and the traversed flat-source regions. In the CASMO-4E implementation of the macro-track and diffusion acceleration methods, each energy group is swept only once between source updates. No attempt is made to converge the in-group scattering sources or the group-to-group scattering sources (including up-scatter) for a given fission source – all sources are updated after each transport sweep.

4. DIFFUSION ACCELERATION TESTING

Previous applications of CASMO-4 [2] to all types of LWR lattices have shown that any simple acceleration technique (coarse mesh rebalance, source extrapolation, Chebyshev acceleration, etc.) can be used to obtain convergence of lattice problems typically in 6-10 MOC transport sweeps. However, the simple methods break down as leakage effects become important. Consequently, the efficiency of the new non-linear acceleration methods is examined here by focusing on LWR core calculations. The first case is a ¼-core model of a 16-bundle BWR critical assembly, with central control rod, surrounded by a water reflector. The MOC representation used eight energy groups, 64 azimuthal angles, three polar angles, and 0.1 cm ray spacing, and the physical geometry is displayed in Figure 7.

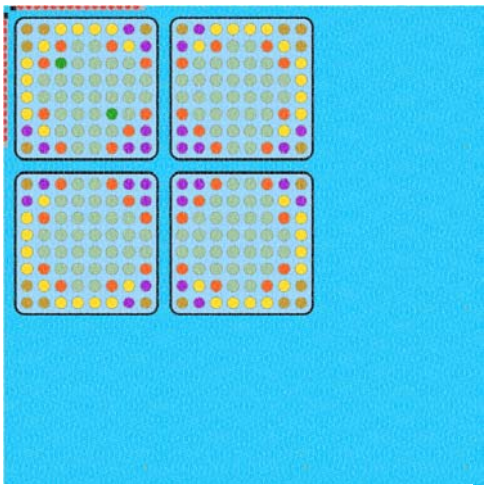

Fig. 7 Critical Assembly Geometry

Table I compares results obtained by power iteration without acceleration and with numerous applications of the diffusion acceleration method. From these results, it can be seen that the diffusion acceleration reduces the number of transport sweeps required by roughly a factor

of 20. The dominance ratio of this diffusion problem is quite small (< 0.9), and convergence is nearly independent of the number of fine-mesh diffusion iterations.

Table I Diffusion Acceleration Results	
Number of Power Iterations	# of MOC Transport sweeps
No Acceleration	436
10 power iterations/MOC transport sweep	23
25 power iterations/MOC transport sweep	23
50 power iterations/MOC transport sweep	24
100 power iterations/MOC transport sweep	23

5. MACRO-TRACK ACCELERATION TESTING

By virtue of the fact that diffusion acceleration has been shown to be very effective in accelerating the MOC convergence in LWR problems, we can deduce that the low-order acceleration method does not require extensive resolution of azimuthal and polar variation of angular fluxes. Consequently, the macro-track acceleration procedure described above has actually been simplified in the CASMO-4E implementation to improve the efficiency of the macro-track transport sweeps. This is done by using only one polar angle for the macro-track sweeps, which implies that changes in flat-source region scalar fluxes are assumed to be proportional to the changes angular flux for the selected polar angle (or equivalently, that all angular fluxes in a flat-source region change by the same fractional amount). The number of azimuthal angles is also allowed to be any subset of the MOC azimuthal angles, motivated by the hypothesis that high-order angular resolution should not be required for the acceleration equations.

Fine-angle boundary fluxes are updated after completion of the macro-track iterations by assuming that all polar angle fluxes change by the same ratio as the single polar angle macro-track boundary fluxes. Interpolation is used to fill in the fine-angle azimuthal angular fluxes from the coarse-angle azimuthal macro-track boundary fluxes.

Table II presents results from parametric studies used to test application of the macro-track acceleration method for the critical assembly test case. From the results presented in the first pair of columns in Table II, one can see that about 100 power iterations per MOC transport sweep are required to converge this problem in the minimum number of transport sweeps. Note the larger number of iterations than was required for the diffusion acceleration method. The reason for this behavior is that the macro-track transport sweeps require convergence of the in-group scattering source (which does not exist in the diffusion model), and this results in a larger observed “dominance ratio” when in-group scattering sources are not fully converged at each power iteration.

