

CONF-960690--4

RECENT ADVANCES IN NEUTRAL PARTICLE TRANSPORT
METHODS AND CODES

Y. Y. Azmy

Oak Ridge National Laboratory*
P.O. Box 2008, MS-6363
Oak Ridge, Tennessee 37831-6363
email: yya@ornl.gov

RECEIVED

JUN 03 1996

OSTI

Full paper for publication in the *Proceedings of the 5th International Conference on Applications of Nuclear Techniques*, Crete, Greece, June 9-15, 1996.

"The submitted manuscript has been authored by a contractor of the U.S. Government under contract DE-AC05-84OR21400. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes."

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

*Research sponsored by the Oak Ridge National Laboratory managed by Lockheed Martin Energy Research Corporation for the U.S. Department of Energy under contract No. DE-AC05-96OR22464.

MASTER
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED *JIC*

Yousry Y. Azmy

Oak Ridge National Laboratory
Oak Ridge, Tennessee 37831
yya@ornl.gov

ABSTRACT

An overview of Oak Ridge National Laboratory's (ORNL) three-dimensional neutral particle transport code, TORT, is presented. Special features of the code that make it invaluable for large applications are summarized for the prospective user. Advanced capabilities currently under development and installation in the production release of TORT are discussed in some detail. These include: multitasking on Cray platforms running the UNICOS operating system; Adjacent-cell Preconditioning (AP) acceleration scheme; and graphics codes for displaying computed quantities such as the flux. Further developments for TORT and its companion codes to enhance its present capabilities, as well as expand its range of applications will be discussed. Speculation on the next generation of neutral particle transport codes at ORNL, especially regarding unstructured grids and high order spatial approximations, will also be mentioned.

Keywords: TORT, DORT, DOS code system, neutral particle transport, multitasking, UNICOS Cray, Adjacent-cell Preconditioning, acceleration, graphics

1. OVERVIEW OF TORT AND PERIPHERAL CODES

Oak Ridge National Laboratory's three-dimensional discrete ordinates neutral particle transport code TORT¹ is recognized worldwide as a leader in the field. The code's versatility, robustness, and computational efficiency allow users unprecedented production level computations in a wide variety of applications. In addition, TORT's reliability and a high level of confidence in its computed results have been established over the years by a multitude of independent research groups working on different applications. At present the largest TORT application employs 3.6 million computational cells, S_{16} angular quadrature, P_3 anisotropic scattering, and 11 energy groups, a monumental challenge to the computational resources available to the vast majority of researchers around the world. This challenge is met by a highly judicious management of the memory and disk space while keeping the execution time well within bounds.

TORT was originally conceived more than ten years ago as a three dimensional extension of its two dimensional predecessor DORT.² The primary application of TORT was the analysis of radiation fields penetrating into large concrete buildings as described in Ref. 3. Since then the applications base of TORT has expanded dramatically beyond the traditional applications in nuclear reactor analysis⁴⁻⁶ and shielding,^{7,8} cross section preparation,⁹ and pressure vessel computations,¹⁰ to include novel medical applications,¹¹ and even charged particle transport.¹²

In the remainder of this Section we will provide a broad overview of prominent features and capabilities of TORT and its peripheral codes. In subsequent sections we will focus on three topics of interest to applications users and along which most progress has been achieved recently. In Sec. 2 we describe the multitasking capability on Cray supercomputers that exhibited Wall Clock speedup approaching a factor of 5 on a time-shared lightly-loaded 8-CPU Y/MP, and moderately large problems, i.e. S_{16} , and over 100,000 cells. In Sec. 3 we briefly present the formalism for the Adjacent-cell Preconditioning (AP) acceleration scheme that is extremely effective in reducing the number of iterations required to achieve tight convergence particularly in optically thick diffusive regimes. Section 4 is dedicated to an exposition of graphics codes that enable a multitude of postprocessing capabilities of final and intermediate data, comprising a rudimentary stage in the development of a comprehensive state-of-the-art graphical user interface for the entire Oak Ridge discrete ordinates codes system. Finally, in Sec. 5 we present our vision for the future of this generation of codes, as well as the next generation presently in the preliminary planning stage.

