

# Compression of 3D pin-by-pin burnup data

Daniele Tomatis<sup>a,\*</sup>, Aldo Dall'Osso<sup>b</sup>

<sup>a</sup>*DEN, Service d'études des réacteurs et de mathématiques appliquées (SERMA),  
CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France*

<sup>b</sup>*Framatome, Neutronics Department  
Tour Areva, 92084 Paris la Défense Cedex, France*

---

## Abstract

Several 3D distributions of pin-by-pin data are produced in nodal core calculations. In particular, burnup distributions are used as input of thermo-mechanical computer codes in one-way coupling in order to verify the compliance to safety constraints, either at operation or in accidental transients. Their storage is required because they are computed with restart from previous calculations. Since one value is needed per fuel pin and per axial level, the amount of data to be stored for one reactor configuration is very high. Data compression techniques are very powerful to reduce the storage need. They are proposed in this work to optimise the amount of data saved and processed, without compromising the final accuracy. Specifically, we recommend the Hotelling transform to achieve higher compression performances.

*Keywords:* Data compression, Reactor analysis, Pin-by-pin data

---

## 1. Introduction

The calculation of pin-wise reaction rates is mandatory in reactor analysis, and their storage takes several Giga-Bytes in memory for the many possible configurations of the reactor. Computed data are archived in external files for different purposes. Their storage allows restart capabilities in complex calculations by accessing data from earlier code execution. Backups allow also data recovery in case of deletion or corruption, and they are fundamental in quality assurance protocols for non-regression testing and version release of computer codes.

The most common application of 3D pin-by-pin data is probably the calculation of the local power and burnup distributions in nodal codes. These

---

\*Corresponding author. Tel.: +33 1 69 08 39 79.

*Email address:* [daniele.tomatis@cea.fr](mailto:daniele.tomatis@cea.fr) (Daniele Tomatis)

computer codes reproduce the physics of neutron diffusion in the core, considering also the thermal feedback on the reactivity by a few group cross sections model. The standard modelling approach in LWRs resolves the neutron balance equation on a coarse mesh whose typical cell size is of the order of the fuel assembly, or of its quarter (Lawrence, 1986). Pin-wise quantities are then reconstructed by a dehomogenization procedure (Koebke and Hetzelt, 1985). Specifically, pin burnups are currently computed by integrating the dehomogenized pin power distributions in time along fuel exposure. Its follow-up needs to store all spatial distributions on disk-file. These distributions are provided to thermo-mechanical computer codes in order to verify the fuel rod performances and the mechanical integrity during load-following operations and accidental transients.

The estimated storage need of pin-wise data corresponds to a large amount of information indeed. For example, the core of a 900 MWe French PWR has 157 assemblies, each containing 264 fuel rods, and with 32 axial nodes we have 1326336 values to store per 3D sample, which means about 5.3 Mega-bytes (in simple precision). This is about twice the size of the full data stored for one reactor state calculation point, which is based on a coarse nodal mesh instead. The use of different types of pin-data other than burnup, like general reaction rates, will increase the storage requirement. Besides, this amount will certainly increase in the future because nodal methods are expected to be replaced by pin-by-pin methods, and in this case nuclide distributions will be stored per fuel cell as well. Taking into account that this information has to be collected in all physical states of interest of the reactor, one can easily go beyond Tera-Bytes. Recently compression techniques have been applied successfully on power form factors, showing a considerable storage gain (Tomatis, 2019). Compression techniques are also used in this work to reduce this amount. We apply these techniques to compress the 3D burnup distributions with a resolution of the fuel pin in the radial direction  $xy$  and 32 axial segments in the  $z$  direction.

## 2. Pin-by-pin Burnup

The core of the reactor is composed of several fuel elements stacked together according to a regular grid. The fuel elements are in turn assemblies of thin and elongated rods surrounded by a coolant for heat removal and still arranged in a similar geometrical grid. Since the characteristic lengths of the assembly grids are usually the same, the whole construction becomes a wide lattice with the same basic geometric cell. Therefore, the spatial mesh upon which the pin burnup values are computed relies as well on a regular geometry. Ordinary grids are either Cartesian or hexagonal.

