

# Parallelisation of MONK with Coupling to Thermal Hydraulics and Gamma Heating Calculations for Reactor Physics Applications

Simon D. Richards<sup>1</sup>, Nigel Davies<sup>1</sup>, Malcolm J. Armishaw<sup>2</sup>, Geoff P. Dobson<sup>1</sup>, and George A. Wright<sup>1</sup>

<sup>1</sup>AMEC Clean Energy - Europe, Kimmeridge House, Dorset Green Technology Park, Winfrith Newburgh, Dorset, DT2 8ZB, United Kingdom

<sup>2</sup>Retired; formerly AMEC Clean Energy - Europe

Monte Carlo methods are increasingly being used for whole core reactor physics modelling. We describe a number of recent developments to the MONK nuclear criticality and reactor physics code to implement parallel processing, mesh-dependent burn-up and coupling to both thermal hydraulics and gamma transport codes. Results are presented which demonstrate the effects of gamma heating in a MONK calculation coupled to the MCBEND Monte Carlo shielding code. Experimental validation of the mesh-dependent tracking and gamma coupling methods is provided by comparison with the results of the NESSUS experiment. The gamma heating calculated by coupled MONK-MCBEND, and the neutron heating calculated by MONK, both compare well against measurement. Finally results are presented from a parallel MONK calculation of a highly detailed PWR benchmark model, which show encouraging speed-up factors on a small development cluster.

**KEYWORDS:** MONK, MCBEND, Monte Carlo, criticality, reactor physics, parallel, HPC, burn-up, gamma heating, neutron heating, validation, benchmark models.

## I. Introduction

MONK<sup>®</sup> is a Monte Carlo code for nuclear criticality and reactor physics analyses. It has a proven track record of application to the whole of the nuclear fuel cycle and is well established in the UK criticality community. Furthermore it is increasingly being used for reactor physics applications, which has led to the requirement for a number of recent developments:

### 1. Parallel processing

Monte Carlo reactor physics calculations require a large number of samples, compared to criticality calculations, in order to achieve acceptable statistical convergence of distributed parameters, such as reaction rates in individual fuel pins. This leads to far greater computational costs and hence the need to implement parallel processing to facilitate running the calculations on high performance computers (HPC).

### 2. Spatially-dependent burn-up

In order to account for the spatial dependence of material burn-up it has in the past been necessary to design models with multiple regions and materials specifically to allow material burn-up to vary spatially. This is very labour intensive and difficult to change at a later stage. We describe a new method which has been developed in MONK to allow a mesh to be superimposed on an existing model in order to account for the spatial dependence of the burn-up.

### 3. Coupled thermal hydraulics

The burn-up of materials depends on temporally and spatially varying reaction rates in the core. These reaction rates in turn depend on temporally and spatially varying material temperatures and densities. Furthermore the reaction rates depend on material compositions which vary both temporally and spatially as the materials burn up. There is therefore an implicit coupling between thermal hydraulics and core neutronics which is typically ignored in Monte Carlo burn-up calculations. We describe how MONK has been coupled to a thermal hydraulics code.

### 4. Coupled gamma transport

It is often assumed that the energy from fission reactions, is deposited in the fuel at the reaction sites. However, typically around 7% of the energy released from fission is associated with prompt and delayed gammas. These gammas can transport energy far from the reaction sites and deposit it in non-fuel materials. MONK does not have the facility to model gamma transport, so we describe how it has been coupled to a sister Monte Carlo radiation transport code, MCBEND, in order to account for the distribution of energy by gamma transport and subsequent energy deposition.

## II. Methods

In this section we describe the parallelisation of MONK using the message passing interface (MPI), the artificial material method for introducing mesh-dependent burn-up, and the coupling to thermal hydraulics and gamma transport.

## 1. MPI Algorithm

The parallelisation of MONK has been effected using the Open-MPI implementation of MPI. In this section we briefly describe the parallel algorithm.

In common with other Monte Carlo criticality codes, MONK divides the calculation into *stages*, or batches of neutrons. A fixed, user-definable number of source neutrons is tracked in each stage from their birth in fission reactions, through migration to their ultimate fate.

MONK uses a technique known as superhistory powering<sup>(1)</sup> in which it tracks the histories of each individual neutron and its fission progeny up to a user-definable maximum number of generations (the default is 10) before storing the excess neutrons to form the starting source for the next stage. This is shown schematically in Figure 1.

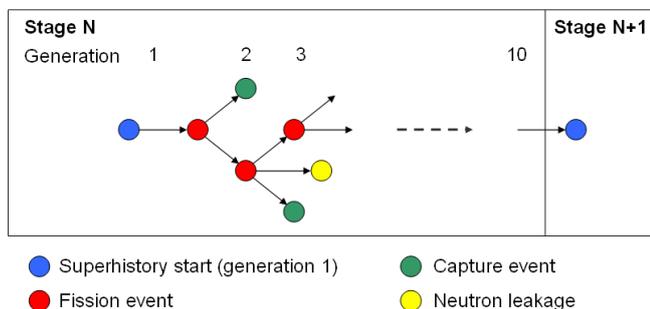


Figure 1: Schematic representation of superhistory powering.

For parallel execution on  $n$  processors one process is designated as the master; this does not track any neutrons. The number of samples to be tracked in each stage is divided equally between the remaining  $n - 1$  slave processes.

MONK does not currently implement any form of domain decomposition so each process must set up a copy of the entire model geometry, material compositions and nuclear data library. The random number generator is given the same seeds in each process to ensure that any elements of the geometry which use random numbers (e.g. some Hole geometries<sup>1</sup>) are identical in each process. However, prior to tracking, the random number generator is nudged by a different amount in each of the slaves so that each process uses a different part of the pseudo-random number sequence to ensure independence of sampling.

The initial birth store of neutrons is formed by sampling randomly in accordance with an initial source guess provided by the user. This may be simply a uniform distribution over the entire problem, or it may be restricted to specific zones and/or materials. Each slave operates on a unique part of the birth store such that each tracks a unique subset of the superhistories in each stage.

During the course of each superhistory  $m$  the following four basic quantities are scored in order to calculate estimators of  $k$ -effective:

- $H_m$  - This is the number of neutrons tracked in the superhistory;

<sup>1</sup>The Hole geometry package in MONK provides an alternative method for modelling common replicating geometries and certain complicated geometries. It uses Woodcock<sup>(2)</sup> (or delta) tracking.

- $R_m$  - This is the real number of fission children produced during the superhistory, scored at every fission event ( $\nu$ );
- $F_m$  - This is the expected number of fission children produced during the superhistory, scored at every collision ( $\nu\Sigma_f/\Sigma_t$ ); and
- $L_m$  - This is the number of neutrons which escape from the system plus the number which are absorbed during a superhistory:

$$E + \frac{\Sigma_a - \Sigma_{n,2n} - 2\Sigma_{n,3n}}{\Sigma_t},$$

where  $\nu$  is the number of neutrons produced in a fission reaction;  $E$  is the number of neutrons escaping in a superhistory;  $\Sigma_t$  is the total macroscopic cross-section; and  $\Sigma_a$ ,  $\Sigma_f$ ,  $\Sigma_{n,2n}$  and  $\Sigma_{n,3n}$  are the macroscopic cross-sections for absorption, fission, (n,2n) and (n,3n) reactions respectively.

Each slave process independently tracks the histories of each neutron within its unique set of superhistories. At the end of each superhistory the four basic tallies are summed into stage tallies, and the fission neutrons arising from the maximum generation are added to a *delay* store which will be used to populate the fission source for the subsequent stage. At the end of each stage the stage tallies are summed into cumulative tallies. A number of derived quantities are also tallied, such as the sums of squares of the basic tallies and their cross-products, which are required to calculate variances and covariances.

