M&C 2017 - International Conference on Mathematics & Computational Methods Applied to Nuclear Science & Engineering, Jeju, Korea, April 16-20, 2017, on USB (2017)

Recent Developments to MONK for Criticality Safety and Burn-up Credit Applications

Simon Richards, Max Shepherd, Adam Bird, David Long, Christophe Murphy and Tim Fry

ANSWERS Software Service, Amec Foster Wheeler, Kings Point House, Queen Mother Square, Poundbury, Dorchester, Dorset, DT1 3BW, United Kingdom.
simon.richards@amecfw.com

Abstract - *MONK®is a Monte Carlo software package for nuclear criticality and reactor physics analyses. The code is actively developed, maintained and supported by Amec Foster Wheeler's ANSWERS® Software Service in line with the ANSWERS vision of providing easy-to-use software that meets the current and emerging needs of the user community. This paper summarizes the current status of MONK and the recent developments which have been carried out to the code and its supporting nuclear data libraries and visualization package since the release of MONK 10A in 2014.*

I. INTRODUCTION

MONK®is an advanced Monte Carlo neutronics code for the solution of criticality safety and reactor physics problems, forming part of the ANSWERS® codes suite [1]. ANSWERS codes are widely used in over thirty countries around the world and on a range of reactor types including: AGR, BWR, CANDU, MAGNOX, RBMK, PBMR, PWR, VVER and many experimental reactors. The codes have already been applied to some of the future reactor technologies (e.g. high temperature and fast breeder reactors in the Generation IV programme) that are being developed not only for electricity generation but additionally for other applications.

MONK has a proven track record of application to the whole of the nuclear fuel cycle and is well established in the UK criticality community as the *de facto* standard Monte Carlo criticality model. It has been in continuous development and use since the 1960s and is developed, distributed and actively supported in use by Amec Foster Wheeler as part of its ANSWERS Software Service.

The current version of MONK is MONK 10A which was released in 2014 after an extensive programme of enhancements over the previous version [2].

In this paper we discuss the further significant developments which have been carried out since that release, many of which will be available to users in a forthcoming MONK 10B release. We also describe the upgrades to the geometry visualization and results display tool Visual Workshop to support the new features of MONK 10B.

The Monte Carlo code MCBEND [3] is a sister code to MONK for general radiation transport applications. MCBEND shares a significant amount of code with MONK including material specification, geometry, nuclear data, neutron collision processing, and some source specfication and scoring functions. Therefore the developments described here for MONK will also be available in the next version of MCBEND where they are based on shared code (particularly the geometry, nuclear data and Visual Workshop improvements). Furthermore some of the geometry modules are also shared with the RANKERN [4] point-kernel program for gamma transport solutions and the CACTUS 3D method of characteristics module in the WIMS [5] modular scheme for neutronics calculations, so developments to the geometry modules in MONK can also be exploited in those codes.

Some of the key developments described in the following sections include:

- a new module to tally the fission matrix and calculate its eigenvalues and eigenvectors, and the dominance ratio;
- improvements to the CAD geometry import function to improve run-time performance and give greater control over material contents;
- the ability to create copies of geometry components with reflections as well as rotations;
- a new elliptical cross-section cylindrical body;
- new nuclear data libraries based on the JEFF3.2 and ENDF/B-VII.1 evaluations;
- improvements to the scattering kernel for scattering near epithermal resonances;
- improvements to the burn-up credit capabilities, including improved robustness and user image and greater control of isotopic content to facilitate actinide-only and actinide plus fission product burn-up credit calculations;
- a new module to score shielded cross-sections in the WIMS 172 group scheme;
- improvements to the parallel capabilities of MONK including a new option to split looping calculations across multiple processors; and
- miscellaneous other improvements.

Additionally some developments to the ANSWERS Visual Workshop tool will be described, including:

- the incorporation of the SPRUCE tool for uncertainty analysis;
- the development of a suite of optimization tools; and
- support for the new features in MONK 10B.

II. FISSION MATRIX

A discretized form of the eigenvalue problem solved by MONK may be written [6]

$$S_i = \frac{1}{k} \sum_{j=1}^{N} F_{i,j} S_j \,, \tag{1}$$

where F is the discrete fission matrix, k and S are the fundamental eigenpair and N is the number of spatial regions. Each element $F_{i,j}$ of the fission matrix represents the number of fission neutrons born in spatial region i resulting from an average neutron born in spatial region j. The spatial discretization must consist of disjoint regions and must contain all of the fissile material in the problem.

A new module has been introduced in MONK 10B which tallies the fission matrix during the Monte Carlo tracking process. Since the size of the fission matrix is $N \times N$ the memory requirement for storing the full fission matrix is prohibitive for practical problems. However, advantage can be taken of the fact that the fission matrix tends to be sparse, albeit with a sparsity pattern which is not known *a priori*.

A number of approaches to tallying the sparse fission matrix have been tried. One method uses geometric intuition [7] to effectively impose an *a priori* sparsity pattern, taking advantage of the knowledge that most fissions are caused by fission neutrons which were born within some local region width. However this approach suffers from the disadvantages of having to apply engineering judgement to decide on the local region width and the need to deal with non-local fissions by tallying them in the nearest region. The approximations which result from these limitations render this approach unsuitable for complex geometries.

These approximations may be removed by using a compressed-row storage scheme such as the one described in [6]. This scheme was initially implemented successfully in MONK, but that implementation was found to be inefficient for calculations with a large tally mesh and a large number of superhistories[1] per stage. An alternative approach based on linked lists was implemented and this has been found to be efficient in both time and memory.

MONK therefore uses a sparse storage scheme which is capable of efficiently inserting elements into the fission matrix during the Monte Carlo tracking process. The tallying of the fission matrix in MONK differs from that in some other Monte Carlo codes in that it does not use the fission bank and includes all of the fission sites arising in each superhistory.