The burnup is a measure of thermal energy released in a given time per unit mass of heavy metal fuel loaded at the beginning of the reactor operation. It is customarily measured with the peculiar unit of MWd per metric tonne (t), where the amount of energy is given by the product between the effective full power in MW sustained along the days of operation (d). The varying neutron flux in space draws uneven burnup distributions in time because of leakage due to the finite domain and for the material differences of the local fuel inventories. Because of the assumptions in the calculations, the elemental values of a distribution are de facto the average burnups computed in the single pins.

We are considering square rectangular fuel assemblies in the following, because they are the most common in western LWRs. The extension to other geometries, like the hexagonal grid typical of VVER, is straightforward. Provided a core comprising  $K$  different assemblies, each made respectively of  $X \times Y$  rods along the  $x$  and  $y$  coordinates, the total amount of pin data to compute is  $X \times Y \times Z \times K$ , where  $Z$  is the number of planes chosen along the  $z$  axial coordinate. A pin burnup is then identified as a positive real-valued scalar quantity by the tensor notation  $b_{xyzk}$ , with  $x = 1, \dots, X$ ,  $y = 1, \dots, Y$ ,  $z = 1, \dots, Z$  and  $k = 1, \dots, K$ . The values are stored in contiguous array, using for instance the FORTRAN column-major ordering with memory-offset as:

$$b[(x, y, z, k)] = b[x + X(y - 1 + Y(z - 1 + Z(k - 1)))]. \quad (1)$$

In this work, we call matricization the specific operation that rewrites a multi-dimensional array as a matrix where two subsets of subscripts are selected out of the original ones. For consistency, we keep the column-major ordering with the subsets too. They designate then rows and columns of a new matrix. The operation of recurrent concatenation of all subscripts in order to obtain an array similar to the one in Eq. 1 is called vectorization instead. Still about Eq. 1, we refer to  $\mathcal{B}$  as the set of all possible matrices  $\mathbf{B}$  retrieved by matricization on  $b$ . Disregarding the order of the subscripts in the subsets, its cardinality is  $|\mathcal{B}| = \sum_{i=1}^3 \binom{4}{i} = 14$ . Here, the binomial coefficients counts the number of subsets obtained by grouping together  $i$  subscripts out of the original four on the first dimension of the target matrix  $\mathbf{B}$ . These matrices are suitable for compression by numerical transforms that implement a linear transformation through matrix multiplication.

After considering the presence of guide tubes, which do not contribute to the global burnup, it is preferable to store the pin burnup using its pin index in the corresponding vectorized matrix for each fuel assembly. This change makes one axis available for the time (or exposure) dependence in

the multi-dimensional representation. We use the Cartesian coordinates to refer to the positions of the fuel assemblies in the core, using letters and numbers for rows and columns respectively. Each position is then uniquely tagged by one letter and by one number, allowing the storage in lexical order on a single axis ( $K$  in the following).

### 3. Data Compression

Compression is about different representations of the given data in more compact forms, achieved by means of redundancy reduction with the goal to keep the initial amount of information. This statement follows as simple observation of practical applications. More rigorous definitions require a wider introduction to information theory and coding, available elsewhere (Shi and Sun, 1999).

Data compression fulfils various needs, like matching a target transmission bandwidth in communication channels or reducing the storage requirements. In this work, we focus our attention only on the second goal. Notably, compression is classified as lossy or lossless according to possible information loss occurring after reconstruction of the original signal from the compressed data. Compression algorithms are based on different steps, summarised in three main parts: a numerical transform, quantization and entropy encoding. Here, we will consider only lossy compression, by truncating out the transformed coefficients of negligible magnitude and without addressing quantization or other subsequent encoding.

The literature about compression offers many numerical transforms aiming at signal decorrelation, that is looking for a signal decomposition in a new linear space whose basis components are much less correlated. These linear transformations provide then the weighting coefficients of the basis components. The Hotelling transform (HT) and the Discrete Cosine Transform (DCT) are introduced hereafter.

#### 3.1. The Hotelling Transform

The Hotelling transform (HT) has high decorrelation capabilities, but it is signal-adaptive and requires to store also the new basis vectors, which depend on the input signal (Shi and Sun, 1999; Penna et al., 2007). The HT is often referred to as the discrete version of the Karhunen-Loève transform.