As previously described, the maximum macro-track length also needs to be chosen. Results presented in the second pair of columns in Table II show that the convergence rate of the macro-track method is relatively insensitive to the selection of the maximum macro-track length. A maximum macro-track length of 1.0 cm appears to be near optimum, and this choice results in an average of about five MOC fine-mesh tracks per macro-track. Note that in the limit of small maximum macro-track length, Eqs. 5-7 require no spatial approximation, and the macro-track method becomes simply a lower-order (i.e., fewer polar and azimuthal angles) transport acceleration on the MOC spatial mesh – much like traditional TSA methods.

Results presented in the third pair of columns in Table II show that the convergence rate of the macro-track method is insensitive to the number of azimuthal angles used in the macro-track sweeping. These results suggest that eight azimuthal angles are sufficient.

It should be noted that, in principle, there is no reason to use all of the parallel rays for the chosen azimuthal angles. However, in the CASMO-4E implementation of the macro-track method, this option has not been pursued because the flat-source regions are very small, and it is necessary that all flat-source regions be traversed by some macro-track.

Table II Performance of Macro-track Acceleration vs. Macro-track Parameters

Azimuthal angles = 8 Macro-track < 1.0 cm		Azimuthal angles = 8 100 power iterations/sweep		Macro-track < 1.0 cm 100 power iterations/sweep	
# Source Iterations/Sweep	# MOC Sweeps	Macro-track <	# MOC Sweeps	# of Azimuthal Angles =	# MOC sweeps
10	49	2.0 cm	19	4	20*
25	22	1.5 cm	16	8	10
50	12	1.0 cm	11	16	9
100	11	.50 cm	10	32	9
200	11	.25 cm	10	64	9

* The four-angle set in CASMO-4 is not symmetric with respect to physical octant symmetry.

This small critical assembly has also been analyzed for cases in which the MOC model has 1, 2, 3, 5, or 7 polar angles. All cases required the same number of fine-mesh MOC transport sweeps. Consequently, the choice of using only one polar angle in the macro-track acceleration appears not to introduce degradation of convergence in the macro-track method.

By selecting one polar angle, 8 azimuthal angles, and 1.0 cm maximum macro-track length, the number of angular fluxes required in the macro-track power iteration has been reduced by a factor of 120 (5 for track length, 8 for azimuthal angles, and 3 for polar angles) relative to the number of angular fluxes used in the basic MOC transport sweep. Consequently, the amount of computational effort consumed by the macro-track acceleration method is dramatically reduced. It is noteworthy that in this difficult problem, the macro-track method converges in less than half the number of MOC transport sweeps than were required for the diffusion acceleration method.

6. LARGE GEOMETRY APPLICATIONS

Both of the non-linear acceleration methods have been applied to a ¼-core model of a 764 assembly BWR at hot zero-void conditions with a water reflector. This problem uses the same detailed flat-source representation and angular flux resolution used in the small critical assembly problem. The physical geometry of this problem is displayed in Figure 8.

The unaccelerated MOC solution, while converging smoothly, fails to achieve convergence in 2500 power iterations. Diffusion accelerated MOC results are displayed in the first pair of columns in Table III. From these results, it can be seen that about a thousand diffusion-theory power iterations are needed to minimize the number of transport sweeps. This, of course, is expected since the dominance ratio of the diffusion iteration is approximately 0.99 in this problem.

Since the dominance ratio of the diffusion problem is known to be independent of the diffusion solution mesh, it is natural to extend the single-level acceleration to a multi-level diffusion acceleration. Formal multi-grid acceleration methods [10] have been applied to the S_n method in the TWODANT code with great success. Cho [11] has used a single-level rebalance scheme, and Kosaka [12-13] has applied a two-level rebalance scheme to successfully accelerate MOC methods for large scale LWR problems. (Cho and Kosaka have not reported divergence problems with rebalance that have been observed by others [5-6].) The non-linear diffusion acceleration method presented here is also directly extendable to a multi-level approach [14]. For simplicity (as well as for edits needed for CASMO-4E) a simple two-level scheme has been adopted in which the coarse-mesh diffusion uses an assembly-size mesh (~15.0 - 20.0 cm mesh). The fine-mesh diffusion is still defined on the "pin-cell"