The stage tallies from each of the slaves are also summed into the stage tallies in the master via MPI so that the master cumulative tallies contain the contributions from all of the slaves.

At this point each of the slaves can compute estimators of  $k$ -effective from the subset of samples it has tracked, whilst the master can compute  $k$ -effective estimators based on the samples tracked by all of the slaves. The stage tallies can then be zeroed in preparation for the next stage.

The individual slave delay stores are gathered from the slaves to populate a delay store in the master process which contains all of the neutrons from all of the slave delay stores. The master process then samples from this master delay store to form the birth store for the subsequent stage.

As an alternative to running a fixed number of scoring stages MONK can optionally stop when the standard deviation of the  $k$ -effective estimator is less than a user-defined target value. The standard deviation calculated by the master will always be lower than that calculated by any of the slaves because it will always be based on more samples. Therefore the master broadcasts its standard deviation to the slaves at the end of each stage so that all processes stop tracking neutrons when the target standard deviation has been reached.

## 2. Artificial material methodology

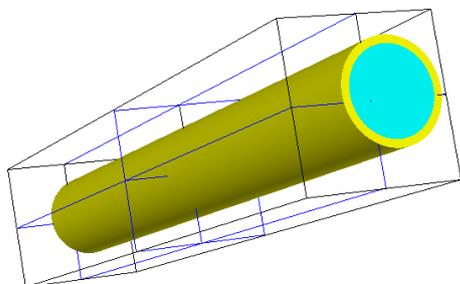
In order to model spatially-dependent burn-up of a specific material it is necessary to represent that material by a unique model material in each spatial zone in which a unique burn-up

should be calculated. In previous versions of MONK this meant that the user needed to design the model carefully to achieve this aim, manually creating the separate zones and copies of materials. This meant that the design and modification of burn-up models was onerous; it was difficult to keep track of which materials were in which zones and it was very hard to convert an existing criticality model to do a burn-up calculation.

To solve these issues a new burn-up methodology has been implemented in MONK 10, introducing the concept of an *artificial material*. This approach automates the process of creating unique materials in each spatial zone requiring unique burn-up to be calculated.

First a Cartesian burn-up (BU) mesh is superimposed over the whole problem. This is achieved using the *unified tally* module,<sup>(3)</sup> which is a new feature in MONK 10 allowing subdivided scoring bodies to be superimposed on an underlying model geometry. The unified tally mesh is completely independent of the underlying geometry, allowing both the model geometry and the burn-up mesh to be changed independently, and also allowing a burn-up mesh to be added to an existing model which had not been specifically developed to allowing spatially-dependent burn-up.

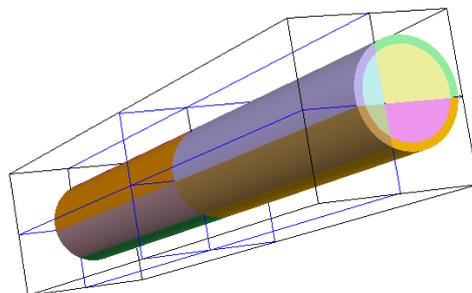
Figure 2 shows a simple geometry example consisting of two materials in nested cylinders. This could represent, for example, a small region of a fuel pin in a larger model. A unified tally mesh has been overlaid on this model; in this example a 2x2x2 mesh has been used.



**Figure 2:** A simple geometry example containing two user-defined materials with an overlaid mesh.

A Monte Carlo sampling algorithm determines which user-defined materials occur in each cell of the BU mesh and defines a unique artificial material for each instance of a user-defined material in different cells. The same algorithm also estimates the volume of each artificial material. This process is parallelised using MPI in the parallel version of MONK. A user option is provided to force the code to generate an artificial material for every defined material in every cell; this ensures that no material gets missed by the initial search algorithm at the expense of possibly storing more artificial material data than are required.

Figure 3 shows the artificial materials which would be generated in the BU mesh overlaid on the model shown in Figure 2. Since each of the two defined materials appears in each of the 8 mesh elements this simple example results in 16 artificial materials.



**Figure 3:** The same geometry example as shown in Figure 2 with sixteen artificial materials resulting from the overlaid mesh.

The Monte Carlo criticality calculation proceeds using these artificial materials to determine reaction rates which are used for the depletion calculation. The mappings between user-defined materials and artificial materials, together with their volumes, are stored in an *archive file* during the first cycle of a burn-up calculation so that they do not need to be recalculated in subsequent cycles.

### 3. Thermal hydraulics coupling

In section I-3 the implicit coupling between thermal hydraulics (TH) and core neutronics was introduced. This leads to the requirement to couple MONK to a suitable TH code. The principle behind this coupling is to allow the effect of temperature and density variations within a core on the calculation of flux, heat generation and burn-up to be modelled in MONK. Furthermore it allows the local heat generation rates calculated in MONK to be used in the TH code.

In order to transfer heating, temperature and density data between MONK and a TH code some means of mapping between the MONK geometry and the TH code geometry is required. This requires some form of spatial meshing (known as the thermal hydraulics, or TH mesh) to be used in MONK, and the scheme adopted is a Cartesian mesh using the unified tally module, as was done for the BU mesh. The coupling methodology in MONK has been designed to be independent of the choice of TH code, so external conversion utilities are used to map the MONK TH mesh to and from the meshing scheme used by the TH code (e.g. a tetrahedral mesh as commonly used in computational fluid dynamics codes).

In a burn-up calculation with TH coupling there are therefore two overlaid meshes: the BU mesh and the TH mesh. These can be identical but they do not need to be. MONK merges

the BU and TH meshes to form a combined (BUTH) mesh and it is this mesh which is used for the artificial material method described in section II-2.

During tracking MONK tallies the reaction rates in each artificial material, and uses these reaction rates to calculate the heating power in each cell of the TH mesh. These powers are normalised by the user-supplied total core power, which can be changed in each cycle of the burn-up calculation.

The heating powers in the TH mesh are then mapped to the spatial mesh used in the TH code by an external translator tool. The output of the TH code will include temperatures and densities in its native mesh, which are then mapped back onto the MONK TH mesh by an external converter utility. The revised temperatures and densities are applied to the artificial materials in the MONK calculation in the subsequent time-step of the burn-up calculation.

#### 4. Gamma coupling

In order to account for the distribution of heating by gamma migration and subsequent energy deposition, coupling between MONK and its sister shielding code, MCBEND, has been developed. This requires MONK to calculate the gamma source by artificial material and MCBEND to do a gamma heating calculation in the same geometry with the same artificial materials.

The reaction rates in each artificial material are tallied during the MONK calculation, and these are combined with gamma production data in the nuclear data libraries to calculate gamma production rates. In order to achieve this the PHODAT module from the WIMS<sup>(4)</sup> deterministic reactor physics code has been incorporated into MONK. For the gamma production calculation the 172 group WIMS energy group scheme is used for neutron interactions, and a standard 22-group scheme is used for the resultant gammas. The PHODAT module within MONK writes a file containing the calculated 22 group gamma spectrum and intensity (in photons  $s^{-1} cm^{-3}$ ) for each artificial material in a format which can be read by the MCBEND unified source module. Note that the group scheme used for the gamma production calculation is independent of that used for neutron tracking; the criticality calculation can either use point energy nuclear data or broad group data.

MCBEND has also undergone the developments necessary to track in artificial materials, and to sample source particles from artificial materials. It reads the unified source data written by MONK in the form of an embedded file. Another embedded file contains the material compositions for the artificial materials, and MCBEND also reads the MONK archive file in order to get the mappings between artificial and user-defined materials. For all cycles after the first cycle the material compositions will have resulted from the solution of the depletion equations in the MONK burn-up module. As MONK and MCBEND share common geometry, material and nuclear data packages the gamma transport in MCBEND is calculated using exactly the same materials and geometry as the neutron transport in MONK.