1.1 Adjustable problem parameters

Perhaps the most daunting factor contributing to the difficulty of transport calculations is the large size of the discrete variable system of equations that must be solved numerically. This is a direct consequence of the high dimensionality of phase space: $2d$ in d -dimensional geometry, d variables representing physical space, $d-1$ representing the

particles direction of motion (discrete ordinates), and one energy variable. [Note we have assumed steady state]. It is critical, therefore, to conserve discrete variables as much as possible without jeopardizing the accuracy and reliability of the solution. This is accomplished in TORT by permitting the user to adjust the level of detail in a variety of problem parameters according to the anticipated rate of change in the solution.

For example, sharp flux and cross section anisotropies are more notable at the high energy end of the spectrum. Hence TORT permits the user to select a high order angular quadrature set and high order P_l expansion of the cross sections in the high energy groups, and lower orders in the low energy groups thereby reducing the problem size without sacrificing accuracy. Another important example which has proven extremely valuable in solving very large problems with substantial geometric detail is the discontinuous mesh option in TORT. This feature allows the user to employ more computational cells where fine geometric structure exists, and fewer cells elsewhere. This is achieved by allowing the user, within some constraints, to adjust the number and boundaries of individual cells within a row, not necessarily coinciding with those in adjacent rows.

Other adjustments to problem parameters not related to phase space size include, for example, setting iteration number by group, etc.

1.2 Memory management

Experience over the lifetime of electronic computers has demonstrated time and again that there is no such thing as *sufficiently large memory!* Problems, in most applications, tend to grow in size with available memory size constantly challenging computer codes and their developers to meet a seemingly insatiable demand. In recognition of this fact, and also to accommodate a variety of platforms with a disparate range of memory sizes TORT attempts to meet a *memory objective* set by the user at run time. First TORT tries to fit the entire problem in the memory objective; if this is not possible it attempts to fit one group at a time within the memory objective and I/O flux and cross section data to scratch files. If this too fails, the code breaks up the geometric configuration into blocks of planes, each of which can fit in the memory objective, and I/O to scratch files is used to maintain and update the data during the calculation. If a single plane does not fit in the memory objective TORT tries to increase the memory objective, assuming it is less than all that is available on the machine, and if this fails it informs the user of its attempts then terminates execution unsuccessfully. Clearly this sequence of attempts is designed to minimize the adverse effect of I/O on performance while enabling the solution of ever larger problems.

1.3 Spatial Approximation

TORT is based on the discrete ordinates approximation of the independent angular variable, and the multigroup discretization of the energy variable. A variety of approximations of the spatial dependence are available in the code to accommodate a broader spectrum of applications. These are

- (a) The θ -weighted method, which is essentially a whole set of methods parametrized by the single parameter θ set by the user at run time in the range $[0,1)$. This is a weighted diamond difference method spanning the range from the diamond difference scheme (optically thin cells) to the Step method (optically thick cells) where the weights are computed to ensure positivity of the outgoing angular flux given positive incoming flux and volumetric source. Due to the relative simplicity of its equations and its low order approximation, the θ -weighted method is the least computationally intensive option among TORT's spatial approximation methods.
- (b) The Linear Nodal (LN) method was installed in TORT to enable using optically thick cells while retaining high accuracy of the computed angular flux. This is achieved by computing the first spatial moment of the angular and scalar flux, in addition to the average quantities computed in the θ -weighted method.^{13,14} Many modifications to the original method have been implemented over the years in order to improve method accuracy and efficiency, and solution positivity to the extent that, at present, LN is the recommended method for most large applications with optically thick regions.
- (c) The Linear Characteristic (LC) method¹⁴ also computes the first spatial moment of the flux on each cell's surfaces and volume using the exact characteristic paths from incoming to outgoing surfaces, then projecting the resulting expression onto the basis functions (constant and linear). As far as accuracy is concerned LC is competitive with LN with each method gaining an edge over the other for some, but not all, problems. LC executes faster than LN on scalar machines, but due to the high level of vectorization of the LN it is about four times faster on Cray computers.¹⁵