Provided a matrix  $\mathbf{B} = (b_{ij}) \in \mathbb{R}^{I \times J}$ , we interpret its column vectors as particular realizations of the random variable  $\vec{b}$  whose mean vector is  $\vec{\mu} = \mathbb{E}[\vec{b}] = (\mu_i = \mathbb{E}[b_i], i \dots, I)^T$ . If the variables  $b_i$  are uniformly distributed, the expected values can be estimated by the arithmetic averages of the  $J$  samples:  $\mu_i = \sum_j b_{ij}/J$ . The covariance between the components  $i$  and

$j$  of  $\vec{b}$  is  $\text{Cov}(b_i, b_j) = \mathbb{E}[(b_i - \mu_i)(b_j - \mu_j)]$ , or in matrix notation  $\mathbf{C}_b = \mathbb{E}[(\vec{b} - \vec{\mu})(\vec{b} - \vec{\mu})^T]$ . By definition  $\mathbf{C}_b$  is diagonalizable, because it is real and symmetric ( $\mathbf{C}_b = \mathbf{C}_b^T$ ), so that its eigen-decomposition yields  $\mathbf{C}_b \mathbf{E} = \mathbf{E} \mathbf{\Lambda}$ , with the orthogonal matrix  $\mathbf{E}$  (i.e.  $\mathbf{E}^{-1} = \mathbf{E}^T$ ) containing the eigenvectors per column and their corresponding eigenvalues in the diagonal matrix  $\mathbf{\Lambda}$ , which are all real. The HT consists in the linear transformation

$$\vec{\beta} = \mathbf{E}^T (\vec{b} - \vec{\mu}), \quad (2a)$$

and its anti-transform is

$$\vec{b} = \mathbf{E} \vec{\beta} + \vec{\mu}. \quad (2b)$$

It is easy to verify that  $\mathbb{E}[\vec{\beta}] = 0$  and that the new covariance matrix in the transformed space is purely diagonal,  $\mathbf{C}_\beta = \mathbb{E}[\vec{\beta} \vec{\beta}^T] = \mathbf{E}^T \mathbf{C}_b \mathbf{E} = \mathbf{\Lambda}$ . The two covariance matrices have the same eigenvalues for the similarity transformation,  $\mathbf{\Lambda} = \text{diag}(\sigma_i^2, i = 1, \dots, I)$ , being exactly the variances of the elements of  $\vec{\beta}$ . It follows that the new variables are fully uncorrelated. After sorting the eigenpairs in decreasing order and retaining only  $L$  terms for the reconstruction, it can be shown that the mean squared error MSE is equal to the sum of the discarded variances:  $\text{MSE} = \mathbb{E}[||\vec{b} - \vec{b}'||^2] = \sum_{i=L+1}^I \sigma_i^2$ .

The spectral factorization of the real symmetric matrix  $\mathbf{C}_b$  can be performed by the QR algorithm, or by the singular value decomposition of the matrix  $(\mathbf{B} - \vec{\mu} \otimes \vec{1}_J)$  when  $I > J$ . In this last case, the eigenvalues of the covariance matrix are simply the squared non-vanishing singular values and the left-singular vectors are their eigenvectors.

### 3.2. The Discrete Cosine Transform

The Discrete Cosine Transform (DCT) expands the input signal as a sum of cosine functions at different frequencies. Its performances in terms of energy compaction capabilities are close to the HT, with the advantage that the basis components are analytical and do not need to be stored (Shi and Sun, 1999). It is very popular in image compression, being a fundamental element of the JPEG standards (Wallace, 1992).

The DCT reproduces a periodic discrete sequence of values. The repetition can be odd or even, with symmetry at the midpoint or at the border of the peripheral cells, for a total of eight different variants. We use here the DCT-II (even at both borders), which is the most used in image and signal processing. Thanks to the tensor representation, the multi-dimensional ortho-normalized DCT-II is:

$$\beta_{lmnp} = \sum_{k=1}^K \sum_{z=1}^Z \sum_{y=1}^Y \sum_{x=1}^X d_{pk} d_{nz} d_{my} d_{lx} b_{xyzk}, \quad (3a)$$

where,

$$d_{lx} = \sqrt{\frac{2 - \delta_{l1}}{I}} \cos \frac{\pi(x - 1/2)}{I} (l - 1) \text{ for } l, x = 1, \dots, X, \quad (3b)$$

and similarly for the other 2D tensors acting on the remaining coordinates<sup>1</sup>. The inverse transformation is exactly the ortho-normalized DCT-III:

$$b_{xyzk} = \sum_{p=1}^K \sum_{n=1}^Z \sum_{m=1}^Y \sum_{l=1}^X d_{kp} d_{zn} d_{ym} d_{xl} \beta_{lmnp}. \quad (4)$$

Both direct and inverse transformations can be efficiently implemented by simple matrix multiplication and subsequent reshape of the partial products.