MCBEND has the facility to score quantities in a *splitting mesh* which is normally used in conjunction with an importance

map to carry out variance reduction using splitting and Russian roulette algorithms. Since the point of coupling MCBEND to MONK is to derive gamma heating results in whole-core burn-up problems the splitting mesh is not required for variance reduction. Therefore we can define a splitting mesh which is coincident with the MONK TH mesh in which to score energy deposition. This has the effect of giving gamma heating data (in  $MeV cm^{-3}$ ) in the TH mesh. These heating rates are returned to the MONK burn-up calculation where they are added to the neutron and fission heating rates to determine the total heating rates in the TH mesh for the thermal hydraulics calculation.

### III. Job control

#### 1. Introduction

With the potential to couple MONK to MCBEND for gamma transport modelling, and a thermal hydraulics code with associated mesh conversion utilities, running coupled calculations becomes more complicated. This is further compounded by the need to run many cycles in a burn-up calculation, with the possibility of changing the core power and step length in each cycle.

To manage the running of such calculations a Perl script has been developed. This script reads a user-supplied input file and automates:

- running the MONK criticality and burn-up calculations for each cycle, with requested core power and step length;
- running a coupled MCBEND calculation for each cycle if gamma coupling is requested;
- running a coupled thermal hydraulics calculation for each cycle if requested;
- iterating the thermal hydraulics and MONK calculations until convergence is achieved if two-way thermal coupling is requested (see section III-2);
- control rod movements;
- searches for critical rod positions;
- branching calculations; and
- special branches to specific xenon and samarium states.

The control script currently supports parallel execution of the MONK criticality calculation.

#### 2. One-way and two-way thermal coupling

If thermal coupling is requested for any cycle, this may either be specified as one-way or two-way. In the case of one-way coupling the temperatures and densities calculated by the TH code are applied to the artificial materials in the MONK calculation, but the heating rates calculated by MONK are not subsequently used in a TH calculation.

If two-way coupling is requested the heating rates generated by MONK are used in the TH calculation and the calculated material temperatures and densities are used by MONK in

the subsequent cycle. Since the heating rates depend on the reaction rates which in turn depend on the material temperatures and densities, and the temperatures and densities depend on the heating, the option is provided to iterate the MONK and thermal hydraulics calculations within a burn-up cycle until they converge. The convergence criterion depends on the standard deviations of the track-length estimates of scalar fluxes in the TH mesh in the MONK calculation.

### 3. Control rod movements

The script allows for automated control rod movements by taking advantage of the facility in MONK and MCBEND to define parameters in the input file, and also to read embedded files.

Based on the control rod movements requested by the user the script writes a small file containing lines of the format

```
@Auto_rod_insert_1 = 95.0
@Auto_rod_insert_2 = 90.0
...
```

Each of these lines defines a sequentially-numbered parameter to have a value which is nominally the percentage insertion for a given control rod or bank of control rods. This file is read as an embedded file at the top of the MONK or MCBEND input file and these parameters may be used in the model geometry to define the positions of control rods or banks of control rods.

### 4. Control rod searches

An additional feature of the job control script is the ability to search for critical rod positions in any given cycle. This is achieved by running a MONK criticality calculation and extracting the value of  $k$ -effective and its standard deviation from the output file. If the value of  $k$ -effective differs from 1.0 by more than a preset number of standard deviations the rod insertion parameters are either increased or decreased, depending on whether the system is supercritical or subcritical respectively, and the MONK calculation is repeated. This process continues until the calculated  $k$ -effective is within the preset number of standard deviations of unity, at which point the burn-up calculation continues. During these rod searches all of the rod insertion parameters are incremented equally.

## IV. Results

### 1. Gamma coupling

In order to test the mechanism of coupling MONK and MCBEND to calculate the redistribution of energy by gamma transport a simple test case was devised. The geometry for this model is shown in Figure 4. This consists of a central rod containing a mixture of fissile material, water and structural material, with four zirconium rods surrounding it. Three of the zirconium rods are surrounded by water, the fourth is bare. All of the rods are of length 1 m and the overall model forms a cube with sides of length 1 m.

An overlaid BUTH mesh has 30 elements in each of the  $x$  and  $y$  directions and a single element in the  $z$  direction. This

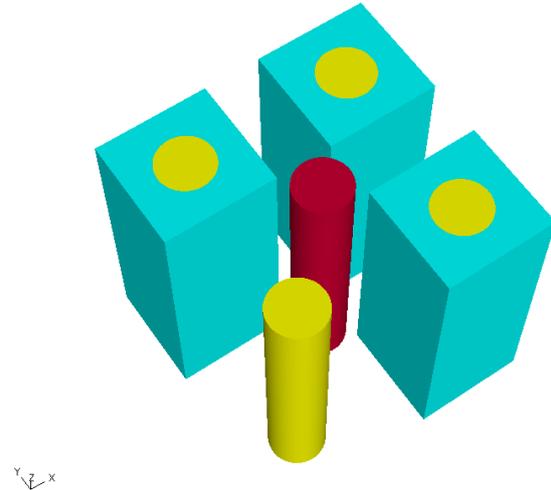


Figure 4: Geometry for a simple test case; the red material is fissile, yellow is structural and blue is water.

yields a total of 900 elements; 2700 artificial materials are therefore generated from the 3 defined materials in the problem. The PHODAT module in MONK calculates the gamma source in each of the artificial materials, as shown in Figure 5. Here we can see that the dominant source of gammas is the fissile rod in the centre of the model, as expected. A much lower gamma source is shown by the zirconium rods, arising from  $(n, \gamma)$  reactions involving neutrons which have travelled from the fissile region. The water surrounding the three flooded zirconium rods shows an even lower gamma source intensity.

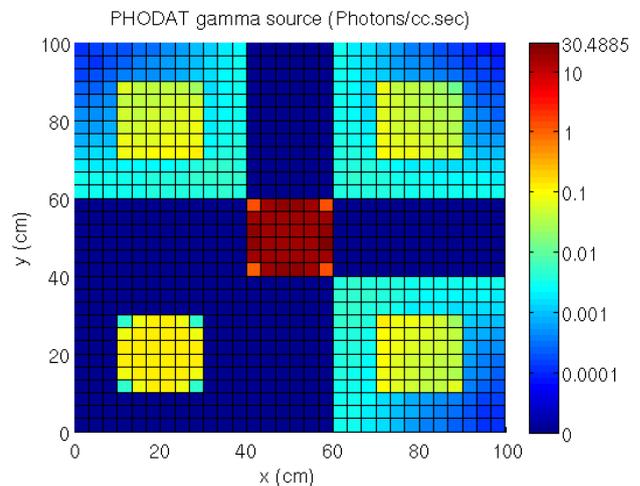


Figure 5: Gamma source in a thermal hydraulics mesh for the simple test case shown in Figure 4.

Figure 6 shows the gamma heating in the mesh resulting from the gamma source shown in Figure 5, as calculated by MCBEND. The central rod containing the fissile material shows the highest level of gamma heating, as expected. However, the zirconium rods and the water surrounding the flooded rods show a much higher level of gamma heating than would be expected from the respective source intensity, demonstrating the redistribution of gamma energy by gamma transport calculated by MCBEND.