On a given mesh LN and LC run longer, require more memory, and consume larger disk space but provide more accurate solutions. However, for a fixed accuracy the θ -weighted method typically requires eight times as many computational cells than either of the linear schemes. Ultimately for the same solution accuracy requirement, the linear methods end up utilizing less computational resources, i.e. CPU time, memory, and disk space, than the θ -weighted method.¹⁵

1.4 DORT: the two-dimensional companion code

While TORT offers a few two-dimensional geometry options, DORT offers a more comprehensive solution for two dimensional problems with more capabilities than is available in TORT. For example DORT includes several cylindrical geometry options that have been tested thoroughly over the years by many users. However, recent developments implemented in TORT, such as LN and LC are not available in DORT. DORT problems are typically smaller in size than TORT problems; thus DORT appeals to more applications users who may be limited by the available computational resources or turn-around time. For example scoping calculations that involve a large number of runs are more likely to employ a DORT model.

1.5 Splicing with TORSED and TORSET

TORT is distributed within the Discrete Ordinates System (DOS) package which includes DORT and several other peripheral codes that are used to pre- and post-process computed data. One important data processing function that is commonly used to solve large problems is *splicing* wherein the problem domain is split in two or more subdomains that are loosely coupled in only one way from the source, *primary*, to the observable, *secondary*, subdomain with no feedback. In order to avoid running the large problem in its entirety a detailed primary computation is performed and a source is computed for the defined secondary subdomain that includes points at which the solution is of interest. Additional geometric detail can be introduced, or new material compositions can be substituted in the secondary configuration which is then solved. The primary geometry can be a DORT *rz* case or a TORT Cartesian case; the secondary geometry is TORT Cartesian case whose axes do not necessarily coincide with the primary axes. Coupling between the primary and secondary calculations is performed by the TORSED^{16,17} (DORT \rightarrow TORT) and TORSET¹⁷ (TORT \rightarrow TORT) codes which extract the boundary angular flux from an interface file, perform angular or spatial interpolations if necessary including coordinate translation and rotation, then write out the resulting source in the proper format for TORT to read.

2. MULTITASKING ON CRAY COMPUTERS

Two multitasking options are available in TORT for execution on multiprocessor Cray computers running the UNICOS operating system: the Octant Parallel (OP), and the Direction Parallel (DP) methods.¹⁸ Both algorithms are based on coarse-grained angular domain decompositions which typically produce good parallel efficiency due to the relatively large computation load to parallelization overhead ratio. It has been observed previously that angular domain decomposition in Cartesian geometry is intrinsic, implying a one-to-one correspondence (within arithmetic precision) between the sequential and parallel intermediate and final results, so that the number of iterations required to achieve convergence is independent of the number of concurrent processes. These considerations combined with the long execution times for large applications motivated equipping TORT with multitasking capability.

The mesh sweeps are performed in TORT one plane at a time; for each plane every row in the *x*-direction is swept to the left then to the right in all angular directions within an octant in angular space. Hence, the OP is implemented by starting a slave task to perform the left sweep for all angles with $\mu < 0$ while the master task completes the right sweep for all angles with $\mu > 0$ then awaits the slave task before proceeding to the next row. Mutually exclusive locks are implemented within the parallel section of the code to avoid memory conflicts in the process of accumulating the angular flux contributions to the new iterate of the scalar flux and its higher angular and spatial moments. After sweeping a row along discrete ordinates within two octants, TORT proceeds to sweep the next row over the same two octants. After both octants are solved in all rows of a plane, TORT solves the other two octants which have the same η -level sign. Upon completing a plane, TORT steps to the next plane; a complete inner iteration consists of a downward and upward sweep through the computational mesh.