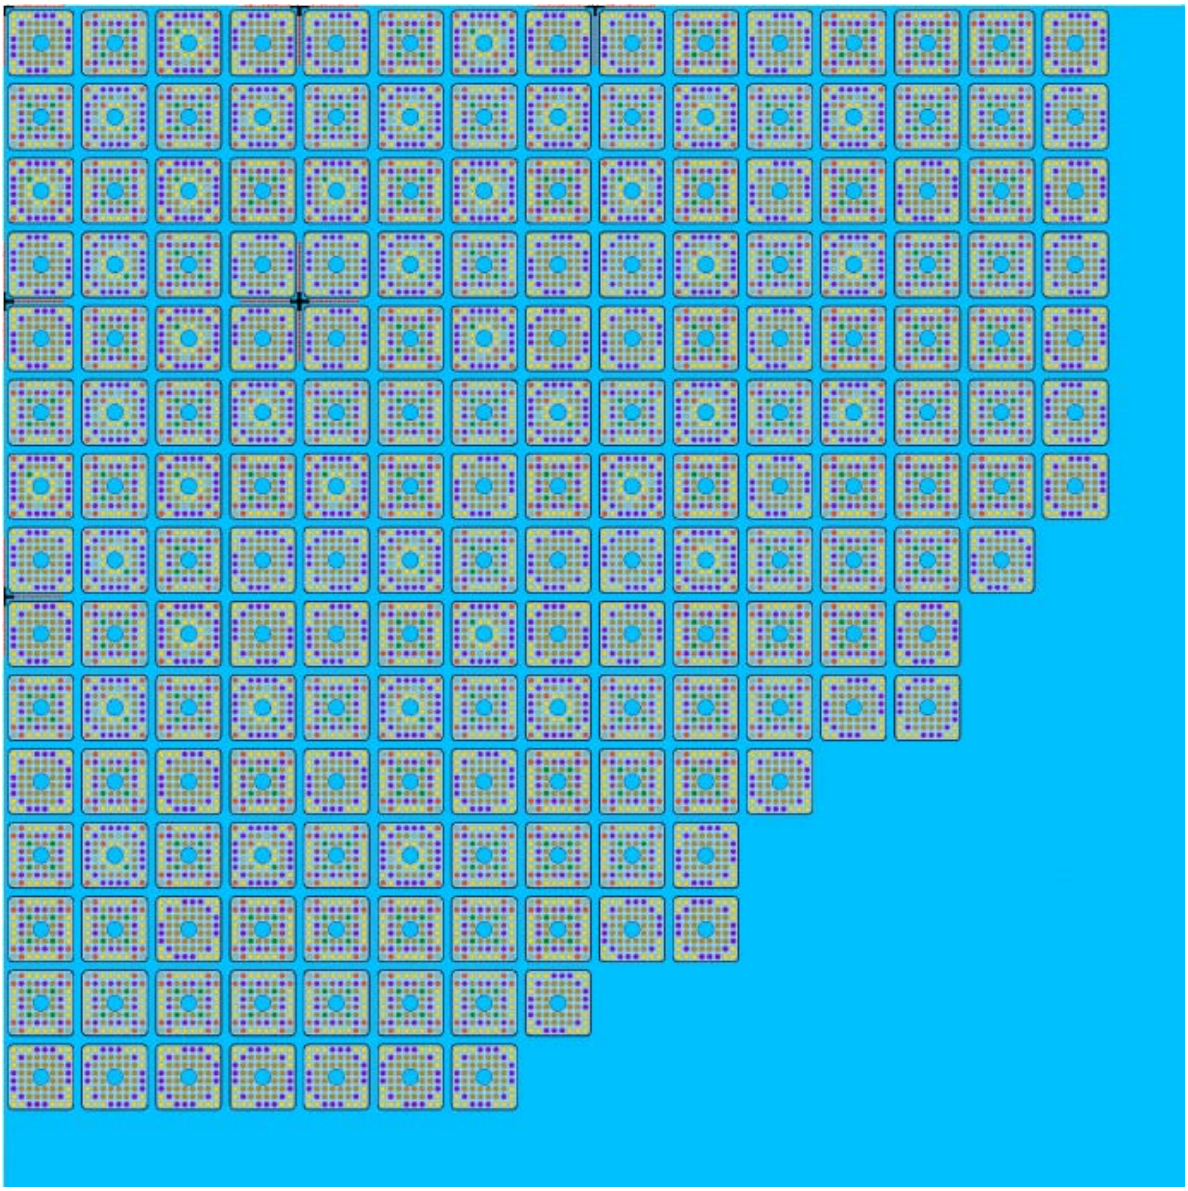

Fig. 8 Quarter-Core BWR Geometry

mesh. Assembly-averaged scalar fluxes and assembly-surface-integrated net currents from the fine-mesh diffusion solution are used in Eq. 3 to define the non-linear coupling coefficients that preserve

the surface-averaged net currents of the fine-mesh diffusion solution. These coarse-mesh diffusion equations are solved by the same source and inner iteration method applied to the fine-mesh diffusion problem.