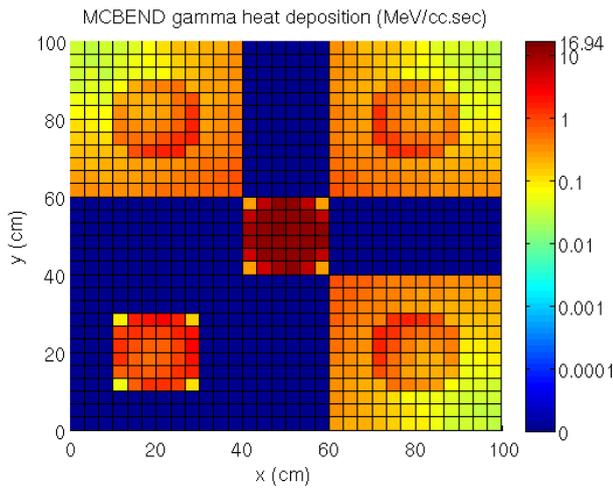


Figure 6: Gamma heat deposition in a thermal hydraulics mesh for the simple test case shown in Figure 4.

## 2. Burn-up

Extensive testing has been carried out to compare the results of the new burn-up methodology, using both point energy and broad group calculations, with calculations carried out using the old burn-up method in MONK, and deterministic calculations using the WIMS reactor physics code. Such comparisons have been carried out for a variety of reactor types including: AGR, CANDU, MAGNOX, PWR, VVER and several types of research and material test reactors. These investigations compared calculations of  $k_{\infty}$  as well as the evolution of number densities of certain key nuclides over very long burn-up periods.

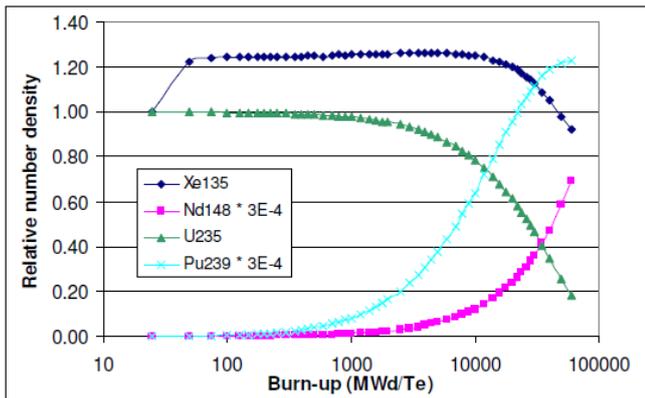


Figure 7: Relative number densities of key nuclides as a function of burn-up (note that the results for Nd148 and Pu239 are on a different scale).

A full discussion of these results is beyond the scope of this paper, but an example result is shown in Figure 7. This shows the variation in the relative number densities of Xe135, Nd148, U235 and Pu239 in a single, infinitely reflected 17x17 array of PWR fuel pins. This compares well with the other calculations discussed above and contributes to the verification of the new methodology.

## V. Experimental validation

### 1. Introduction

The NESSUS Experiment was carried out as part of a project which involved both measurement and calculation of nuclear heating.<sup>(5)</sup> It was performed on the NESTOR 30 kW source reactor at the United Kingdom Atomic Energy Establishment Winfrith, in the late 1970s and early 1980s.

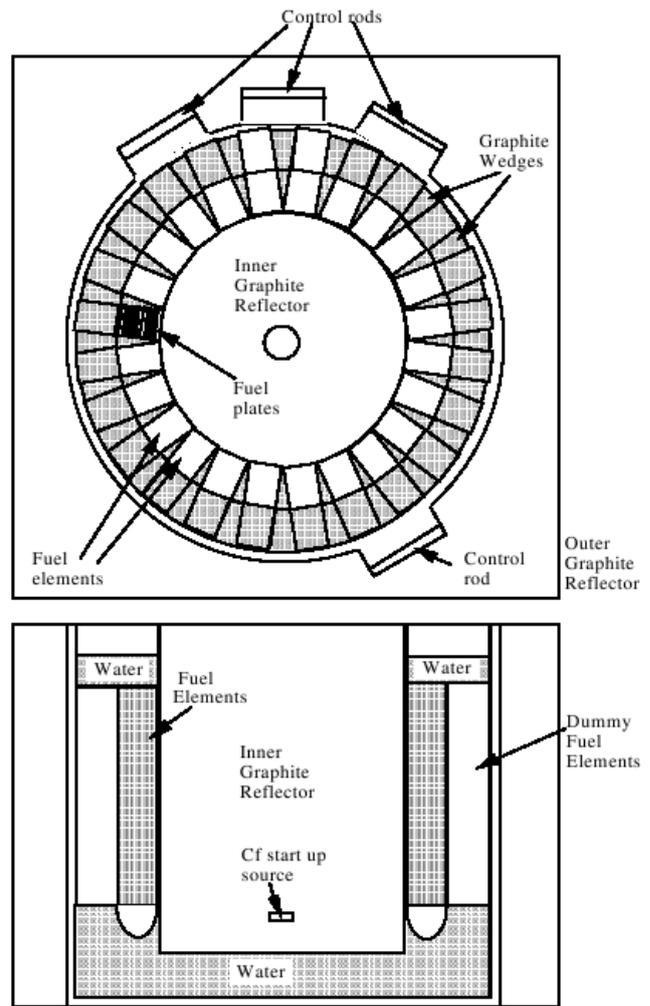


Figure 8: Horizontal and vertical cross-sections of the core of the NESTOR reactor.

NESTOR was a light water cooled, graphite and light water moderated reactor fuelled with enriched (80%) uranium fuel elements. Figure 8 shows a schematic of the reactor. The reactor core was contained in the region formed by two coaxial cylindrical aluminium vessels. The outer vessel was enclosed by an outer graphite reflector which was 5ft x 5ft x 4ft high. The inner vessel enclosed the central graphite thermal column of diameter 2ft. The fuel elements could be loaded into the annular core in two rows. A fuel element consisted of an assembly of 16 MTR fuel plates with spacers. The spaces between fuel elements were filled with aluminium clad graphite wedges. The inner fuel row contained the full complement of 24 fuel elements whilst the outer row contained only two fuel elements. The remaining fuel positions were filled with

aluminium clad graphite blocks whose external dimensions were the same as those of the fuel assembly. The control rods and reactor instrumentation were positioned in the outer graphite reflector leaving the central column clean of neutron absorbers.

The NESSUS facility provided easy access to the radiation environment that existed along the axis of the central graphite column of the NESTOR reactor. An aluminium thimble of wall thickness 4 mm extended from the top of the rotating shield down into the central hole of the graphite inner reflector, displacing a 100 mm diameter graphite plug from the reflector. A guide tube linked the thimble to the access hole in the upper biological shield. Samples or instruments were carried down into the thimble inside a special sample holder made of graphite which was suspended by a cable from a hoist located on the upper biological shield.

In the NESSUS experiments the core power was in the range 1 kW to 5 kW and the heating rates at the centre of the reactor were very low, of the order of  $\mu\text{Wg}^{-1}$ . Thus calorimeters with the necessary sensitivity were required.<sup>(5)</sup> A number of different materials were used in the calorimeters: graphite; aluminium; iron; and a hybrid iron/graphite version. In the current work only the graphite calorimeter is of interest since this material minimises the contribution to heating from gamma photons originating in the calorimeter itself. Thus the graphite calorimeter predominantly measures heating due to gamma photons that have originated in the fuel and elsewhere in the reactor as well as a component from neutrons originating in the fuel. In addition to the calorimeter measurements of total heating, measurements were made using both ionisation chambers and thermo-luminescent detectors (TLDs) in order to estimate the gamma heating. Measurements were also made using fast neutron dosimetry detectors, primarily to monitor NESTOR power but also to infer the neutron heating.

The results of the heating measurements for the graphite calorimeter are shown in Table 1. It can be seen that the measurement of total heating is consistent with those of the heating components. The latter will only be used as a guide in this work: the aim is to reproduce the total heating result.