The OP strategy suffers two drawbacks. First, only two tasks, corresponding to the right and left sweep directions, can execute concurrently, thus bounding the potential speedup from above by 2. Second, only fixed source (including vacuum) and periodic boundary conditions are permissible along the spatial dimension in which the sweep


concurrency is realized, the x -dimension in TORT. OP is selected by the user at run time by setting the $ncpu$ parameter in the 62\$\$ input array to -1, or ≤ -2 , to execute with the master, or one master plus one slave, tasks, respectively. [Setting $ncpu$ to 0 selects the original sequential/vectorized algorithm].

The DP algorithm is based on sweeping each row of cells in the mesh concurrently along the discrete ordinates in an octant of angular space. As such it eliminates the two drawbacks of OP mentioned above thus providing a large number of independent tasks to be executed simultaneously and, since octants are processed successively, allowing all types of boundary conditions. Its main drawback is the typically limited number of tasks available, essentially the number of discrete ordinates per octant, compared to the very large number of computational cells in a spatial domain decomposition. Of course the number of concurrent tasks sets an upper bound on the potential for speedup, except in case of distributed memory architectures where the aggregate memory of the participating processors can reduce or eliminate the need for I/O. Typically for coarse-grained platforms supporting at most a few tens of supercomputer class CPUs such as the Cray Y/MP, the real bound on speedup is set by the number of processors.

DP is selected by the user at run time by setting the $ncpu$ parameter to a positive value implying execution on one master and $ncpu - 1$ slave tasks. The present implementation of DP creates as many slave tasks as specified by the user every time subroutine *row* is called, then the angles within the right and left octants in angular space are dynamically scheduled to the available tasks. This feature of DP reduces the potential for load imbalance across tasks which is highly possible in the time sharing environment typical of Cray computers.


Several test problems were used to measure and characterize the parallel performance of the multitasking options in TORT. Among these is sample problem 6 (TP6) distributed with the code which we executed on an 8-CPU Y/MP computer running UNICOS 8.0. TP6 is a two group, S_6 , 104,247 cell problem that uses LN spatial approximation and converges both groups in 17 iterations. In order to examine the scaling of parallel performance with angular quadrature we also solved this problem using S_{16} . The higher quadrature order benefits parallel performance in two ways. First, it provides a larger pool of independent processes which for DP improves the potential for speedup on a platform supporting more CPUs and enhances the load balance across tasks. Second, since in OP and DP implementation the tasks are created only once per row, then the higher quadrature results in a larger computational load per task thus diluting the effect of the task starting penalty. For each run we compared the iteration history and the converged solution for each $ncpu$ case with the sequential, i.e. $ncpu=0$, output to verify the parallel implementation. We observed differences between the two outputs due to roundoff errors that manifest themselves primarily in very small values of the solution, of the order 10^{-20} or smaller. The measured Wall Clock speedup for these cases is plotted in Fig. 1 as a function of the input parameter $ncpu$ with values ranging from -2 to 8, the number of processors on the target platform.

Figure 1. Measured Wall Clock Speedup for TP6 on Typically Loaded Machine.



The nondedicated environment of these runs implies dependence of the parallel performance on machine loading at the time of execution, hence the non-monotonic speedup with increasing n_{cpu} . The results depicted in Fig. 1 were observed in single runs for each case shown (except when unusually high machine loading occurred) so they are representative of performance on a typically loaded machine. In order to characterize the parallelization itself we repeated the run for each value of n_{cpu} , until a sufficiently large and monotonically increasing machine utilization occurred. The measured Wall Clock speedup plotted in Fig. 2 is representative of machine performance on a lightly loaded machine. More importantly, it comes closer to characterizing performance in a dedicated environment, hence judging the success of the parallel algorithms employed in ideal circumstances.

Figure 2. Measured Wall Clock Speedup for TP6 on Lightly Loaded Machine.



Figures 1 and 2 exhibit good speedup of the respective computations which appears to increase with the order of angular quadrature and CPU availability. These results demonstrate the beneficial effect of large problem size on parallel speedup in general and on lightly loaded machines in particular where the Wall Clock speedup approaches a factor of 5.