### 3.3. Truncation and Performance Parameters

The numerical transforms rewrite the original data through a weighted sum, where in general many coefficients can be neglected still providing a good approximation by anti-transformation. Concretely, the coefficients whose magnitude is lower than a given threshold are zeroed. This truncation process is called threshold coding. About the HT, it applies on the sorted variances, thus discarding implicitly the associated eigenvectors. Instead with the DCT, the zeroed coefficients do not follow a fixed order, forcing to keep track of their positioning, which is sometimes disadvantageous. It must be noted however, that lower frequencies take usually the higher coefficients, suggesting to save only the coefficients in a given spectral zone. This last truncation scheme is known as zonal coding.

The error after truncation is simply determined by  $e = b - b'$ , where  $b'$  indicates the reconstructed value. Performances of compression techniques are determined by the compression ratio CR, defined as the ratio between the amounts of uncompressed and compressed data. Since we always deal with real floating-point burnup values, these amounts are expressed by the number of initial values in the dataset and the number of transform coefficients retained after truncation. The memory saving is then defined as  $\eta = 1 - 1/\text{CR}$ .

## 4. Results

### 4.1. Characterization of the dataset

The studied configuration refers to a 900 MWe French reactor (157 fuel assemblies) during an equilibrium fuel cycle (Coppolani, 2012; IRSN, 2008).

---

<sup>1</sup>By the Kronecker function, it is  $\delta_{l1} = 1$  if  $l = 1$ , 0 otherwise.

The core loading consists of 125 UO<sub>2</sub> assemblies enriched at 3.25% w/o <sup>235</sup>U and 32 MOX assemblies with 5.3% and 5.4% average Pu content. The MOX assemblies have 3 zones, with the central one showing highest Pu content. The plan of the south-east core quadrant is reported in Fig. 1. The assembly type, its batch number, and the average burnups at BOC and EOC (Beginning/End Of Cycle) are specified per each FA. The core presents the half- $\pi$  rotational symmetry, which means that the total core configuration is given by rotating three times the given quadrant of 90°. The pin-by-pin burnup values are calculated by the nodal code SMART, belonging to code suite SCIENCE of Framatome (Girieux et al., 2001). The calculation geometry is composed of 1 central assembly having spent 3 cycles, 52 assemblies having spent 2 cycles, 36 assemblies having spent 1 cycle and 36 fresh assemblies in the periphery of the core.

Fig. 2 shows the axial distributions of the pin burnup for the average fuel rods of four assemblies from different batches. The average is carried on the 264 fuel rods and the coloured areas around the curves indicate the values within two standard deviations. The distributions are plotted at different times along the cycle with the average core burnup values in the legend. The figure enlightens the out-in loading pattern. In fact, fresh fuel at the position A08 has larger radial variations of the pin burnups because of the fading neutron flux near the reflector. This trend increases during the cycle exposure. As expected, higher axial values are observed in the lower part of the core for the enhanced neutron thermalization and fission due to the denser cooling water. Assemblies from other batches are exposed to a flatter flux for being farther from the boundary. Radial variations of pin burnup are smoothed out in time, since the generation of the power is favoured in the zones where the pin burnup is lower and so the concentration of fissile material is relatively higher. This is noticed in the same figure, and verified in the assembly H08 which is located at the core center. Each batch counts 52 FAs, and this last assembly is the only one to incur an additional cycle for the unequal size three batches per cycle loading. The 3D pin burnup maps are available at 20 different times along the core exposure and at 32 axial planes evenly distributed on the active height. It must be considered that the compression of the current datasets is quite challenging because the numerical values span almost ten decades in magnitude.