Detector	Particle type	Heating rate
Graphite calorimeter	Neutron and gamma	$25.2 \mu\text{Wg}^{-1} \pm 3\%$
BeO TLD	Gamma	$21.4 \mu\text{Wg}^{-1} \pm 3\%$
Ionisation chamber	Gamma	$21.6 \mu\text{Wg}^{-1} \pm 3.3\%$
Fast neutron detectors	Neutron	$4.1 \mu\text{Wg}^{-1} \pm 3.9\%$

**Table 1: Experimental measurements of NESSUS heating rates.**

The heating due to fission neutrons in the calorimeter is not included in the MONK-MCBEND route. The reason for this is that the non-gamma heating in each artificial material (and hence burn-up mesh) is calculated as

$$P = R_f \bar{Q} - P_\gamma^0,$$

where  $R_f$  is the fission reaction rate,  $\bar{Q}$  is the average energy released per fission and  $P_\gamma^0$  is the gamma heating power in the source positions from MONK.

The gamma heating from MCBEND is added to the non-gamma heating. This means that for meshes that do not contain

fissile material, such as the calorimeter, only gamma heating will be scored, with no neutron heating. The MONK-MCBEND method does not account for the heating that is lost from the meshes containing fissile material by neutrons migrating to other meshes with subsequent deposition of recoil energy. However, it does include energy lost from these meshes via inelastic and capture gamma photons arising from neutron interactions. For those meshes that contain fissile material this is a small approximation (a few percent at most), and so is acceptable. However, for non-fissile meshes it means that a significant proportion of the heating will not be scored. Thus, whilst the MONK-MCBEND method may be considered appropriate for meshes containing fissile material, it is not correct for other meshes. This is a limitation of the method as it currently stands and future developments of the method will include neutron heating in all regions. A separate approach has been taken to determine the neutron heating, as discussed in section V-5.

The heating due to prompt fission gammas is included in the method. The fission reaction rate calculated in MONK is combined with PHODAT data to produce the gamma photon source for MCBEND.

Gamma heating due to decay of pre-existing fission products is not included but is not required since this forms a background which is subtracted from the experimental heating results.

Gamma heating due to the decay of fission products that are produced during irradiation are required, both for long-lived and short-lived fission products. These are included by combining the fission reaction rate from MONK with PHODAT data. These PHODAT data are associated with the actinides that produce the fission products rather than with the fission products themselves. Thus the fission products do not need to be present in the calculation for this component to be included. This means that in the analysis performed here this component will be present in the cycle 1 heating (which has no fission products) as well as the cycle 2 heating. The model results will be for fission products in secular equilibrium, whereas equilibrium was not reached in the experiment. Inspection of the data in the PHODAT library shows that all of the expected energy arising from decay of the fission products appears to be accounted for. Thus the issue of short-lived fission products, which was a problem in the previous analysis,<sup>(5)</sup> appears to have been solved satisfactorily.

Heating due to secondary gamma photons arising from capture and inelastic scatter in non-fission product nuclides is included in the method. Reaction rates (total for non-fissile and non-resonance energy ranges; fission, scatter and capture for fissile and resonance energy ranges) are combined with PHODAT data to produce the gamma photon source for MCBEND.

Heating due to activation product decay gammas from activation of non-fission products is included in the MONK-MCBEND route. Capture reaction rates from MONK are combined with relevant data from PHODAT to produce the gamma source. The only important nuclide for this is Al28. The MONK-MCBEND route will give the gamma source for activation products in secular equilibrium whereas the experimental results would not have been obtained at equilibrium.

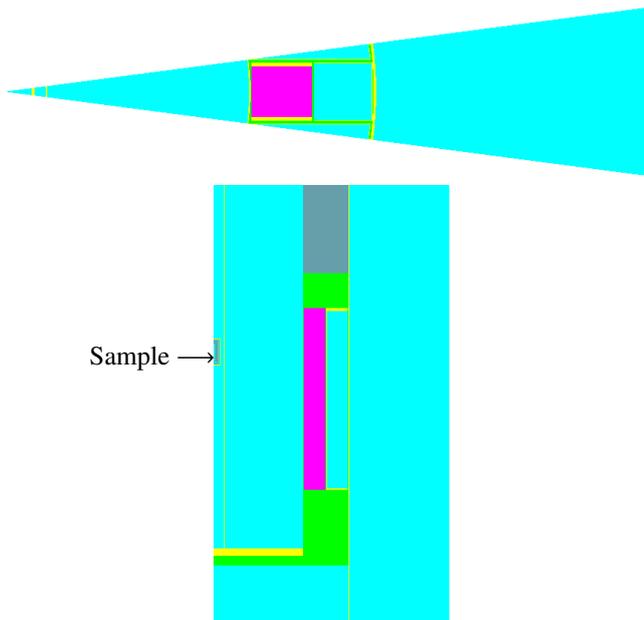
Heating due to secondary gammas and activation gammas arising from neutron interactions with existing fission products

is not included. This is because these fission products are not included in the model rather than due to the method. The component due to neutron interactions with fission products arising from the irradiation will be included in cycle 2 but due to the very low irradiation (20 minutes at 1 kW) is expected to be negligible and indistinguishable.

## 2. Geometry, materials and BUTH mesh

The NESSUS experiment forms part of the MCBEND Validation Database. The MCBEND model from the database was used as the basis for the current MONK-MCBEND modelling of the experiment to facilitate comparisons with earlier analyses.

The model, shown in Figure 9, consists of a 1/24th sector of NESTOR with reflecting lateral boundaries. The angle of the sector is thus 15 degrees. The actual reactor contains 24 fuel elements on one radius and two on a slightly larger radius: these outer elements were not included in the model and were accounted for by increasing the MCBEND source strengths. The material composition and densities used in this model are shown in Table 2.

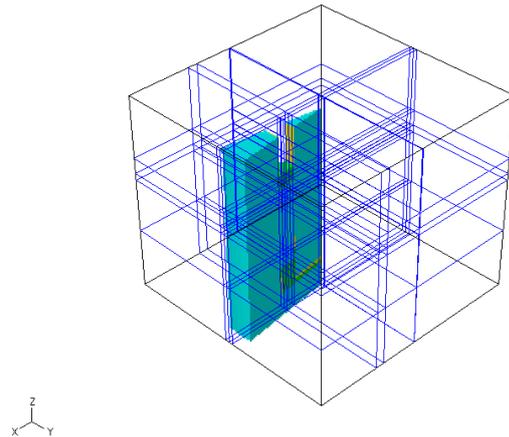


**Figure 9: Horizontal and vertical cross-sectional views through the fueled region of the NESSUS model in MONK and MCBEND (not to scale). Blue regions are graphite; yellow is aluminium; green is water; magenta is fuel/water smear; and grey is void (air).**

Material	Density ( $\text{gcm}^{-3}$ )	Element	Proportion by weight
Graphite	1.65	C	1.0
Aluminium	2.70	Al	1.0
Water	1.00	H	0.112
		O	0.888
Fuel/water smear	1.33	Al	0.356
		U	0.029
		H	0.069
		O	0.546

**Table 2: Material densities and compositions used in the model.**

A 5x5x5 BUTH mesh, shown in Figure 10, was used in which the fuelled regions as well as the calorimeter region were matched as closely as possible. This mesh necessarily encloses the whole of the problem geometry, including regions of absorbing and reflecting boundaries. As described in section II-2, an artificial material is created for each user-defined material in each mesh cell, resulting in 500 artificial materials for this case, although not all of these are required.