Further insight into the parallel performance has been gained by constructing and verifying a CPU-performance model. This was motivated by the unexpected dependence of the CPU time reported by the system on machine loading. Experiments involving setting the CPU-hold time to zero indicate that the reason for this dependence, and a large fraction of the parallelization penalty results from the operating system deciding to keep a CPU on hold for the application until a task becomes ready to run. Of course the lower the machine load the longer the system can afford to keep a CPU on hold, hence the observed dependence. The performance model also showed that certain redundancies in the angular loop and task creation overhead results in penalties that severely burden the speedup factor. Research is presently underway to alleviate the last two sources of inefficiency.

3. ADJACENT-CELL PRECONDITIONER ACCELERATION

The recommended method for accelerating the iterative convergence of the inner iterations in TORT is the Partial Current Rebalance (PCR) method. This method is based on reinforcing the balance of neutrons over each computational cell using the cell-surface partial currents resulting from the latest mesh sweep. The discrete variable equation resulting from PCR has the same cell-coupling stencil as a discretized cell-centered diffusion equation but does not necessarily possess some of its important features like diagonal dominance, etc. The PCR matrix equation is solved via a Successive Over-Relaxation (SOR) scheme with the relaxation factor computed numerically from the SOR iterates.


For most TORT applications PCR has proven robust and efficient. However, recent advances in the analysis of the spectral properties of iterative procedures for solving the transport equation have provided the basis for powerful

acceleration operators that, at least theoretically, exceed the performance of PCR. The most notable example of such new methods is Diffusion Synthetic Acceleration (DSA) whose spectral radius is bounded from above by 0.25 for model problem configurations, i.e. homogeneous material composition and uniform mesh.¹⁹ The most serious limitation of the class of unconditionally stable DSA operators as far as large three dimensional applications is concerned is that they are edge-centered. Since there are many more surfaces than computational cells the DSA matrix equation can be prohibitively large; this difficulty is compounded further in high order methods, i.e. LN and LC, if the first spatial moments of the flux are to be accelerated also.

This motivated us to pioneer a new class of cell-centered acceleration schemes based on preconditioning the transport operator.²⁰ In this approach we introduce a preconditioning step after the mesh sweep stage, then we consider a specific coupling stencil, the adjacent cell stencil which is identical to a cell-centered diffusion coupling scheme. Upon performing a spectral analysis using a model problem configuration and computing the eigenvalue of the AP iterations we show that a diffusion-like relation between the diagonal and off-diagonal elements of the preconditioner must be satisfied to guarantee stability of the flat mode. Another relationship between the preconditioner elements is obtained by setting the limit of the eigenvalue at the origin in Fourier space to zero, thus implying immediate convergence of the flat mode. From these two relationships we determine the diagonal and off-diagonal elements of the AP operator. For practical, non-model problem configurations these are augmented by Larsen's treatment of the boundary conditions,¹⁹ and reciprocal-averaging for material and mesh discontinuities. Extensive testing of the spectrum of the AP demonstrated its unconditional stability and high efficiency in reducing the spectral radius of the iterative process even for cells with very small aspect ratios. Furthermore, the spectral radius for model problems actually vanishes in the limit of infinite cell optical thickness.

The spectral properties of the AP scheme were thoroughly tested using Burre's Suite of Test Problems (BSTeP) whose configuration is depicted in Fig. 3. BSTeP is comprised of 144 cases that cover a wide range in parameter space including configurations with sharp mesh and material, namely total cross section, discontinuities.

Fig. 3. Geometric Configuration of the Burre's Suite of Test Problems (BSTeP).



We solved all cases in BSTeP using AP and PCR as implemented in TORT with the Step spatial approximation, and using TWODANT's²¹ DSA with the adaptive diamond difference scheme. For each case we observed the number of iterations required to achieve 10^{-4} relative pointwise convergence by each method as shown in Table 1. Examining the AP performance in Table 1 we note significant deterioration in the spectral radius in cases with sharp material discontinuity. [This deterioration is worst in PCR and DSA, as well as other methods not shown here]. In particular, the

Table 1. Number of Iterations Required to Achieve Convergence for Burre's Suite of Test Problems Using the Symmetric AP (top), TORT's Rebalance (middle), for the Step Method in a Special Version of TORT, and DSA (bottom) for Adaptive DD in TWODANT 3.0, on IBM RS6000.