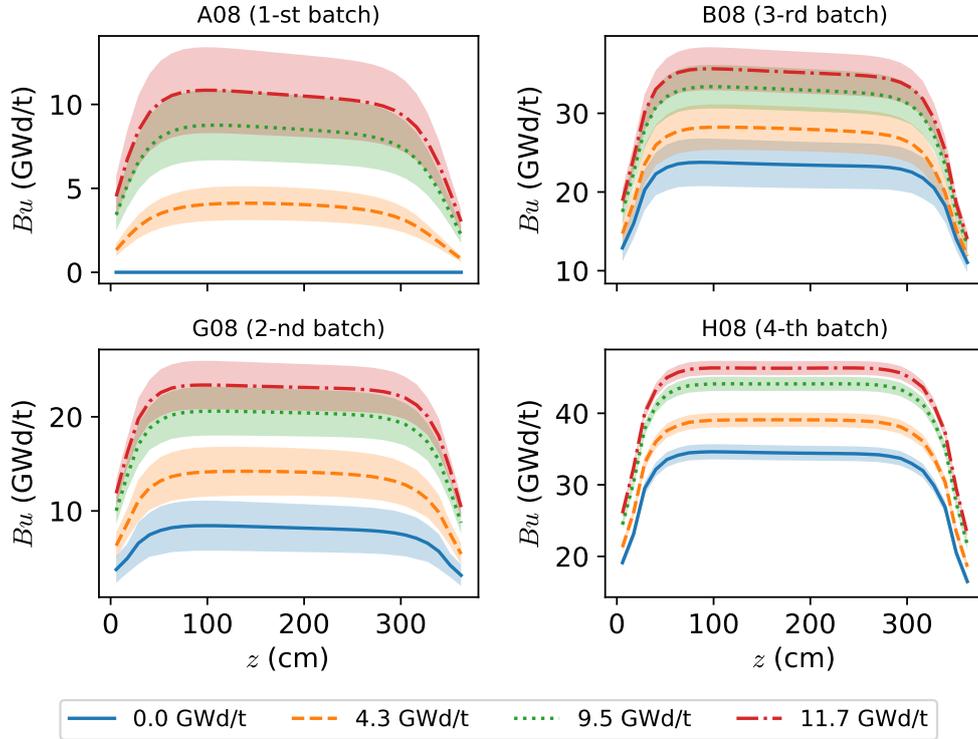
#### 4.2. Application of numerical transforms

The error observed by applying the (ortho-normalized) DCT-II on the whole dataset is shown in Fig. 3a up to a saving of  $\eta = 80\%$  (CR = 5). The interest in the DCT is motivated by the theoretical advantages introduced in section 3.2. The dotted green line marks the mean error, while the 5%, 50%

	H	G	F	E	D	C	B	A
08	UO2 (4) 31818 43028	UO2 (2) 7400 21354	UO2 (3) 19889 31885	UO2 (3) 26170 37814	UO2 (2) 8944 22532	UO2 (3) 21917 33857	UO2 (3) 21733 32215	UO2 (1) 0 9361
09	UO2 (2) 7400 21354	UO2 (3) 21786 33693	MOX (2) 11648 24973	UO2 (2) 7217 21024	UO2 (3) 26214 38018	UO2 (2) 8079 21580	MOX (1) 0 12974	UO2 (1) 0 7595
10	UO2 (3) 19889 31885	MOX (2) 11665 24984	UO2 (3) 26139 36967	MOX (3) 11754 25460	UO2 (2) 10198 23486	UO2 (3) 24216 35495	UO2 (1) 0 11030	
11	UO2 (3) 26170 37814	UO2 (2) 7216 21030	MOX (2) 11742 25517	UO2 (2) 9876 23288	UO2 (3) 21480 33155	MOX (1) 0 12663	UO2 (1) 0 7874	
12	UO2 (2) 8944 22532	UO2 (3) 26262 38052	UO2 (2) 10185 23527	UO2 (3) 21377 33127	UO2 (2) 7392 19688	UO2 (1) 0 8734		
13	UO2 (3) 21917 33857	UO2 (2) 8088 21560	UO2 (3) 24176 35477	MOX (1) 0 12702	UO2 (1) 0 8759			
14	UO2 (3) 21733 32215	MOX (1) 0 13004	UO2 (1) 0 11046	UO2 (1) 0 7890				
15	UO2 (1) 0 9361	UO2 (1) 0 7615						

FA (b)
BOC $Bu$
EOC $Bu$

Figure 1: Fuel loading plan of the south-east core quadrant (FA, fuel assembly type; b, batch; BOC and EOC  $Bu$ , average FA burnups).