**Figure 10: 3D ray-trace of the NESSUS model showing the 5x5x5 BUTH mesh**

## 3. Modelling parameters

MONK can be run within the coupled MONK-MCBEND route using either BINGO continuous energy nuclear data or WIMS 172 group nuclear data. Both modes were used in this work. The MONK calculations were performed with 21 settling stages and 200 scoring stages, each with 1000 superhistories per stage.

The standard MONK-MCBEND route does not employ any variance reduction techniques in the MCBEND calculation. In the NESSUS benchmark the sample region is very small (1 cm diameter and 5.8 cm long) so achieving a low stochastic error on the gamma heating in that region would require very long calculations. As heating results were only required in this region the file written by MONK giving the splitting mesh and importance map was not used; instead a rudimentary importance map, involving splitting/Russian roulette in space, but not energy, was used. This gave a factor of 20 improvement in calculation efficiency and so was used throughout the MCBEND analyses. MCBEND was run for 80 million samples. The stochastic uncertainty on gamma heating in the sample mesh is then around 3-4% at the one standard deviation level.

After MCBEND has calculated the gamma heating power in each mesh this is added to the non-gamma heating power calculated by the MONK burn-up module and normalised to a user-specified power. In this case the power was normalised to 1/24th kW since a 1/24th sector of the reactor is modelled and the experiments were performed with a full core power of 1 kW.

#### 4. Gamma heating results

Result	MONK calculation type	
	BINGO	WIMS
Cycle 1 heating	24.1 $\mu\text{Wg}^{-1}$	24.5 $\mu\text{Wg}^{-1}$
Standard deviation	3.4%	2.7%
Standard error	1.0%	0.8%
Cycle 2 heating	23.3 $\mu\text{Wg}^{-1}$	24.1 $\mu\text{Wg}^{-1}$
Standard deviation	3.9%	4.5%
Standard error	1.2%	1.3%
Cycles 1 and 2 combined	23.7 $\mu\text{Wg}^{-1}$	24.3 $\mu\text{Wg}^{-1}$

**Table 3: Calculated gamma heating results in the NESSUS calorimeter sample.**

As stated above, the MONK-MCBEND results for mesh cells which do not contain fissile material only include gamma heating. The average values of the gamma heating in the mesh produced are shown in Table 3, together with the standard deviation and the standard error on the mean. These values have been converted from MW, as in the output of the MONK burn-up module, to  $\mu\text{Wg}^{-1}$ . The mass of graphite in the scoring mesh is 0.313 g. It can be seen that the standard deviations on the distributions, ranging from 2.7% to 4.3%, are consistent with the MCBEND standard deviation on the gamma heating of between 3% and 4%. Thus the stochastic uncertainty from MONK is not increasing the overall uncertainty significantly. Using the standard error on the mean gives stochastic uncertainties on the results of up to 1.3%.

The gamma heating results for cycle 2 appear to be slightly lower than those for cycle 1. However, the differences are not statistically significant. Physically one might expect a very small increase due to additional gamma photons from capture and inelastic scatter with the fission products created at the end of cycle 1. However, the irradiation time (20 minutes) and power (1 kW over the whole reactor) are so small that any difference is expected to be negligible. Since the gamma heating from the two cycles is expected to be virtually the same it is considered to be reasonable to combine their results. If this is done then the gamma heating is 23.7  $\mu\text{Wg}^{-1}$  in BINGO mode and 24.3  $\mu\text{Wg}^{-1}$  in WIMS mode, with a standard error on the mean of 0.2  $\mu\text{Wg}^{-1}$  in both cases.

#### 5. Neutron heating calculations

As noted above, the coupled MONK-MCBEND route does not produce neutron heating results in any burn-up meshes which do not contain fissile material. Thus no neutron heating is included in the results for the sample mesh. The PHODAT calculation in MONK does, however, calculate neutron heating. In principle this result can be normalised to the required power using the total heating rate given in the MONK output. However, the stochastic uncertainty on the neutron heating result in the sample region will be very large because the sample region is so small. In order to address this a larger sample burn-up mesh element was used, incorporating some of the sample holder graphite as well as the sample itself. Additionally the number of scoring stages was increased to 400. This gave a result of 2.3  $\mu\text{Wg}^{-1}$  with a standard error of 2.0%.

Thus an approximate value for the neutron heating of 2.3  $\mu\text{Wg}^{-1}$  with a standard error of 2.0% has been produced outside of the standard MONK-MCBEND route. If this is added to the average cycle 1 and cycle 2 gamma heating results then the total heating is 26.0  $\mu\text{Wg}^{-1}$  with a standard error of 0.2  $\mu\text{Wg}^{-1}$  when MONK is run in BINGO mode, and 26.6  $\mu\text{Wg}^{-1}$  with the same standard error when MONK is run in WIMS mode.

#### 6. Discussion

The total neutron and gamma heating in the calorimeter evaluated using the coupled MONK-MCBEND route for the gamma heating and the PHODAT results for the neutron heating is, as stated above,  $26.0 \pm 0.2 \mu\text{Wg}^{-1}$  and  $26.6 \pm 0.2 \mu\text{Wg}^{-1}$  for BINGO and WIMS modes, respectively. The uncertainty here is purely stochastic. These total heating results compare well against the measured total heating of  $25.2 \pm 0.8 \mu\text{Wg}^{-1}$ , with calculation within 6% of measurement. Thus the MONK and MCBEND codes and their associated nuclear data appear to give accurate heating results, albeit with the neutron heating calculated outside the standard route. However, it is pertinent to consider the individual heating components and also to identify possible shortcomings in the method/analysis in order to assess the validity of the method more accurately. If one compares the gamma heating results against the gamma heating measurements given in Table 1<sup>(5)</sup> then the calculated results are around 10% to 13% too high, with the neutron heating being around half of the value inferred from detector measurements. These differences tend to cancel, of course, when the total heating is calculated. A number of possible shortcomings in the current method/analysis have been identified, in addition to the lack of neutron heating in the standard MONK-MCBEND route. These are discussed below:

- The saturation of fission products in the experiment was  $0.81 \pm 0.15$  whereas in the analysis they are assumed to be in secular equilibrium. This would make a difference of around 20% to the gamma heating arising from the decay of fission products. This component has not been isolated in the current analysis but was calculated as around 4.0  $\mu\text{Wg}^{-1}$  in previous analyses. Thus, if the previous analyses were accurate, this correction might reduce gamma heating by around 0.8  $\mu\text{Wg}^{-1}$ .
- The saturation of Al28 in the experiment was  $0.91 \pm 0.10$  whereas in the analysis here it was assumed to be in secular equilibrium. This would make a difference of around 10% to the gamma heating arising from the decay of activation products. Again this component has not been isolated here, but from previous analyses is around 1.0  $\mu\text{Wg}^{-1}$ , so the difference would be around 0.1  $\mu\text{Wg}^{-1}$ .
- The model of the NESSUS assembly is approximate since it is a sector of 15 degrees and does not include features such as control rods. Two fuel elements actually reside in an outer annulus and they are not included. In the MONK calculation this would mean that some fluxes and reaction rates in the inner annulus of fuel elements and inboard of that region may be slightly underestimated. The size of this effect is unknown though is likely to be reasonably

small as there are only two elements in the outer annulus and 24 in the inner annulus. It would be relatively straightforward to construct MONK and MCBEND models of the entire NESTOR core, but in the present work the model used for the MCBEND Validation Database was retained to facilitate comparisons with previous analyses.

- The model does not include any pre-existing fission products since to include them would necessitate burn-up calculations from the start of life of NESTOR. This might affect the transport of thermal neutrons, with less absorption in the model than in reality. It would not affect the gamma tracking. It also means that gamma heating arising from capture and inelastic scatter in the pre-existing fission products is not included in the analysis.