Results obtained by using this simple two-level scheme (with 500 coarse-mesh power iterations between fine-mesh diffusion solutions) are displayed in the second pair of columns in Table III. From these results, it is easy to see that the coarse-mesh diffusion iterations effectively eliminate the most slowly converging modes, and that very few fine-mesh diffusion iterations are needed to minimize the number of basic MOC transport sweeps. Note that the coarse-mesh (assembly-level) diffusion problem requires so little CPU time that it is essentially negligible.

Table III Acceleration Results for a Large ¼-Core BWR Problem

Single-Level Diffusion Acceleration		Two-Level Diffusion Acceleration		Macro-Track Acceleration	
# Fine-mesh Diffusion Iterations/ MOC Sweep	# MOC Sweeps	# Fine-mesh Diffusion Iterations/ MOC Sweep	# MOC Sweeps	# Macro-track Power Iterations/ MOC Sweep	# MOC Sweeps
0	>2500	-	-	0	>2500
25	30	1	21	100	37
50	20	2	12	250	16
100	12	5	10	500	12
200	10	10	10	1000	12

Results obtained by application of the macro-track acceleration method to this large BWR problem are summarized in the third pair of columns in Table III. These results show that the macro-track acceleration method is also very effective in reducing the number of fine-mesh MOC sweeps needed to converge this large BWR problem. As pointed out previously, the macro-track method requires a larger number of power iterations than does the diffusion method because of the necessity of converging the in-group scattering sources. The observed dominance ratio for the transport power iteration is approximately 0.998 for this problem.

7. DISCUSSION

The two-level diffusion acceleration has been adopted as the default acceleration model in CASMO-4E. It has been applied to many tens of thousands of single-assembly LWR problems, to thousands of multi-assembly LWR problems, and to hundreds of quarter-core and full-core LWR problems. This experience leads us to observe the following:

- The method is very well behaved and routinely achieves convergence of both small and large LWR transport problems in 10 or fewer transport sweeps.

- Accumulation of the net-currents needed for the diffusion acceleration introduces about 25% overhead in the MOC transport sweep.

- The computational overhead for solving the two-level diffusion problem requires only about 15% of the computational effort required to perform the MOC sweeps.

- The overall speedups achieved in large LWR problems can be greater than a factor of 100 relative to unaccelerated MOC power iteration. (e.g., 2500 unaccelerated MOC sweeps were reduced to 10 MOC sweeps, with roughly 50% increased CPU requirements for diffusion parameter collection and iterative solution.)

The macro-track acceleration experience leads us to observe the following:

- The basic MOC transport sweep requires computation of macro-track parameters only for the macro-track angles. (Whereas, the diffusion acceleration method requires computation to compute net currents when sweeping each angle.)

- The method can be applied to any MOC tracking algorithm (in arbitrarily complex geometry) without need for the rectangular acceleration mesh used in the diffusion acceleration.

- More macro-track power iterations are required to achieve convergence than are required with diffusion power iteration.

- In large LWR problems, the computational effort to solve the macro-track iterations is about one half of that required for the MOC transport sweeps, (e.g., 100 macro-track iterations required for a factor of 120 fewer tracks than for the fine-mesh MOC sweeps, with no net current accumulation overhead).

- Approximation of the flat-source region fluxes and sources requires about twice the amount of computational effort as the macro-track iterations. (The fine-mesh construction is required at each macro-track iteration.)

- Approximation of macro-track equivalent sources also requires about twice the amount of computational effort as the macro-track iterations. (The fine-mesh construction is also required at each macro-track iteration.)

- In rare cases in which the diffusion acceleration method has difficultly converging (e.g., requiring more that 15 transport sweeps - as in the critical assembly problem of this paper), the macro-track acceleration is more effective in reducing the number of required MOC transport sweeps.

- Overall speedups of 20-50 relative to unaccelerated MOC power iteration are routinely achieved in large LWR problems.

The observation that the macro-track method requires more power iterations than the diffusion acceleration method leads one to the obvious conclusion that acceleration of the macro-track iteration is needed to further improve computational efficiency.

