

# Singular Value Decomposition of Adjoint Flux Distributions for Monte Carlo Variance Reduction \*

Elliott D. Biondo<sup>a,1</sup>, Thomas M. Evans<sup>a,1</sup>, Gregory G. Davidson<sup>a,1</sup>, Steven P. Hamilton<sup>a,1</sup>

<sup>a</sup>*Oak Ridge National Laboratory, 1 Bethel Valley Rd., Oak Ridge, TN 37831 U.S.A.*

---

## Abstract

Monte Carlo (MC) shielding calculations often use weight windows (WWs) and biased sources formed from a deterministic estimate of the adjoint flux to improve the convergence rate of tallies. This requires a significant amount of computer memory, which can limit the memory available for high-resolution tally output. A new method is proposed for reducing these memory requirements by using singular value decomposition (SVD) in linear or logarithmic space to approximate the adjoint flux. This method's performance is evaluated using the Shift and Denovo codes for streaming and diffusion base case problems, followed by problems using the Westinghouse AP1000 and the Joint European Torus. The log SVD reduced WW memory requirements by an order of magnitude in all cases without a significant performance penalty. Additionally, the linear SVD reduced biased source memory requirements by an order of magnitude, but further investigation is needed to account for

---

\*Notice: This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

\*Corresponding Author. Tel: +1 865 241-4431

*Email addresses:* [biondoed@ornl.gov](mailto:biondoed@ornl.gov) (Elliott D. Biondo), [evanstm@ornl.gov](mailto:evanstm@ornl.gov) (Thomas M. Evans), [davidsongg@ornl.gov](mailto:davidsongg@ornl.gov) (Gregory G. Davidson), [hamiltonsp@ornl.gov](mailto:hamiltonsp@ornl.gov) (Steven P. Hamilton)

<sup>1</sup>Radiation Transport Group, Reactor and Nuclear Systems Division

observed limitations.

*Keywords:* Monte Carlo radiation transport, singular value decomposition, variance reduction

---

## 1. Introduction

The design and operation of nuclear systems often requires obtaining computational estimates of radiological quantities in highly shielded regions for regulatory, environmental, and occupational safety applications. The accuracy required for these calculations necessitates the use of Monte Carlo (MC) radiation transport that—unlike deterministic radiation transport—employs a continuous representation of spatial, energy, and angular phase space. Sophisticated MC variance reduction (VR) techniques are generally required to obtain converged tally results in shielded regions in reasonable computer processor times.

It has long been known that VR can be optimized for a given detector response function via importance sampling, i.e., biasing MC transport with the adjoint flux that results from using the detector response function as an adjoint source [1]. Deterministic methods are ideally suited to obtain estimates of the adjoint flux for VR due to their ability to quickly calculate approximate solutions and because variations in the magnitude of the flux do not directly impact convergence. Biasing MC using deterministic estimates of the adjoint flux has historically been impractical due to software development limitations. Instead, slowly converging methods for obtaining the adjoint flux [2], or even biasing with respect to another function entirely [3, 4], have commonly been employed.

The Consistent Adjoint Driven Importance Sampling (CADIS) method [5] uses a deterministic estimate of the adjoint flux to produce MC weight windows (WWs) [6] and corresponding biased sources. Forward-Weighted (FW)-CADIS [7] is a variant in which transport is optimized with respect to multiple detector response functions simultaneously using an additional deterministic forward (i.e., not adjoint) transport calculation. In recent years CADIS and FW-CADIS have been implemented in a variety of production-level software suites, including Exnihilo [8], ADVANTG [9], MC21 [10], and MAVRIC [11], and they have gained traction in the nuclear analysis community.

The performance of CADIS depends on how closely the deterministic

estimate of the adjoint flux matches the true importance function. When using CADIS, users must select the discretization scheme for the deterministic transport step, including the spatial mesh and the energy group structure via supplied nuclear cross sections. This ultimately determines the size of the weight windows and biased source meshes that must be stored in memory for the duration of the MC simulation. This creates a trade-off: increasing the resolution of the deterministic transport mesh and energy group structure may result in improved tally convergence rates, but it will decrease the amount of memory available for tally results.