			σ_1											
			.01			.1			1			10		
			dy ₁	.01	.1	1	.01	.1	1	.01	.1	1	.01	.1
σ_2	.01	dy ₂	7	5	5	8	7	7	10	13	9	24	32	13
		.01	4	5	5	6	7	10	9	13	18	19	58	55
			15†	6†	5	16†	9†	6†	15†	16†	13	48†	259†	500
	.1		5	5	5	8	7	8	12	13	15	15	28	21
			5	5	5	5	6	10	11	12	17	17	26	149
			6†	5†	5	8†	8	7	15†	15†	21	168†	474†	500
	1		5	5	5	8	8	8	12	12	12	13	14	12
			5	5	5	8	6	10	13	12	13	106	59	108
			5	5	5	7	7†	8	12†	14	15	246†	255	42
	.1	.01	8	6	5	8	7	7	8	7	7	14	12	6
			5	5	5	6	7	10	9	13	18	18	53	53
			18†	6†	5	14†	8†	6	13†	10†	6	60†	241†	233
.1	.1	7	6	5	7	7	7	9	8	8	10	14	9	
		5	6	5	7	7	10	13	11	28	15	55	143	
		8†	8†	5	8†	7†	6	7†	10	8	39†	165†	43	
1		7	7	7	6	6	7	10	10	9	9	11	9	
		8	8	6	8	8	11	13	15	12	76	59	72	
		6†	6†	7	6	6†	7	8†	10	10	443	125	404†	
1	.01	13	8	5	9	7	7	7	8	7	10	6	5	
		6	5	5	6	6	10	8	12	17	17	35	49	
		16†	9†	5	13†	9†	6†	10†	9†	6	39†	231†	308	
.1	.1	11	14	10	8	9	8	8	8	7	9	7	6	
		9	11	8	8	11	8	15	9	13	32	37	114	
		8†	8†	7†	11†	14†	9	9	10†	6	22	335†	204	
1		9	10	13	8	8	8	8	8	8	8	7	6	
		11	10	9	9	10	13	9	8	10	72	77	59	
		7†	29†	14	18	11†	9†	8	8†	6	217†	500	500	
10	.01	32	17	11	17	12	9	7	8	7	10	8	6	
		12	10	9	12	10	7	11	17	16	13	25	43	
		18†	19†	11	25†	18†	9†	20†	15†	8	54†	148†	274	
.1	.1	24	19	13	9	11	9	8	9	8	8	9	7	
		18	13	12	23	13	9	14	15	21	16	16	18	
		60†	67†	23	50†	43†	18	55†	40†	18	162†	265†	500	
1		12	12	8	8	8	8	8	9	8	8	8	5	
		18	20	18	23	16	13	12	21	22	47	17	19	
		71†	67†	60	71†	65†	54†	106†	89†	40†	500	500	500	

DSA: 500 ⇒ not converged

† Faster convergence obtained with *w damp*=2.0

upper right corner cases of Table 1, i.e. $\sigma_1 = 10$, $\sigma_2 = .01$, converge in far more AP iterations than can possibly be explained by the spectral analysis. Spectral analysis of a problem configuration with a single periodic interface across which mesh and material discontinuities are permitted provides evidence that the diffusion coupling stencil imposed by AP, and shared by many other acceleration schemes, is responsible for this poor behavior. This should not be too surprising because the diffusion approximation is known to be invalid near sharp material discontinuities. Alternative coupling schemes capable of faster convergence, and comparison of AP performance to that of other acceleration methods is under investigation.