After the fission matrix has been tallied the eigenvalues and eigenvectors for a user-requested number of eigenmodes can be computed. MONK uses the ARPACK subroutines [9] based on the Implicitly Restarted Arnoldi Method (IRAM) to compute the eigenvalues and corresponding eigenvectors. Since the elements of the fission matrix are non-negative the Perron-Frobenius theorem asserts that it has a unique largest real eigenvalue and that the corresponding eigenvector can be chosen to be strictly positive. However there is no *a priori* reason why the IRAM method should return an eigenvector having all positive components as opposed to one have all negative components. The matrix of eigenvectors is therefore normalized such that the fundamental eigenvector is non-negative.

The eigenvalues are ordered by monotonically-decreasing magnitude and, if at least the first two eigenmodes are requested, an estimate of the dominance ratio (the ratio of the second largest to the largest eigenvalue) is calculated. The fundamental eigenvector of the fission matrix represents the stationary fission source distribution, and the convergence rate of the power iteration method is give by the dominance ratio. Therefore the fission source converges slowly in problems having a dominance ratio which is close to unity.

Experiment No. 89 in the MONK Validation Database [10] has been used here to demonstrate the fission matrix capability in MONK 10B. This experiment is LEU-COMP-THERM-048 in the ICSBEP Handbook [11], and consists of a circular array of 3 wt% uranium dioxide fuel rods on a square lattice with light water moderation and reflection. The fission matrix was tallied in a $54 \times 54 \times 1$ mesh over the array, giving one region per pincell. The MONK geometry for this case is shown in Fig. 1 and the first nine eigenvectors of the fission matrix are shown in Fig. 2.

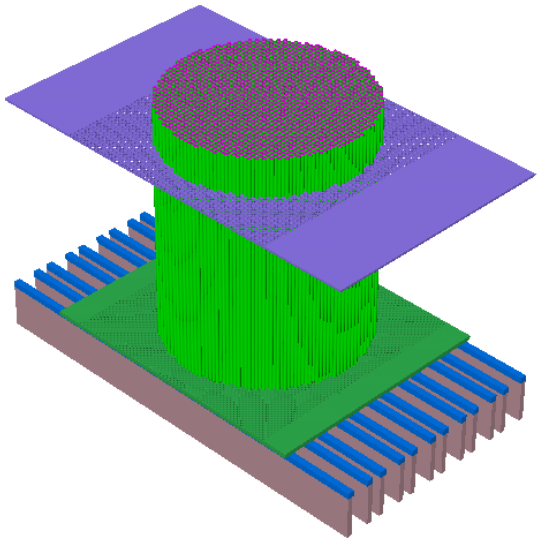

Fig. 1. Visual Workshop 3D ray-trace of the MONK geometry for Experiment 89.

The corresponding eigenvalues for the first nine eigenmodes are shown in Table I. The fundamental eigenvalue is very close to 1, as expected for this critical experiment. The dominance ratio for this system is determined to be 0.991796.

Work is currently underway to utilize the fission matrix in MONK both to accelerate the fission source convergence and to determine automatically when the fission source has converged. It is attractive to consider using the fission matrix in this way because each element of the fission matrix describes a property of the system which is independent of the global properties of the system (namely, for neutrons causing fissions in a given cell it captures information about which

[1] A superhistory comprises the histories of a source neutron and all of its fission progeny up to some maximum number of generations, L (typically 10). Superhistory powering [8] was introduced in MONK to eliminate biases in eigenvalue estimators and variance estimators.

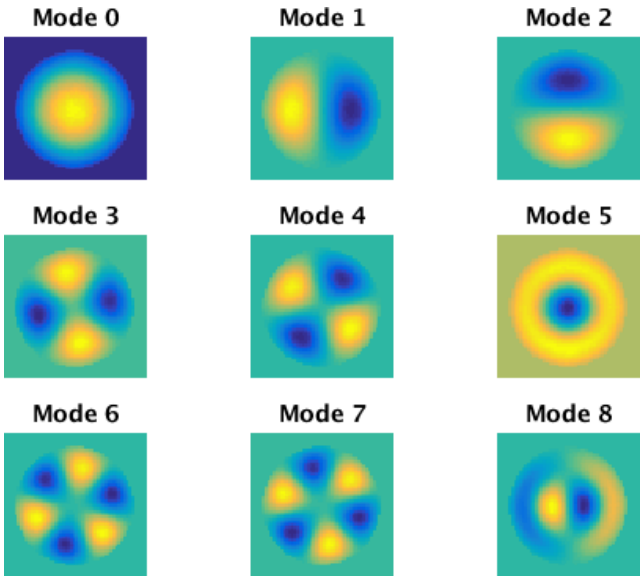

Fig. 2. The first nine eigenvectors of the fission matrix for Experiment 89, tallied in a $54 \times 54 \times 1$ mesh.

Mode	Eigenvalue	Mode	Eigenvalue
0	0.999993	5	0.979549
1	0.991789	6	0.972299
2	0.991595	7	0.972246
3	0.982376	8	0.967388
4	0.982209		

TABLE I. The first nine eigenvalues of the fission matrix for Experiment 89.

alternative cell they were born in). This means that, for an appropriately-sized mesh, the value of each element of the fission matrix will not be subject to a bias as a result of an inadequately or inappropriately converged source. Conversely the eigenvalues and eigenvectors of the fission matrix contain information from all cells of the mesh to give results which are characteristic of the system as a whole.

A number of different parameters can be calculated from the fission matrix after each stage (or batch), including:

- the absolute difference between the MONK estimate of k_{eff} and the fundamental eigenvalue of the fission matrix;
- the dominance ratio; and
- the Shannon entropy of the fundamental eigenvector of the fission matrix.

The efficacy of using these parameters as diagnostics for the convergence of the fission source distribution is currently being investigated. When suitably filtered they appear to provide a good measure of source convergence.

Since the elements of the fission matrix should converge locally, independently of the global convergence of the source distribution, they have the potential to provide additional information about the relative contributions of different regions prior to the source being fully converged. This therefore leads to the consideration of using the fission matrix to accelerate source convergence [12, 13]. There are various ways in which the information can be exploited, some of which include weighting the fission bank neutron source distribution based on:

- the fundamental eigenvector of the fission matrix;
- the diagonal elements of the fission matrix; or
- the source distribution tallied as part of the process of normalizing the fission matrix.