Notwithstanding the above comments, the results produced by the coupled MONK-MCBEND route (plus neutron heating from PHODAT) compare well against measurement. The total heating, albeit with some possibly cancelling errors, is predicted to within 6% and gamma heating to within 13% of measurement. Correction of the identified shortcomings would not be expected to alter these figures so much that the general validity of the method and data would be compromised.

## VI. Benchmark model

### 1. Introduction

The Benchmark for Evaluation and Validation of Reactor Simulations (BEAVRS) model<sup>(6)</sup> provides a highly-detailed PWR specification with two cycles of measured operational data that can be used to validate high-fidelity core analysis methods. The key parameters are summarised in Table 4.

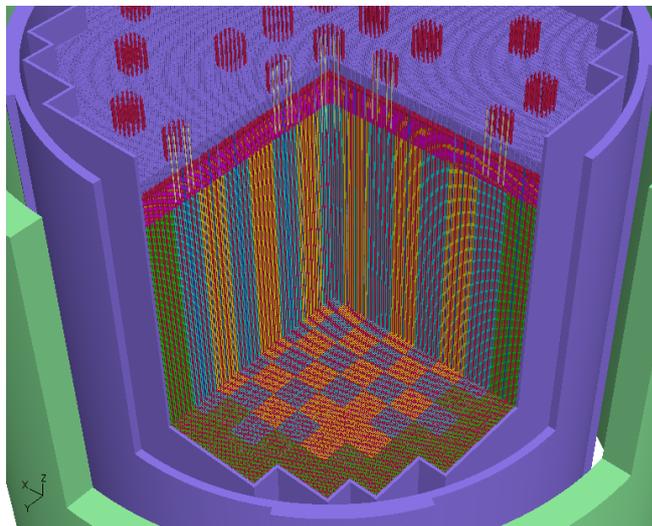
Core lattice	
No. fuel assemblies	193
Loading Pattern	
Region 1	1.60 wt% U235
Region 2	2.40 wt% U235
Region 3	3.10 wt% U235
Fuel Assemblies	
Pin lattice configuration	17 x 17
Active fuel Length	365.76 cm
No. Fuel rods	264
No. Grid Spacers	8
Control	
Control Rod Material	Ag-80%, In-15%, Cd-5%
No. Control Rod Banks	57
No. Burnable Poison Rods in Core	1266
Burnable Poison Material	Borosilicate glass, 12.5 wt% B <sub>2</sub> O <sub>3</sub>

**Table 4: Summary of the key parameters of the BEAVRS model<sup>(6)</sup>.**

As the BEAVRS benchmark provides such a highly-detailed model and requires a large number of samples to achieve acceptable statistics on distributed parameters for comparison with measured data (e.g. axial profiles of fission rates within individual fuel elements) it has been chosen as a suitable test case for parallel MONK.

## 2. MONK modelling

A MONK model has been developed from the benchmark specification. Whilst the specification itself includes some geometric approximations and simplifications, no further approximations were introduced when modelling this in MONK. The MONK model therefore includes all of the geometric detail and the material compositions specified in the benchmark model.



**Figure 11: 3D cut-away ray-trace image of the MONK model for the BEAVRS benchmark (water has been excluded to aid clarity).**

Figure 11 shows a 3D cutaway ray-trace representation of the MONK model as displayed by the MONK-MCBEND visualisation package, Visual Workshop.<sup>(7)</sup> A key feature of the Visual Workshop ray-trace is that it uses exactly the same geometry package as MONK and MCBEND so the geometry displayed is an exact representation of the geometry used in the physics codes. Water has been excluded from this image so that the structures in the model are more easily seen. The core lattice is visible, and the loading pattern of the three fuel enrichments can be seen. Also visible are the control rods penetrating the upper nozzle; the steel baffle, core barrel and neutron reflector panels; and the reactor pressure vessel. Note that the Moiré patterns seen in the figure are an aliasing artefact of the image pixelisation, and are not present in the geometry.

Figure 12 shows a close-up 3D ray-trace of the upper nozzle region near the central fuel element, showing the modelled structures in greater detail. This shows the individual clad fuel pins with plenums and end plugs; guide tubes; and top grid spacers. This figure also shows the partial insertion of the Bank D control rods into the active region, and the other rods which are fully withdrawn. Parameters in the MONK model control the movement of four separate banks of control rods, each with 228 steps between fully inserted (step 0) and fully withdrawn (step 228). Alternatively a single parameter ranging from 0 (all rods out of the core) and 574 (all rods fully inserted) can be used to move all four banks in a predetermined insertion sequence.<sup>(6)</sup> In this figure the Bank D rods are at step 213 and the other three control rods are at step 228 (fully withdrawn).

The BEAVRS model was run in MONK on a small development cluster, consisting of thirty-two cores on eight quad-core nodes. Inter-node networking on this cluster is not optimised

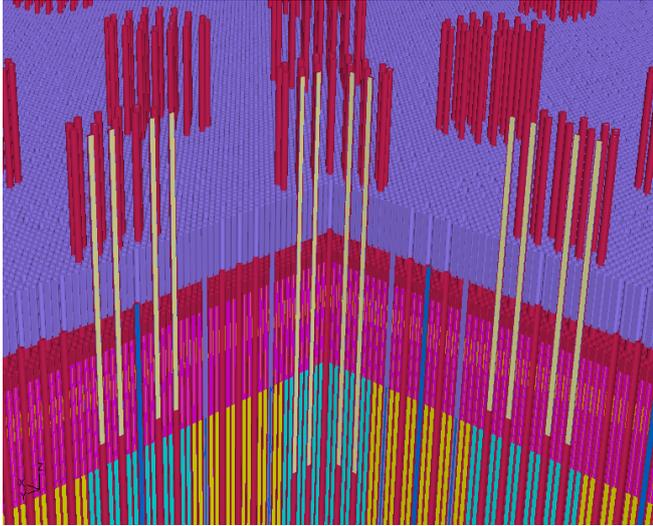


Figure 12: 3D ray-trace image of the upper nozzle region of the model, showing the level of modelling detail.

for parallel processing. At the time of this test only twenty-four cores were available for use. A small scaling study was therefore carried out by running the same calculation on 1, 4, 8, 12, 16, 20 and 24 cores. For this test 10 settling stages and 10 scoring stages were used, each consisting of 40,000 superhistories. The maximum generation number in each superhistory was 10, so each calculation tracked approximately 4 million neutrons.

For this calculation the Bank D control rods were at step 207 and all other rods were fully withdrawn. The point energy criticality calculation used the BINGO collision processor and nuclear data library based on JEFF 3.1 data. This allowed runtime Doppler broadening<sup>(8)</sup> to be performed. The benchmark specification lists an inlet coolant temperature of 560 °F for the hot zero power configuration. In the calculations presented here all materials were therefore uniformly Doppler broadened to a temperature of 566.483 K (560 °F).

### 3. Results and discussion

The results of this study are shown in Table 5. On a single processor the calculation took around 19.5 hours, reducing to around 1 hour on 24 cores. Two sets of run times and speed-up factors are reported in the table: *tracking* refers to the time which MONK reported was spent during the tracking phase of the calculation; and *elapsed* refers to the real time elapsed during the course of the calculation. Generally the elapsed time is a few minutes longer than the tracking time because of the time required to set up the model geometry and materials, and process the nuclear data. The speed-up factor for elapsed time is slightly lower than that for tracking alone because only the tracking part of the code is parallelised. These results are encouraging given the non-optimal nature of the development cluster.