One natural way to achieve this would be to use a two-level acceleration scheme with the coarse level being the same non-linear assembly-mesh diffusion model used in the previously-discussed diffusion acceleration method. This scheme has yet to be tested.

8. CONCLUSIONS

Both acceleration methods presented here are effective at accelerating convergence of LWR applications of the method of characteristics. These methods permit acceleration meshes to be far coarser than the basic solution mesh, without restriction on geometrical complexity. Both accelerations methods achieve reductions in the number of MOC transport sweeps that are independent of the number of energy groups, angular quadrature, ray spacing, and flat-source resolution of the basic MOC LWR model.

The diffusion acceleration method is robust and reduces CPU times by more than a factor of 100 in large LWR problems. This method can be used to accelerate MOC methods that have arbitrarily complex geometries, provided a rectangular acceleration mesh is super-imposed on the physical mesh.

The macro-track acceleration method presented here is also robust and effective in achieving acceleration of MOC solutions. It has the advantage of requiring no regular geometrical acceleration mesh, and it further reduces the number of MOC transport sweeps required in difficult LWR problems.

Multi-cycle, full-core, 2-D PWR and BWR calculations, complete with fuel shuffling, rotation, and depletion, are now practical using the acceleration methods for the method of characteristics in CASMO-4E. Quarter-core problems (in eight energy groups, starting from flat flux) require about 20 CPU minutes on a 1.8 GHz Pentium 4 PC, and complete cycle depletions typically require a few CPU hours.

9. REFERENCES

1. M. Edenius, *et al.*, “CASMO-3 User’s Manual, Rev. 3,”STUDSVIK/NFA-89/3, Studsvik of America (1993).
2. D. Knott, *et al.*, “CASMO-4 Methodology Manual,” STUDSVIK/SOA-95/02, Studsvik of America (1995).
3. M. J. Halsall, “CACTUS, A Characteristics Solution of the Neutron Transport Equations in Complicated Geometries,” AEEW-R-1291, U.K. Atomic Energy Authority (1980).
4. G. S. Lee, *et al.*, “ Acceleration and Parallelization of the Method of Characteristics for Lattice and Whole-Core Heterogeneous Calculations,” *Proc. Int. Mtg. Advances in Reactor Physics and Mathematics and Computations into the new Millennium*, Pittsburgh, PA, 2000, p 2301, CD-rom, American Nuclear Society (2000).
5. M. L. Adams and E. W. Larsen, “Fast Iterative Methods for Discrete-Ordinates Particle Transport Calculations,” *Progress in Nuclear Energy,* Vol. 40. No. 1 pp3-159 (2002).
6. M. R. Zika and M. L. Adams, “Transport Synthetic Acceleration for Long-Characteristics Assembly-Level Transport Problems,” *Nucl. Sci. Eng*. 134, 135, (2000).
7. K. S. Smith, “Nodal Method Storage Reduction by Nonlinear Iteration,” *Trans. Am. Nucl. Soc.*, 44, 265 (1984).

8. C. Lee and T. J. Downar, "Nonlinear Formulation for Multigroup Simplified P3," *Trans. Am. Nucl. Soc.*, 83, 297 (2000).
9. J. M. Aragonés and C. Ahnert, "A Linear Discontinuous Finite-Difference Formulation for Synthetic Coarse-Mesh Few Group Diffusion Calculations," *Nucl. Sci. Eng*. 94, 309, (1986).
10. R. E. Alcouffe, *et al.,* "User's Manual for TWODANT: A Code Package for Two-Dimensional, Diffusion-Accelerated, Neutral-Particle Transport," Los Alamos National Laboratory Report LA-10049-M (1984).
11. N. Z. Cho, *et al.*, "Whole-Core Heterogeneous Transport Calculations and Their Comparisons with Diffusion Results," *Trans. Am. Nucl. Soc.*, 83, 292 (2000).
12. S. Kosaka, "Transport Theory Calculation of a Heterogeneous Multi-Assembly Problem by Characteristics Method with Direct Neutron Path Linking Technique," *Journal of Nuclear Science and Technology,* Vol. 37, No. 12 p. 1015 (2000).
13. S. Kosaka, Personal Communication, (2002).
14. K. S. Smith and J. D. Rhodes, "CASMO-4 Characteristics Method for Two-Dimensional PWR and BWR Core Calculations," *Trans. Am. Nucl. Soc.*, 83, 294 (2000).