In order to summarize the comparison presented in Table 1 across parameter space we devise the following scoring scheme: For each case in the table we assign one point to the fastest converging method; in case of a two- (three-)way tie each involved method is assigned 1/2 (1/3) point, respectively; otherwise a method is assigned a 0. Summing each method's score across Table 1, AP comes first with 79 2/3 points (55%), followed by PCR with 38 2/3 (27%), then TWODANT's DSA with 25 2/3 (18%). Summing all iterations (bounded from above by 100 per case) for each method across the entire table, AP consumes the least iterations, 1,364, followed by PCR, 2,877, then TWODANT's DSA, 5,044. Focusing on the hard cases with $\min(\sigma_1, \sigma_2) \geq 1$, AP's advantage gets bigger, with 71 2/3 points (66%), to PCR's 22 2/3 points (21%), and DSA's 13 2/3 points (13%). The total number of iterations (bounded from above by 100 per case) for these is 1,128, 2,638, and 4,774, respectively.

4. GRAPHICS ROUTINES

Earlier development of graphics capabilities for DORT utilized a commercial library to provide elementary constructs. More recently ORNL has opted for public-domain-based development of our neutronic codes graphics capabilities in order to facilitate distribution of the package in a self-contained form for the users' convenience. For this purpose we adopted Sandia National Laboratory's RSCORS graphics library which is distributed with the latest release of our codes.

The user of TORT can generate plots of flux or activities from formatted or unformatted files generated by the code using the XTORID and ISPL3D codes.²² Two dimensional plots are generated for selected planes in the TORT geometry, and can be viewed on-screen for a variety of platforms, or in hardcopy. Several options are available for displaying the plotted quantity: color and greyscale shading, symbol and line contours. In addition, the geometric configuration of the plane can be overlaid on the plot as shown in the sample plot in Fig. 4.


5. FUTURE OF TORT AND NEXT GENERATION OF TRANSPORT CODES

In the past decade since TORT was conceived many developments have occurred on the theoretical, software, and hardware fronts that make a complete overhaul of the neutral particle transport codes system rather tempting. Hence our vision of the future of our codes is based on two time scales.

In the short term, 5-10 years, TORT will continue to be developed and maintained to meet the immediate needs of scientists and engineers worldwide. Foreseeable developments include construction of a Graphical User Interface and advanced visualization tools to simplify input preparation and debugging, output interpretation and manipulation, code interfacing, etc. This activity will build on the presently available graphics codes described above. Enhancement of iterative acceleration schemes will continue especially for problems with sharp material discontinuities, and advanced methods for accelerating outer iterations will be considered. Multiprocessing capability for message passing architectures that is based on the evolving MPI standard also will be provided. Several of the peripheral codes, e.g. TORSET, and methods, e.g. splicing, will be updated and upgraded, and interface files coupling these codes will be standardized.

For the long term, plans for the next generation of neutron transport codes are already underway. Specific features of the new code system include: Arbitrarily High Order Transport (AHOT) formulation potentially with adaptive order to preserve positivity or meet a prespecified accuracy objective; unstructured grids to reduce the adverse effects of geometric approximations on solution accuracy and computational resources; development of unconventional iteration and acceleration strategies better suited to take full advantage of multiprocessor architectures; advanced CAD-like GUI to aid the user in modeling complex geometric configurations and viewing the results.

Fig. 4. Neutron Flux with Energy > 0.1 MeV for the Three-Element Core Design of the Advanced Neutron Source.²²



6. REFERENCES

1. W. A. Rhoades, "The TORT three-dimensional discrete ordinates neutron/photon transport code," *ORNL/TM-13221*, to be published, 1996.
2. W. A. Rhoades and R. L. Childs, "An updated version of the DOT 4 one- and two-dimensional neutron/photon transport code," *ORNL-5851*, 1982.
3. W. A. Rhoades, R. L. Childs, M. B. Emmett, and S. N. Cramer, "Application of the three-dimensional Oak Ridge transport code," *Proc. Am. Nucl. Soc. Topical Mtg. on Reactor Physics and Shielding*, pp. 225-238, Chicago, IL, September 17-19, 1984.
4. A. J. J. Bos, J. E. Hoogenboom, "Calculations and measurements of neutron fields in a facility for big sample neutron activation analysis at the research reactor at Delft," *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*, pp. 584-590, No. Falmouth, MA, April 21-25, 1996.