Amdahl's law is commonly used to predict the speed-up factor for parallel execution of a code which is not wholly parallelised. It may be written

$$T_n = T_1 \left( B + \frac{1 - B}{n} \right),$$

where  $T_n$  is the execution time for  $n$  processes,  $T_1$  is the execution time for one process and  $B$  is the fraction of the algorithm which is strictly serial. This leads to the following expression for the maximum theoretical speed-up factor,  $S_n$ , on  $n$  processes:

$$S_n = \frac{1}{B + (1 - B)/n}.$$

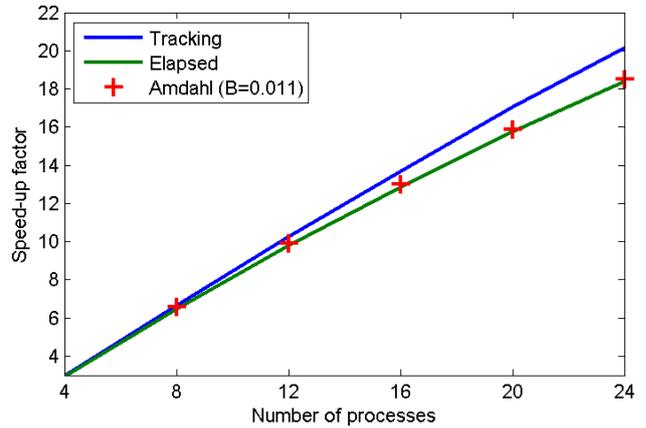


Figure 13: Parallel scaling of the speed-up factor with number of processes.

Figure 13 shows the two speed-up factors plotted against the number of processes. Also shown is the maximum speed-up predicted by Amdahl's law for a serial fraction of  $B = 0.011$ , which was found to give good agreement with the observed speed-up in elapsed time in this case. Simple extrapolation of the above equation would suggest a speed-up factor of around 53 on 128 cores and 84 on 1024 cores. However this is based on the assumption that the run times observed on up to 24 cores are close to optimum, which is almost certainly not the case for the development cluster. Furthermore, achieving acceptable statistics on distributed parameters such as fluxes and reaction rates in the BEAVRS benchmark would require significantly more samples than used in this scaling study, which would significantly reduce the serial fraction of the code and thus improve the speed-up factors on larger numbers of cores. For example increasing the number of stages by a factor of 10 might be expected to reduce the serial fraction by a similar factor. This would suggest that a speed-up of around 500 would be achievable on 1024 cores, with further improvements possible with faster inter-node communications. Scaling studies on a large, well-configured production cluster are planned, along with detailed comparisons with the real plant data published with the BEAVRS benchmark specification.<sup>(6)</sup>

The estimates of  $k$ -effective for the parallel calculations all agree with the result of the serial calculation to within three standard deviations, and all but one agree within two standard deviations, giving confidence in the parallel algorithm. Combining all of the results gives a mean of  $\bar{k} = 0.99781 \pm 0.00012$  and a reactivity of  $\rho = -222$  pcm. This suggests that the reactor is subcritical by a statistically-significant margin. It might reasonably be expected that the published benchmark specification is for a critical system, and the reasons for this apparent discrepancy have not yet been investigated in detail. Some possible

Processes	Run time (minutes)		Speed up		$k$	$\sigma$	$(k - k_1)/\sigma$
	Tracking	Elapsed	Tracking	Elapsed			
1	1169.00	1172.07	1.00	1.00	0.99805	0.00031	
4	396.32	402.02	2.95	2.91	0.99845	0.00031	1.29
8	175.77	181.26	6.65	6.45	0.99767	0.00031	-1.23
12	114.10	119.47	10.25	9.78	0.99752	0.00031	-1.71
16	86.54	91.06	13.67	12.84	0.99771	0.00031	-1.10
20	68.59	74.20	17.04	15.75	0.99769	0.00031	-1.16
24	58.00	63.52	20.16	18.40	0.99738	0.00031	-2.16

**Table 5: Parallel run times, speed-up factors and  $k$ -effective estimates for the MONK model of the BEAVRS benchmark.**

explanations might include:

- the benchmark specifies a genuinely subcritical system;
- a modelling error in the MONK input;
- uncertainties in the nuclear data;
- a systematic error in the MONK calculation; or
- the effect of uniformly Doppler broadening to 566.483 K.

## VII. Conclusions

The MONK Monte Carlo code for nuclear criticality and reactor physics analyses has undergone a number of recent developments aimed at improving its utility to reactor physicists.

A new spatially-dependent burn-up capability has been introduced, based on the concept of generating artificial materials in an overlaid mesh. This has been shown to work well, greatly simplifying the production, modification and maintenance of burn-up models.

In order to account for energy which is transported from fission sites by gamma transport, and deposited in other materials, coupling to the Monte Carlo shielding code MCBEND has been developed. Results have been presented which demonstrate this process in a simple test model. Additionally, coupling between MONK and a thermal hydraulics code has been implemented to account for the implicit coupling between thermal hydraulics and core neutronics, typically ignored in Monte Carlo reactor physics codes. This coupling has been tested and reported elsewhere.<sup>(9)</sup>

Experimental validation of some of the new developments was carried out using results of the NESSUS experiment carried out on the NESTOR reactor. This used the artificial material methodology in the new burn-up method, together with coupling to MCBEND for gamma transport and heating calculations, and the method of combining the gamma and localised fission heating into a thermal hydraulics mesh. Additionally it used data from the PHODAT module in MONK to determine neutron heating in non-fissile meshes. A number possible improvements to the coupled MONK-MCBEND route were identified, but nevertheless the calculation compared well against measurement.

MONK has successfully been parallelised using OpenMPI. A highly-detailed PWR benchmark model (BEAVRS) was used to verify that the parallelised version of MONK produced results consistent with the serial version. Scaling studies on a small development cluster were encouraging.

## Acknowledgments

The authors would like to thank David Hanlon of AMEC Clean Energy for carrying out the burn-up calculations presented in section IV-2.

The authors gratefully acknowledge the support of the United Kingdom Ministry of Defence and Rolls-Royce.

## References

- 1) R. J. Brissenden and A. R. Garlic, "Biases in the Estimation of  $k_{\text{eff}}$  and its Error by Monte Carlo Methods," *Ann. Nucl. Energy*, **13**, 2, 63–83 (1986).
- 2) E. R. Woodcock, T. Murphy, P. J. Hemmings, and T. C. Longworth, "Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry," *Proc. Conf. on the Application of Computing Methods to Reactor Problems*, ANL-7050, p. 557–579, 1965.
- 3) S. D. Richards et al., "MONK and MCBEND: Current Status and Recent Developments," *Proc. this conference*, Paris, France, October, 2013.
- 4) ANSWERS Software Service, *WIMS. A Modular Scheme for Neutronics Calculations. User Guide for Version 9*, 2004.
- 5) I. J. Curl, *Energy Deposition in the NESTOR Reactor*, PhD thesis, Imperial College of Science and Technology, 1982.
- 6) N. Horelik, B. Herman, B. Forget, and K. Smith, "Benchmark for Evaluation and Validation of Reactor Simulations (BEAVRS), v1.0.1.," *Proc. Int. Conf. Mathematics and Computational Methods Applied to Nuc. Sci. & Eng.*, Sun Valley, Idaho, USA, 2013.
- 7) <http://answerssoftwareservice.com/visualworkshop>.
- 8) C. J. Dean et al., "Validation of Run-time Doppler Broadening in MONK with JEFF3.1," *Proc. International Conference on Nuclear Data for Science and Technology*, Korea, 2010.
- 9) N. Davies, M. J. Armishaw, S. D. Richards, and G. P. Dobson, "Point Energy MONK(BINGO) TH Coupling and MONK/MCBEND Coupling," SERCO/TCS/004749.00/002, Serco Technical Services, Winfrith (2012), RESTRICTED.