**Figure 2: Axial distributions of the average pin burnup for four different FAs, selected at different core cycle burnup.**

and 95% quantiles are plotted by dashed lines. The values of the pin burnup vary on several decades over the entire core, that is between zero and a few dozens of GWd/t. Although the 95% of the errors are smaller than 0.08%, the maximum error given by the solid red line on the right axis is quite high and is always located on the fresher pins of pin burnup less than 1 MWd/t. Since these values are on the tail of the distribution and they are very low, close to 0, relative errors are high. Although the average absolute error is of a few hundreds of kWd/t, its maximum jumps rapidly above the MWd/t. We note that unfortunately the error distribution has large tail.

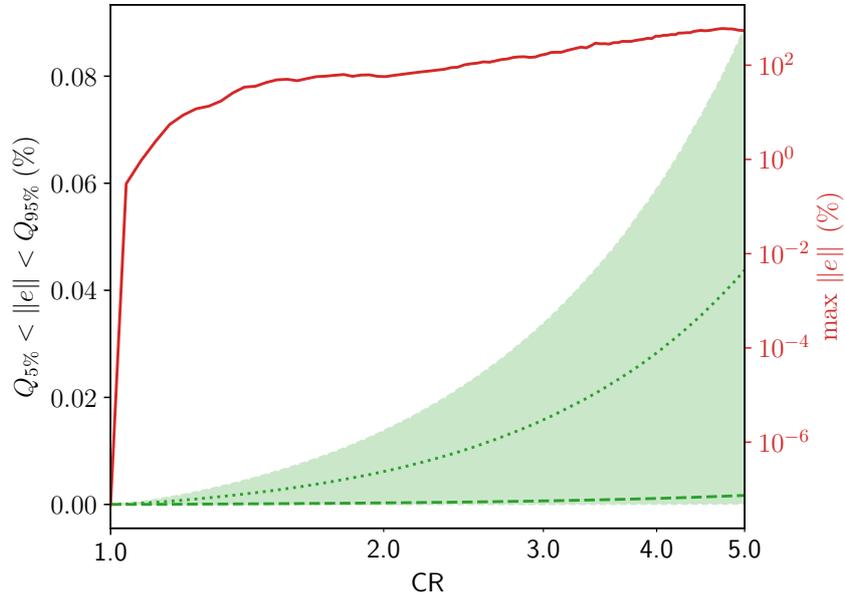
Indeed the CR should be halved to account for the additional storage for the positions of the retained coefficients, for which an unsigned char data type is assumed, except for the dimension of the number of fuel rods that demands 16 bits. An option to get rid of this disadvantage is zeroing terms by the zig-zag scheme (Shi and Sun, 1999), but no significant improvement was noticed in this case. This particular scheme selects all elements belonging to a given (multi-dimensional) plane perpendicular to the main diagonal.

When applying the transform separately per FA, it can be observed that the higher errors are due to the FAs from batches 1 and 2 whose spatial distributions of pin burnup are less uniform, see Fig. 3b. Consequently, the acceptable savings depend on the FA position in the core, with 20% as maximum storage gain for some assemblies from batch 2 still aiming at errors smaller than 1 MWd/t. However, separate compression by FA can achieve more interesting scores for about a third of the total FAs.

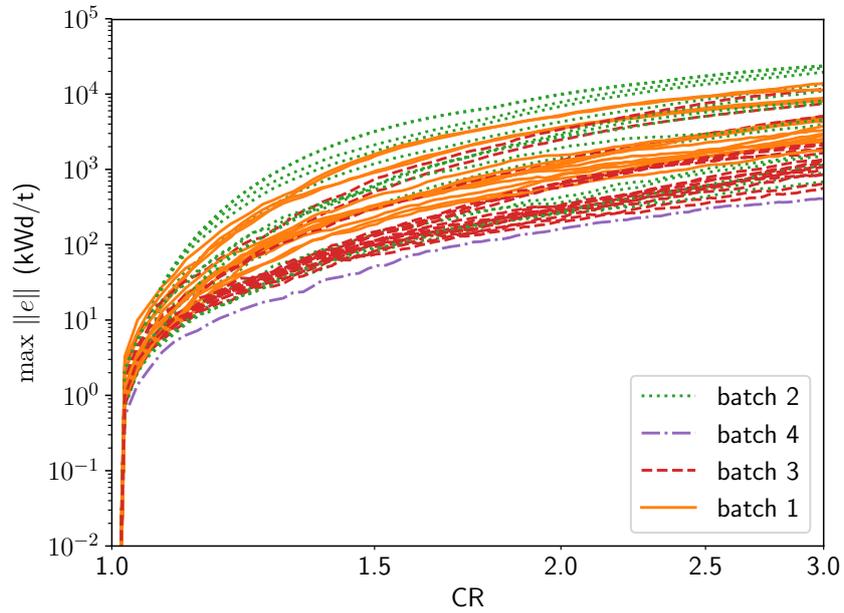
As suggested by the JPEG format, the DCT can also be applied on sub-blocks of the input image, after partitioning the whole pixel space (Wallace, 1992). This means in practice that each DCT application elaborates a local bi-dimensional signal using only a few frequencies. We have tried to reproduce the same application on our multi-dimensional arrays without remarking significant improvement. In general, the results improve with hyper-cubic samples of increasing size, i.e. 4, 8, 16, up to other powers of 2, but the unwanted large tails in the error distributions remain. The global trend is still observed when truncating the transform coefficients.