Preliminary results of testing these approaches have shown some promising results, but also some difficulties relating to convergence and stability issues, as have been found by other studies (e.g. [13]).

III. GEOMETRY

1. CAD Import

The MONK CAD (Computer Aided Design) import facility enables a user to import a geometry model in IGES (Initial Graphics Exchange Specification) format and perform Monte Carlo radiation transport calculations directly in that geometry without the need to convert it to a MONK geometry format.

Tracking in imported IGES geometries is generally less efficient than tracking in the native geometry packages in MONK: Fractal Geometry (FG), which uses simple body tracking; and Hole Geometry, which uses Woodcock [14] tracking. In MONK 10B the IGES import option has therefore been further developed to give:

- improved run-time performance and acceleration, using various optimizations including a voxel grid approach to reducing intersection calculations;
- increased user control of material specification;
- user control over diagnostic output;
- increased integration with the MONK code structure (re-entrant loading allowing for looping calculations and error tracking with sensible messages before exiting the code); and
- improved visualization in Visual Workshop.

Table II demonstrates the improvements made in run-time performance by comparing the speed of tracking neutrons in an IGES model with that of tracking in an equivalent FG model.The performance metric used here is the average number of samples per second, i.e. the numbers in the table represent the number of samples per second completed when tracking in the FG model divided by the number of samples per second completed when tracking through the IGES model. The k_{eff} estimators for these cases agree for both IGES and FG geometries to within expected statistical limits.

In terms of the user input, the main new feature is the use of *subfigures*. The material naming of model components

Model	MONK 10A (FG/IGES)	MONK 10B (FG/IGES)	Speed-up factor
Transport Flask	39.0	15.0	2.6
Fuel Flask	262.8	52.9	4.95
Barrel Room	21.7	12.1	1.8

TABLE II. Improvements in run-time performance for imported CAD models.

within an IGES model can take one of two forms; *default* or *updated*. The default is for each subfigure instance (multiple occurrences of a unique component) to have an identical material name as the subfigure definition (unique component) which defines its geometric shape. However the option exists to update the material names in order that different subfigure instances can use different materials. As an example Fig. 3 shows a cut-away view of a model of a barrel-filled room, which has been imported in IGES format. In the original model each barrel contains a mixture of UO_2 and concrete, but subfigures have been used here to replace the contents of some of the barrels with void.

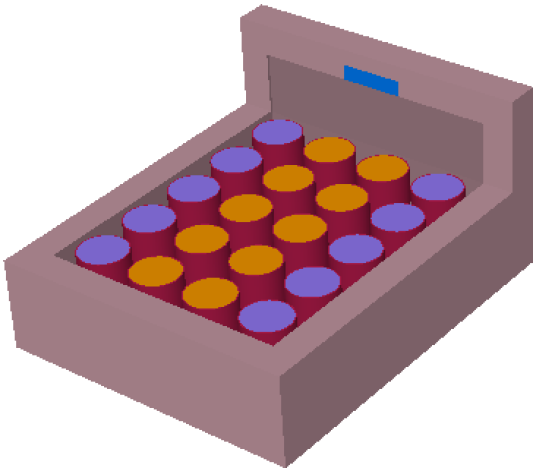

Fig. 3. An example of an imported CAD model with material substitution used to replace the contents of some of the barrels.

2. Part Reflection

MONK features a hierarchical geometry capability in which components (or *parts*) can be defined once and used many times in the model, with an optional rotation or sequence of successive rotations. To achieve this the internal data structure defining each geometric body includes a 3D rotation matrix which is initialized to the identity matrix. If the user-specified input data for a body include a rotation then the specified rotation data are used to derive a rotation matrix. The rotation matrix associated with that body is then multiplied by this new rotation matrix. By default the body is rotated about its origin, but options exist to rotate bodies about other points, such that the resulting transformation includes a displacement of the body origin. If multiple successive rotations are requested then the body rotation matrix is multiplied by the matrix for each rotation in turn such that the resulting body rotation matrix represents the final orientation of the body. If the body contains subsidiary geometric detail this will all be transformed with the parent body.

This powerful feature reduces the effort required to set up models, minimizes the potential for errors and reduces the QA burden. For some applications an ability to reflect bodies proves useful, leading to further reductions in time taken to set up and verify models. While some reflections can be achieved by combinations of rotations others cannot, but since reflections can also be represented as matrices, the mechanism used for rotations in MONK may easily be extended to include reflections.

MONK 10B therefore extends this capability to include reflections in a plane. Bodies can be reflected in any of the three orthogonal planes: XY, XZ or YZ. By default the plane of reflection is the plane containing the body origin, so the origin is not displaced by the reflection. However a body may optionally be reflected in any XY, XZ or YZ plane by specifying the displacement of the plane along the normal axis. A variation of this allows a body to be reflected about a plane passing through the geometric centre of the body, for bodies which have an unambiguous centre (not cones, prisms or hemispheres). Generally these additional options result in the displacement of the body origin, unless the original body origin lies in the plane of reflection. Fig. 4 shows an example in which a range of MONK simple bodies, each of which contains subsidiary geometric detail, have been reflected in the XY plane.

Reflections can be combined with rotations and further reflections to give a powerful general transformation capability.

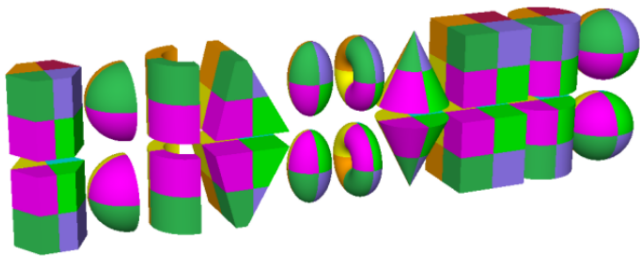

Fig. 4. Reflections applied to a variety of MONK simple body types, each of which contains subsidiary geometric detail.