5. E. B. Brodtkin, A. L. Egorov, V. I. Vikhrov, S. M. Zaritsky, "The determination of the VVER-1000 surveillance neutron fluence using the ^{54}Mn activity measurements and TORT neutron spectra calculations," *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*.
6. Akira Shono, Hirokazu Tsunoda, Morio Takemura, and Hiroyuki Handa, "JASPER experiments and analyses of IHX sodium activation and gap streaming mockups," *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*, pp. 872-879, No. Falmouth, MA, April 21-25, 1996.
7. I. K. Abu-Shumays and C. E. Yehnert, "Angular interpolations and splice options for three-dimensional transport computations," *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*.
8. Arnold H. Fero, Robert C. Sterritt, and James P. Tarzia, "Analysis and measurements of neutron streaming shields at Seabrook," *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*, pp. 329-336, No. Falmouth, MA, April 21-25, 1996.
9. I. Kodeli, L. Giancarli, "Update of the SUSD cross section sensitivity and uncertainty package and extension to 3D analysis," *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*, pp. 1104-1108, No. Falmouth, MA, April 21-25, 1996.
10. Sergey I. Belousov, Krassimira D. Ilieva, "ASYNT - Adjoint synthesis method for neutron flux evaluation onto VVER/PWR pressure vessel", *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*, pp. 11-18, No. Falmouth, MA, April 21-25, 1996.
11. D. W. Nigg, *et. al.*, "Demonstration of three-dimensional deterministic radiation transport theory dose distribution analysis for boron neutron capture therapy," *Medical Physics*, **18**, pp. 43-53, 1991.
12. Clifton R. Drumm, "Multidimensional electron-photon transport with standard discrete ordinates codes," *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*, pp. 398-405, No. Falmouth, MA, April 21-25, 1996.
13. W. F. Walters, "Augmented weighted-diamond form of the linear-nodal scheme for Cartesian coordinate systems," *Nuclear Science and Engineering*, **92**, pp. 192-196, 1986.
14. R. L. Childs and W. A. Rhoades, "Theoretical basis of the linear nodal and linear characteristic methods in the TORT computer code," *ORNL/TM-12246*, 1993.
15. R. L. Childs, "Numerical methods for the flux solution," *Workshop for the DORT and TORT Radiation Transport Codes*, No. Falmouth, MA, April 21, 1996.
16. W. A. Rhoades, "The TORSED method for construction of TORT boundary sources from external DORT flux files," *ORNL/TM-12359*, 1993.
17. W. A. Rhoades, "The TORSED and TORSET codes for coupling three-dimensional TORT calculations," *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*, pp. 551-557, Arlington, TX, April 24-28, 1994.
18. Y. Y. Azmy, D. A. Barnett, and C. A. Burre, "Multitasking the three-dimensional transport code TORT on Cray platforms," *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*, pp. 613-621, No. Falmouth, MA, April 21-25, 1996.
19. E. W. Larsen, "Unconditionally stable diffusion-synthetic acceleration methods for the slab geometry discrete ordinates equations. Part I: Theory," *Nuclear Science and Engineering*, **82**, pp. 47, 1982.
20. Y. Y. Azmy, "Adjacent-cell preconditioners for accelerating multidimensional neutron transport methods," *Proc. Am. Nucl. Soc. Topical Mtg. on Radiation Protection and Shielding*, pp. 390-397, No. Falmouth, MA, April 21-25, 1996.
21. R. E. Alcouffe, R. S. Baker, F. W. Brinkley, D. R. Marr, R. D. O'Dell, and W. F. Walters, "DANTSYS: A diffusion accelerated neutral particle transport code system," *LANL-12969*, 1995.
22. C. O. Slater, "The XTORID and ISPL3D codes for plotting TORT activities," *ORNL Internal Memo*, March 25, 1996.