The application of the HT needs a former matricization on the given dataset, and there are apparently as many different combinations of  $i$  subscripts out of the original four as  $\sum_{i=1}^3 \binom{4}{i} = 14$ . Indeed the true number is rather the half as shown in Fig. 4, because the singular values neither change by transposing the original matrix, nor by swapping the order of the indices grouped on the same dimension of the matrix. The small differences between these corresponding curves are mainly due to the different vectors containing the average values per row of the tested matrices. According to the maximum relative errors and to the 95% quantiles plotted in the figure, the selection XYK-Z (or Z-YXK) is the best option with potential savings higher than 60%. This is due to the fact that the XY burnup shapes are quite similar along the Z axis, and redundancy is high. Briefly, we remind that X, Y, Z, K represent the different FAs, the points in time during cycle exposure, the fuel rods of a given assembly and the axial planes, respectively.

The storage of the new basis components in  $\mathbf{E}$  and of the average values in  $\vec{\mu}$  causes an additional cost with respect to the DCT. After applying the HT on a matrix of shape  $I \times J$  and keeping only  $c$  variances (and the corresponding eigenvectors), the storage reduces to  $S = Ic + Jc + I$ , as expressed by Eq. 2b. Since our goal is  $S < (IJ)$ ,  $c$  cannot exceed the threshold  $I(J - 1)/(I + J)$ . The maximum achievable CR is  $(IJ)/(2I + J)$ . The aspect ratio of the matrix affects then the performances on two aspects. The computational cost of the eigen-decomposition needed by the HT requires a small  $I$ , which limits however the maximum achievable CR. Instead,  $I > J$  can yield higher savings but with a more expensive HT computationally. The