3. EROD Body

The FG component of the MONK geometry package is used to define volumes of space (known as *zones*) using a range of simple geometric bodies and their unions and intersections. Many of these simple body shapes are shown in Fig. 4, each containing further geometric detail in that example.

In MONK 10B a new EROD body has been added to the set of simple bodies. This is a finite-length cylinder with an elliptical cross section, as shown in Fig. 5. It is defined by the position of the origin at the centre of the elliptical base, the major and minor radii (R_1 and R_2), and the height (length) of the cylinder (H). There are three variants (XEROD, YEROD and ZEROD) corresponding to the direction of the axis (before any reflections or rotations are applied).

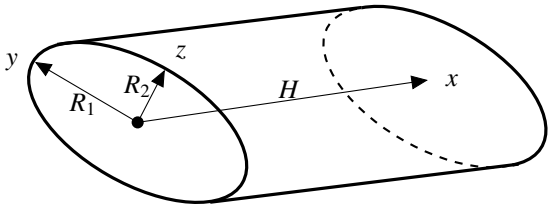

Fig. 5. Definition of an XEROD Body

IV. NUCLEAR DATA

MONK supports nuclear data libraries in the BINGO point energy format with continuous energy collision processing, the DICE hyper-fine 13,193 group format with continuous energy collision processing, and the WIMS 172-group format with broad group collision processing. The BINGO collision processor and libraries released with MONK 10A also support run-time Doppler broadening allowing accurate representation of all material temperatures in a model without being restricted to tabulated library temperatures.

MONK data libraries are available based on a wide range on nuclear data evaluations. In addition to the JEF2.2 based libraries, which have long been the *de facto* standard in the UK criticality community, libraries based on JENDL-3 and ENDF/B-VI have also been available for many years. Additional libraries released with MONK 10A included BINGO and WIMS format libraries based on JEFF3.1, JEFF3.1.1, JEFF3.1.2, ENDF/B-VII.0 and CENDL3.0.

In addition to the previously-released libraries, the MONK 10B package will also include BINGO and WIMS libraries based on the latest JEFF3.2 and ENDF/B-VII.1 evaluations.

The latest version of the JEFF evaluated file, JEFF3.2, was released in March 2014 and included updates to a fairly large range of nuclides. The main changes to the library were:

- new evaluations in the fast energy region for ^{235}U, ^{238}U, ^{239}Pu and ^{240}Pu;
- new evaluations for minor actinides, including: ^{232}Th, ^{233}U, ^{234}U, ^{241}Am, ^{242}Cm, ^{243}Cm, ^{244}Cm, ^{245}Cm, ^{246}Cm and ^{248}Cm;
- the evaluation from JENDL-4.0 was adopted for C-natural to correct an underestimation of the thermal capture cross-section;
- new evaluations for a number of structural materials, including Cr, Fe, Cu, Fr and W; and
- approximately 150 nuclides were adopted from the TENDL-2012 evaluated library, mainly for nuclides of lesser importance, but including some Zr isotopes.

The Hf evaluations and gamma production data for the 89 WIMS fission products, introduced in JEFF3.1.2, were retained.

The ENDF/B-VII.1 library was released in December 2011 and is the latest recommended US evaluated nuclear data file. It incorporates advances made in the five years since the release of ENDF/B-VII.0. These advances focused on neutron cross sections, covariances, fission product yields and decay data. The principle changes were:

- an increase in the breadth of neutron reaction cross section coverage, extending from 393 nuclides to 423 nuclides;
- covariance uncertainty data for 190 of the most important nuclides;
- R-matrix analyses of neutron reactions on light nuclei, including isotopes of He, Li and Be;
- resonance parameter analyses at lower energies and statistical high energy reactions for isotopes of Cl, K, Ti, V, Cr, Ni, Zr and W;
- modifications to thermal neutron reactions on fission products (isotopes of Mo, Tc, Rh, Ag, Cs, Nd, Sm, Eu) and neutron absorber materials (Cd, Gd);
- improved minor actinide evaluations for isotopes of U, Np, Pu and Am (not including ^{235}U, ^{238}U and ^{239}Pu) and adoption of JENDL-4.0 evaluations for isotopes of Cm, Bk, Cf, Es, Fm and some other minor actinides;
- fission energy release evaluations;
- fission product yield advances for fast fission-spectrum neutrons and 14 MeV neutrons incident on ^{239}Pu; and
- a new decay sublibrary.

Following a recent review of the IAEA Transport Regulations [15] the UK Office for Nuclear Regulation (ONR) is asking licence applicants to confirm that the effects on neutron physics of lowered or raised temperatures in transport flasks will not lead to the criticality safety criterion being exceeded. Paragraph 679 of [15] states that packages shall be designed for an ambient temperature range of -40°C to +38°C unless specified otherwise. It is noted that significantly higher temperatures could be realized inside a flask which is subjected to fire. MONK 10B is capable of modelling materials at any temperature covered by this requirement, provided that the base temperature of the BINGO library is at or below the lowest required temperature. The generation of BINGO libraries with a base temperature of -40°C is possible for free-gas nuclides, but the $S(\alpha,\beta)$ data required for thermal scattering from bound nuclides (such as hydrogen in water or polythene) are currently only available at a limited number of temperatures.

The development of bound thermal scattering data at low temperatures is currently the subject of international research effort, and preliminary data are beginning to be made available. MONK 10B has been developed and tested with prototype low temperature data and will therefore be capable of utilizing low temperature BINGO nuclear data libraries when they are available.

Code	Standard		DBRC		$10^5\Delta k$	$10^5\Delta\sigma$
	k_∞	σ	k_∞	σ		
MONK (2012)	1.31316	0.00008	1.30958	0.00008	-358	11.3
MONK (2016)	1.31276	0.00006	1.30938	0.00006	-338	8.49
MCNP	1.31137	0.00006	1.30791	0.00006	-346	8.49

TABLE III. Comparison of standard and DBRC scattering kernels in MONK and MCNP for an infinite array of PWR pincells

V. SCATTERING KERNEL

The scattering kernel determines the probability of a neutron scattering from energy E to E' and from angle Ω to Ω', i.e.

$$\sigma_s\left(E \longrightarrow E', \Omega \longrightarrow \Omega'\right) . \tag{2}$$

At sufficiently high neutron energies the motion of the target nucleus can be neglected, leading to the *asymptotic kernel*. At thermal energies the effects of crystalline structure and molecular rotation and vibrations are important, and the secondary momentum and energy are determined from $S(\alpha,\beta)$ data available for certain bound nuclides. At epithermal energies target motion cannot be ignored but molecular effects are less important. In these cases, and for nuclides for which $S(\alpha,\beta)$ data are not available, a monatomic free gas model is used. The non-zero target velocity results in a small change to the relative neutron energy, and the resultant change in $\sigma_s(E)$ should be taken into account in the scattering kernel. However it is common to neglect this effect and assume the cross-section is constant over that small energy range.

This constant cross-section approximation becomes inaccurate in the region of a neutron resonance. Neutrons scattering at a thermal or epithermal resonance may have secondary energies close to the resonance and may also upscatter. If there is also a significant capture resonance a neutron with incident energy just below the resonance may experience a higher probability of upscattering into the capture resonance. This is most significant for heavy nuclides ($A > 200$) with epithermal resonances involving significant scatter and capture components, such as ^{238}U for example.

The monatomic free gas model in the scattering kernel has been modified in MONK 10B to take account of this effect, using the Doppler broadening rejection correction (DBRC) treatment [16].

After sampling the target nuclide velocity by rejection an additional rejection test is applied. The target velocity is accepted if

$$\rho < \frac{\sigma_s(E,0)}{\sigma_s^{\max}(E_\xi,0)} , \tag{3}$$

where ρ is a random number between 0 and 1, and $\sigma_s^{\max}(E_\xi,0)$ is the maximum cross section at zero Kelvin in the energy range E_ξ, given by

$$E_\xi = \left(\frac{k_B}{A}\right)\left(\sqrt{\frac{AE}{k_B T}} \pm 4\right)^2 , \tag{4}$$

where k_B is Boltzmann's constant, A is the mass of the target nucleus and T is the temperature in Kelvin. The interval ± 4 in Eq.(4) is consistent with that used in the SIGMA1 Doppler broadening module in NJOY [17], in which the broadening integral is evaluated numerically within limits chosen such that the excluded parts of the integrand have a negligible effect on the integral.

This method requires the elastic scattering cross-section $\sigma_s(E,0)$ at a temperature of zero Kelvin in order to carry out the necessary Doppler broadening. These data are not currently part of a standard BINGO nuclear data library so it will be necessary to produce a library containing these data in order for the improved scattering kernel to be used in MONK.

The DBRC method was originally implemented in a development branch of MONK in 2012, and implemented in the MONK 10B trunk in 2016. Table III shows the results of one of the verification tests for the MONK implementation. This case is an infinite array of 4 wt% UO_2 pincells with a fuel temperature of 1200 K. This example is one of the test cases used for testing the DBRC scattering kernel in MCNP [16] and has been chosen here so that the MONK results may be compared with the MCNP results. The effect of the DBRC kernel on the reactivity in this case shows very good agreement between MONK (both the 2012 development version and the current version) and MCNP, giving confidence that the method has been correctly implemented in MONK. We note that there is a small difference in the absolute values for k_∞ between MONK and MCNP, which may be due to modelling or nuclear data differences (the MONK calculations were performed using JEFF3.1 data, while the source of the nuclear data used in the MCNP calculations is not given in [16]).

The effect of epithermal resonance scattering on criticality is only expected to be significant at high temperatures. The magnitude of the effect is found to be of the order of only 10 pcm for reactor criticality calculations at room temperature [18], suggesting that the resonance scattering effect does not need to be considered for criticality calculations at room temperature.

However for criticality calculations involving burn-up credit, although the final criticality calculation may be performed at room temperature, the effects of resonance scattering in the burn-up calculation which determines the irradiated fuel compositions may be important. As an example we consider the OECD/NEA burn-up credit criticality benchmark phase I-B [19]. The aim of this benchmark was to carry out inter-code comparisons of isotopic predictions as well as comparisons with high quality chemical assay measurements. The model used in this benchmark is a single UO_2 pincell repre-

senting a particular fuel assembly in a PWR. Operating history data (burn time, down time and boron concentration) are provided for four operating cycles. The benchmark specifies a fuel temperature of 841 K and three final burn-up values corresponding to different axial locations in the fuel assembly. Here we consider the case with a final burn-up of 44.34 GWd/t. To represent a burn-up credit calculation the final isotopic composition of the irradiated fuel following the four operating cycles and final cooling period of 1870.0 days is used in a room temperature criticality calculation, the results of which are presented in Table IV. These results indicate that the effect of resonance scattering at the reactor operating temperature in this case results in a reactivity increase of 210 pcm in a subsequent room temperature criticality calculation involving the irradiated fuel. This is most likely to be a result of the increased ^{239}Pu content of the irradiated fuel as a consequence of increased resonance capture.

Scattering kernel	k_{eff}	σ	$10^5\Delta k$
Standard	0.8252	0.0001	-
DBRC	0.8273	0.0001	+210

TABLE IV. Results of a room temperature criticality calculation using irradiated fuel compositions arising from burn-up calculations at 841 K with the standard and DBRC scattering kernels in MONK.

VI. BURN-UP CREDIT

The spatially-dependent, point-energy burn-up methodology introduced for evaluation in MONK 10A has been upgraded to a full-release feature in MONK 10B. The Unified Tally module is used to introduce spatial discretization into the model geometry so that spatially-dependendent material burn-up can be calculated without the need to specify different model materials in different spatial zones manually. This makes burn-up models easier to set up, maintain and verify.