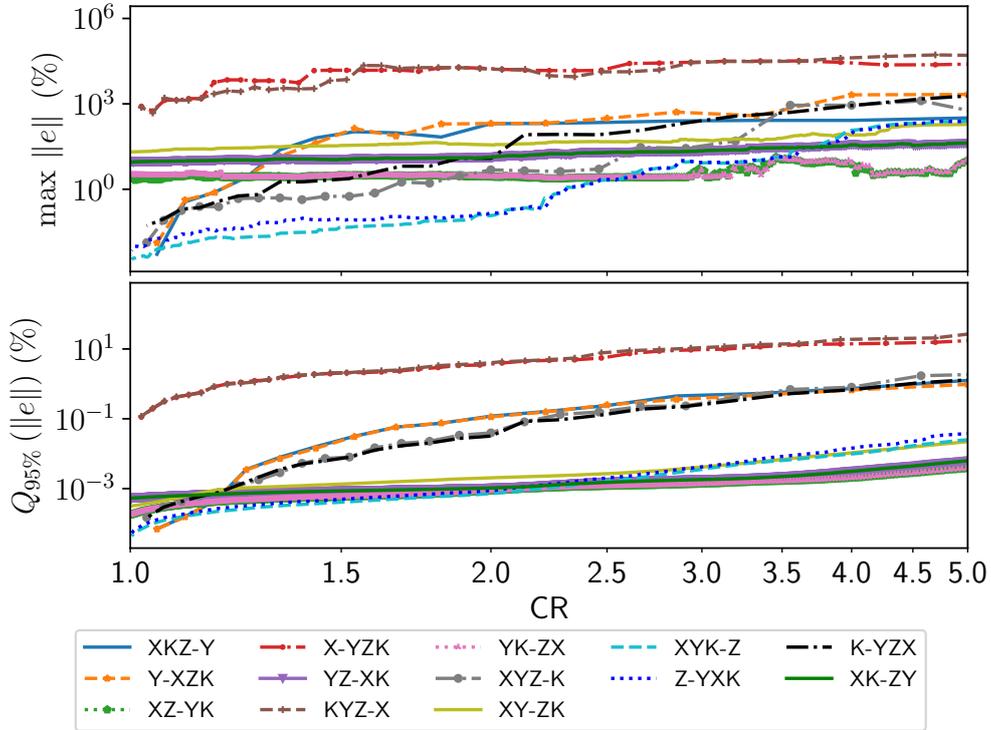


(a) Range within 5% and 95% quantiles (left) and its maximum value (right) by DCT on the whole dataset.



(b) Maximum absolute error with DCT on separate FAs.

Figure 3: Distribution of the relative percent error in absolute value after DCT-II and progressive truncation.



**Figure 4: Maximum relative error and the 95% quantile computed after HT and progressive truncation on the different matrixization options.**

choice XYK-Z uses as secondary axis the Z-axis, which is the one with more values, thus attaining potentially higher CR. Moreover, the decomposition of this matrix is rather cheap.

## 5. Conclusion

The applications of the HT and of the multi-dimensional DCT on 3D pin-by-pin burnup distributions are presented in this work. The numerical transforms are briefly recalled to support the numerical results. The results from different configurations of a PWR unit are analyzed all along a typical equilibrium cycle. The given dataset shows values spanning over ten decades at least, becoming a challenging test for compression algorithms based on numerical transforms. In addition, the accepted relative and absolute tolerances on the reconstructed values were of the order of kWd/t. These are consistent with the accuracy of nodal diffusion codes. Provided the strict tolerances, the HT showed compression ratios of about 60% with a partic-

ular matricization of the original data. Although several DCT applications were tested, none could achieve comparable performances in compression to the HT indeed.

We encourage the use of numerical transforms for compression purposes in core reactor analysis also for other 3D pin-by-pin distributions. For instance, a method to take into account the local history at pin level (Dall’Osso and Hobson, 2017) needs the pin-by-pin distribution of the ratio between the densities of  $^{239}\text{Pu}$  and  $^{238}\text{U}$ . Since nodal methods are expected to be replaced by pin-by-pin methods, the distributions of the whole set of nuclides will be stored in next generation codes. In this case compression techniques will be unavoidable.

## References

- Coppolani, P., 2012. La chaudière des réacteurs à eau sous pression. EDP Sciences. In French.
- Dall’Osso, A., Hobson, G., 2017. Modeling Local History Effects in the Nodal Code ARTEMIS<sup>TM</sup>, in: Proc. Int. Conf. on Mathematics & Computational Methods Applied to Nuclear Science and Engineering M&C2017, Jeju, Korea.
- Girieud, P., Daudin, L., Garat, C., Marotte, P., Tarlé, S., 2001. SCIENCE version 2: the most recent capabilities of the Framatome 3D nuclear code package, in: Proc. Int. Conf. on Nuclear Engineering ICONE-9, Nice Acropolis, France.
- IRSN, 2008. Le point de vue de l’IRSN sur la sûreté et la radioprotection du parc électronucléaire français en 2007. Technical Report. IRSN. RAPPORT DSR N° 271, in French.
- Koebke, K., Hetzelt, L., 1985. On the reconstruction of local homogeneous neutron flux and current distributions of light water reactors from nodal schemes. Nuclear Science and Engineering 91, 123–131.
- Lawrence, R., 1986. Progress in nodal methods for the solution of the neutron diffusion and transport equations. Progress in Nuclear Energy 17, 271–301.
- Penna, B., Tillo, T., Magli, E., Olmo, G., 2007. Transform coding techniques for lossy hyperspectral data compression. IEEE Transactions on Geoscience and Remote Sensing 45, 1408–1421.
- Shi, Y.Q., Sun, H., 1999. Image and video compression for multimedia engineering: Fundamentals, algorithms, and standards. CRC press.

Tomatis, D., 2019. Compression of assembly power form factors for core calculations. *Nuclear Science and Engineering* 193, 622–637.

Wallace, G.K., 1992. The JPEG still picture compression standard. *IEEE transactions on consumer electronics* 38, xviii–xxxiv.