The method is based on the concept of *artificial materials* [20]. The geometry model can be set up without regard for the spatial dependence of material burn-up, and a Cartesian *burn-up mesh* is then overlaid on the underlying model geometry. A Monte Carlo sampling algorithm identifies which materials are present in each cell of the mesh and creates a new artificial material for each, such that there is a unique mapping between each artificial material and a corresponding user-defined material in a unique cell of the mesh. The burn-up calculation is then carried out using these artificial materials to effect the required spatially-dependent burn-up. The artificial materials need only be determined in the first cycle of the burn-up calculation, and they are then stored in an *archive file* for use in subsequent cycles or indeed subsequent calculations.

In upgrading this capability to a full-release feature a number of improvements have been carried out in MONK 10B including:

- user-defined tallying of fluxes and reaction rates in artificial materials;
- improved volume estimates and material searches for artificial materials;
- improved robustness of the new burn-up methodology; and
- user image improvements.

1. COWL Improvements

The COWL method in MONK is used to transfer material compositions from a *donor* model to a *receiver* model. This can be used in burn-up credit calculations to simplify the transfer of irradiated fuel materials from a reactor model to a transport or storage criticality safety model, for example. To use the method in a receiver model, the user specifies which user-defined materials are required from the donor model, and from which cells of the burn-up mesh they should be transferred. The code determines which artificial materials correspond to the requested materials in the specified cells and extracts the relevant material compositions from the archive file generated by the burn-up calculation.

In MONK 10B COWL has been improved to give users more control over the nuclides which are imported into receiver models from donor models in order to expedite actinide-only and actinide plus fission product burn-up credit calculations. The default behaviour is to import all nuclides from donor materials, but the additional controls allow the specification of lists of nuclides (fission products and/or actinides) to *include* or *exclude* from donor materials.

Burn-up calculations can be carried out in MONK either using 172-group WIMS data or BINGO point energy data. There is not a one-to-one mapping of nuclides between the two nuclear data library formats so COWL has been further developed to include nuclide mappings such that irradiated material compositions from a WIMS data donor model can be seamlessly imported into a receiver model which uses point energy BINGO nuclear data.

VII. CROSS-SECTION TALLYING

The Monte Carlo methods used in MONK, coupled with the continuous energy treatment provided by the BINGO nuclear data libraries and collision processing mean that neutron transport can be calculated very accurately, taking resonance shielding in the resolved resonance range into account exactly. In the unresolved resonance range there are a small number of nuclides for which an enhanced resonance treatment is required, and a subgroup pre-shielding treatment is applied to the BINGO data for these nuclides.

By contrast deterministic codes, such as the ANSWERS general purpose reactor physics code WIMS [5], need to approximate the resonance shielding using approaches such as equivalence theory or the subgroup method. While these approaches are successful in many cases, they can be inadequate for some reactor geometries.

MONK 10B therefore introduces a new XT module which tallies the fluxes and reaction rates during the continuous energy Monte Carlo calculation, and uses these to calculate appropriately-shielded cross-sections in the WIMS 172 energy

group scheme for subsequent use in the WIMS code.

At each collision with nuclide n in material m the XT module tallies the collision density estimator of the neutron flux $\phi_{n,m,g}$ and the reaction rate $R_{n,m,r,g}$ where g is the energy group of the incident neutron in the WIMS 172 group scheme and r is the reaction, which can be one of:

- fission;
- n,γ;
- n,2n;
- n,3n; or
- scatter.

Also tallied are the currents, scattering matrix and, in the case of fission, the fission spectrum and the number of fission secondaries in order to estimate the average number of neutrons per fission, $\overline{\nu}$.

The relationship between the reaction rate and the macroscopic cross-section $\Sigma_{n,m,r,g}$ is given by

$$R_{n,m,r,g} = \phi_{m,g}\Sigma_{n,m,r,g}\,, \tag{5}$$

where $\phi_{m,g}$ is the material flux summed over all nuclides in the material,

$$\phi_{m,g} = \sum_{n} \phi_{n,m,g}\,. \tag{6}$$

It follows from Eq.(5) that the effective, resonance-shielded macroscopic cross-sections can be obtained from the tallied reaction rates and fluxes as follows

$$\Sigma_{n,m,r,g} = \frac{R_{n,m,r,g}}{\phi_{m,g}}\,. \tag{7}$$

The corresponding microscopic cross-sections, $\sigma_{n,m,r,g}$ are obtained from

$$\sigma_{n,m,r,g} = \frac{\Sigma_{n,m,r,g}}{N_{n,m}}\,, \tag{8}$$

where $N_{n,m}$ is the number density of nuclide n in material m.

The quantities tallied by the module are written to a binary file in a format which can be read by the WIMS code. Additionally the sums of the squares of these quantities are tallied and written to the output file in order to estimate the stochastic error in the Monte Carlo estimates of the derived cross-sections. Optionally the tallies and derived cross-sections can be requested in the MONK output file.

Fig. 6 shows an example of cross-section data tallied in the XT module. This was generated in a model consisting of an infinitely-reflected uranium/water pincell.

WIMS has a corresponding XTBUILD module which reads the file generated by the XT module in MONK and calculates the cross-sections from Eq.(7) and Eq.(8), writing them to a WIMS interface for subsequent use by other WIMS modules.

The XT module can therefore be used to generate accurate resonance-shielded cross-section data for use in WIMS in geometries for which the approximate resonance shielding methods are inadequate.

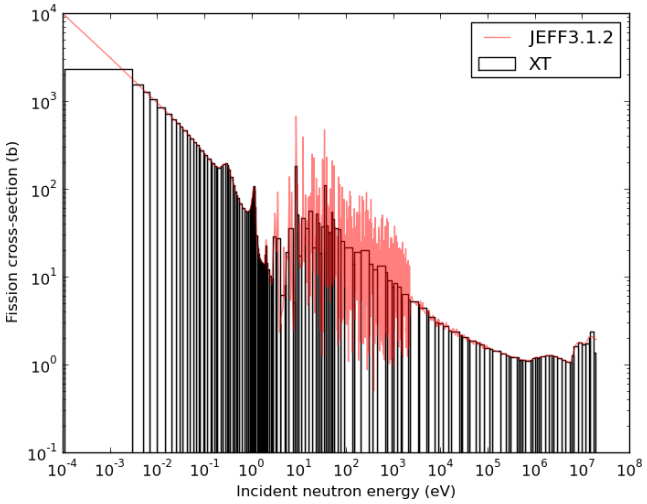

Fig. 6. Fission cross-section for ^{235}U tallied by XT in the WIMS 172 group scheme in a simple pincell, compared with the original JEFF3.1.2 cross-section data.

VIII. PARALLEL PROCESSING IMPROVEMENTS

Parallel processing was released in MONK 10A using OpenMPI to effect distributed memory parallelism in the 64-bit Linux version. While this did not include support for all MONK features in parallel, it did include the default outputs in a standard MONK criticality calculation and the reaction rates required for the depletion calculation in a BINGO point energy burn-up calculation. In MONK 10B support has been added for sensitivity and uncertainty (covariance) in parallel calculations, and all newly-added features are supported in parallel calculations, including those described in this paper.

A number of other improvements have been made to the running of parallel calculations, such as improving support for various additional output files including the birth store dump files and error files, and allowing the calculation to terminate cleanly and generate results in the event that the user requests the calculation to stop.

A powerful feature of MONK is its parameterization and its ability to *loop* over a set of parameter values, or carry out nested loops over several parameters. This allows sophisticated parameter surveys to be carried out with a single MONK execution from a single input file. MONK 10B introduces the ability to select specific loop iterations to run, including the ability for multiple iterations to be distributed across multiple processors to run concurrently, allowing efficient use of multiprocessor computer systems. This is unrelated to the MPI-based distributed memory parallelism described above and can be used on any multiprocessor platform by:

- launching individual iterations or groups of iterations as separate jobs;
- using simple scripts to distribute iterations across processors; or
- using Visual Workshop - see Section X.3.

IX. MISCELLANEOUS IMPROVEMENTS

1. Support for Large Calculations

Increases in computer performance, and particular the use of parallel processing on high performance computers, means that it has become practical to run far more samples than has historically been possible. While calculations involving very large numbers of samples are certainly possible with earlier versions, MONK 10B includes better support for such calculations, with the option to display results to greater precision, and the replacement of 32-bit integer counters with alternatives which are not limited to $2^{31} - 1$ samples.

2. Interim Tallies

Developments to the Unified Tally module allow interim flux tallies to be produced at any stage(s) of the calculation, in addition to the final flux tallies which are produced at the end of the calculation.

This feature also allows users to specify a target standard deviation for Unified Tally fluxes, and for the calculation to terminate when this standard deviation is reached on all of the requested flux tallies.

X. VISUAL WORKSHOP

Visual Workshop is the ANSWERS tool for visualizing and verifying model geometries using 2D and 3D ray-trace views and a 3D wire-frame display. A particular benefit of this tool is that the ray tracing engine is built from the same geometry source code as the physics models giving maximum possible confidence that the model geometry as displayed is the same as that sampled by the physics codes.

In addition to displaying the model geometry, Visual Workshop has a comprehensive capability for displaying model results, superimposed on the model geometry in many cases. It also contains an input file editor, and a tool for running calculations, thus forming a complete, integrated development environment for MONK and other ANSWERS products.

1. Uncertainty Analysis

A major new feature in the next version of Visual Workshop (version 3C) will be a tool known as SPRUCE, shown schematically in Fig. 7. This is designed for uncertainty analysis and allows multiple calculations to be run in which parameters defining the model geometry and material compositions are sampled from statistical distributions.

Parameters can be sampled from a range of different probability density functions, including:

- normal;
- truncated normal;
- uniform; and
- beta distributions

using a number of difference sampling strategies, including:

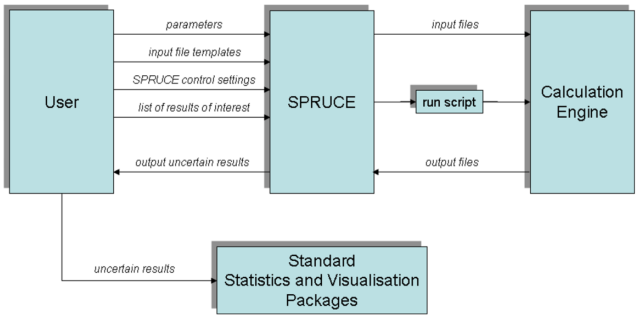

Fig. 7. Schematic of the SPRUCE tool for uncertainty analysis.

- Monte Carlo;
- stratified; and
- Latin hypercube.

Based on the selected parameters, probability distributions and sampling strategy SPRUCE automatically launches the required calculations (optionally distributed across multiple processors). The tool then extracts the requested uncertain results and allows the user to plot them against the sampled parameters or export them for further analysis, greatly simplifying these kinds of uncertainty and sensitivity analyses.

Uncertainty quantification is an imortant research area in criticality analysis and it is likely that future developments to Visual Workshop will include methods for more formal, rigourous calculations to support bias estimation in Monte Carlo criticality calculations. In particular the use of maximum likelihood estimators and extreme value theory is currently being investigated [21].

2. Optimization Tools

A suite of optimization tools is being developed for Visual Workshop, including:

- a goal-seeking algorithm [22] for determination of critical size and safe mass limits; and
- various optimization algorithms to search for the most critical configuration.

3. Support for MONK 10B

The next release of Visual Workshop will include support for new features of MONK 10B including:

- correct display of reflected bodies in 2D/3D ray-trace and 3D wireframe views;
- display of EROD bodies;
- display of the eigenvectors of the fission matrix; and
- a *split loops across multiple processors* option to run multiple loops concurrently.

XI. CONCLUSIONS

The MONK Monte Carlo code for nuclear criticality safety and reactor physics analyses is in continuous development to improve accuracy, performance, ease of use and productivity.

In order to meet these aims a significant number of developments have been incorporated into the code since the release of MONK 10A in 2014. This paper has described some of those developments, many of which will be available to users in the forthcoming release of MONK 10B.

We have also described some developments to Visual Workshop, including the uncertainty analysis tool SPRUCE, a suite of optimization tools, and the developments to support the new features in MONK 10B. These will be available in version 3C of Visual Workshop which will be released after MONK 10B.

Some of the MONK 10B developments, notably the geometry improvements, have been carried out in modules which are shared with other ANSWERS codes including the Monte Carlo general radiation transport code MCBEND, the point kernel shielding code RANKERN and the CACTUS 3D method of characteristics module in WIMS. These shared developments may therefore also be available in the next releases of those codes.

REFERENCES

1. P. SMITH, J. LILLINGTON, and C. MIDDLEMAS, "Radiation transport modelling and the ANSWERS codes suite," *Nuclear Future*, **7**, *2*, 44–49 (March/April 2011).
2. S. D. RICHARDS, C. M. BAKER, P. COWAN, N. DAVIES, G. P. DOBSON, T. C. FRY, A. KYRIELEIS, and P. N. SMITH, "MONK and MCBEND: Current Status and Recent Developments," *Annals of Nuclear Energy*, **82**, 63–73 (2015).
3. ANSWERS Software Service, *MCBEND: A Monte Carlo Program for General Radiation Transport Solutions. User Guide for Version 11*, 2nd ed. (2015), ANSWERS/MCBEND/REPORT/008.
4. ANSWERS Software Service, *RANKERN: A Point-Kernal Program for Gamma-Ray Transport Solutions. User Guide for Version 16A* (2016), ANSWERS/RANKERN/REPORT/006.
5. ANSWERS Software Service, *WIMS: A Modular Scheme for Neutronics Calculations. User Guide for Version 10* (2014), ANSWERS/WIMS/REPORT/014.
6. F. B. BROWN, S. E. CARNEY, B. C. KIEDROWSKI, and W. R. MARTIN, "Fission Matrix Capability for MCNP Monte Carlo," in "Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013," Paris, France (October 2013).
7. S. E. CARNEY, F. B. BROWN, B. C. KIEDROWSKI, and W. R. MARTIN, "Fission Matrix Capability for MCNP Monte Carlo," Tech. Rep. LA-UR-12-24533, Los Alamos National Laboratory (2012).
8. R. BRISSENDEN and A. GARLICK, "Biases in the Estimation of k_{eff} and its Error by Monte Carlo Methods," *Annals of Nuclear Energy*, **13**, 63–83 (1986).
9. http://www.caam.rice.edu/software/ARPACK.
10. ANSWERS Software Service, *MONK: A Monte Carlo Program for Nuclear Criticality Safety and Reactor Physics Analyses. User Guide for Version 10* (2014), ANSWERS/MONK/REPORT/009.
11. "International Handbook of Evaluated Criticality Safety Benchmark Experiments," Tech. Rep. NEA-1486/14, NEA Nuclear Science Committee (September 2015).
12. T. KITADA and T. TAKEDA, "Effective Convergence of Fission Source Distribution in Monte Carlo Simulation," *Journal of Nuclear Science and Technology*, **38(5)**, 324–329 (2001).
13. J. DUFEK, "Accelerated Monte Carlo Eigenvalue Calculations," in "XIII Meeting on Reactor Physics Calculations in the Nordic Countries," Västerås, Sweden (March 2007).
14. E. WOODCOCK, T. MURHPY, P. HEMMINGS, and S. LONGWORTH, "Techniques used in the GEM Code for Monte Carlo Neutronics Calculation," in "Proc. Conf on the Applications of Computing Methods to Reactor Problems," (1965), ANL-7050, pp. 557–579.
15. *Regulations for the Safe Transport of Radioactive Material 2012 Edition. Specific Safety Requirements No. SSR-6*, IAEA Safety Standards (2012).
16. B. BECKER, R. DAGAN, and G. LOHNERT, "Proof and Implementation of the Stochastic Formula for Ideal Gas, Energy Dependent Scattering Kernel," *Annals of Nuclear Energy*, **36**, 470–474 (2009).
17. D. E. CULLEN, "Program SIGMA1 (version 79-1) Doppler Broaden Evaluated Cross Sections in the Evaluated Nuclear Data File / Version B (ENDF/B) Format," Tech. Rep. UCRL-50400 Volume 17, LLNL (1979).
18. E. E. SUNNY, F. BROWN, K. B.C., and W. MARTIN, "Temperature Effects of Resonance Scattering for Epithermal Neutrons in MCNP," in "PHYSOR 2012 - Advances in Reactor Physics - Linking Research, Industry and Education," Knoxville, Tennessee, USA (April 2012).
19. M. DEHART, C. PARKS, and M. BRADY, "Burnup Credit Criticality Benchmarks - Phase I-B: Final Report for Isotopic Prediction," Tech. Rep. NEA/NSC/DOC(1996)6, NEA Nuclear Science Committee (1996).
20. S. RICHARDS, N. DAVIES, M. ARMISHAW, G. DOBSON, and G. WRIGHT, "Parallelisation of MONK with Coupling to Thermal Hydraulics and Gamma Heating Calculations for Reactor Physics Applications," in "Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013," Paris, France (October 2013).
21. M. SHEPHERD, G. DOBSON, B. LINDLEY, C. BAKER, S. RICHARDS, and P. SMITH, "Maximum Likelihood Estimators and Extreme Value Theory for Uncertainty Quantification in Criticality Analysis," in "International Conference on Mathematics & Computational Methods Applied to Nuclear Science & Engineering," Jeju, Korea (April 2017).
22. D. PUTLEY, "Development and Testing of ICASPA: a Goal-Seeking Algorithm for Critical Size Determination," in "9th International Conference on Nuclear Criticality Safety," Edinburgh, UK (September 2011).