This is problematic for high-fidelity analysis; due to memory limitations, it is not uncommon for analysts to perform the same MC simulation multiple times with different tallies to obtain all desired tally results. This may be due to complex geometry descriptions (especially for fusion applications [12, 13]), a significant number of unique material compositions (especially for depletion problems [14, 15]), or mesh tallies that span large spatial regions (global VR problems [16]). This memory issue is exacerbated by the fact that many MC codes are parallelized with Message Passing Interface (MPI) using domain replication (DR), so each MPI process stores its own full copy of the problem data in memory. Furthermore, trends in high-performance computing are driving the development of MC codes on advanced compute architectures such as graphical processing units (GPUs) [15], which typically have less memory than CPU-based computing resources.

The memory requirements of the WWs and biased sources can be reduced by coarsening the spatial mesh or energy group structure or by using domain decomposition (DD). Spatial coarsening can be labor intensive and limited in its possible extent, especially in the presence of fine geometric features such as streaming channels. Energy coarsening is often impractical because deterministic transport codes typically use pre-built cross section libraries, and creating custom cross section libraries is a difficult and time-consuming process. In addition, if the spatial mesh becomes too coarse relative to the particle mean free path, or if the energy groups become too coarse relative to the shape of the cross section and energy loss per collision, then VR performance may deteriorate significantly. Unlike DR, DD involves dividing the phase space of the problem into subdomains, each of which is stored by a single MPI process. While DD can significantly reduce memory requirements, it can also negatively impact the parallel scaling behavior of MC due to load balancing issues [17]. Nonetheless, DD can be used in concert with the new method proposed herein.

This work proposes a new method for reducing the memory requirements of WWs and biased sources without generating a new spatial mesh or cross section library, and without sacrificing performance. This method uses the singular value decomposition (SVD) matrix factorization technique in linear or logarithmic space to produce approximations of the adjoint flux. This method has been implemented in the Exnihilo software suite within SCALE [11] using the Shift MC code, the Denovo 3D  $S_N/SP_N$  code [18, 19], and the Omnibus front end [20]. Preliminary analysis has been performed to demonstrate this technique for streaming and diffusion base cases, as well as two realistic problems: the Westinghouse AP1000 pressurized water reactor (PWR) [21], and the Joint European Torus (JET) experimental fusion reactor [22]. It is shown that this method can reduce the memory requirements of WWs—and in some cases biased sources—by an order of magnitude without significantly impacting the convergence rates of tallies.

The remainder of this document is organized as follows. Sections 2 and 3 summarize CADIS and SVD, respectively. Section 4 describes how the SVD is used in this work to reduce storage requirements for WWs and biased sources. Section 5 describes the design choices for the implementation of this method in the Exnihilo software suite. Section 6 explores the performance of this method. Finally, Section 7 provides concluding remarks and planned future work.

## 2. The CADIS Method

With the CADIS method, a deterministic estimate of the adjoint flux is used to produce optimal WWs and biased sources with respect to a detector response function  $\sigma_d(\vec{r}, E)$ . This is typically a single tally or a group of spatially clustered tallies. With FW-CADIS, an additional deterministic estimate of the forward flux is obtained to form a detector response function that can be used to optimize transport for multiple tallies simultaneously. Both methods are described in detail in the literature [5, 7]. For simplicity, this work focuses on CADIS, with the expectation that any findings will also apply to FW-CADIS.

To carry out the CADIS procedure, the detector response function is first discretized onto the deterministic mesh and energy group structure. The discrete form of the detector response function is then assigned to be the

adjoint source,

$$q_c^{\dagger g} = \sigma_c^g, \quad (1)$$

where  $c$  denotes the mesh cell index, and  $g$  is the energy group index. A deterministic adjoint transport calculation is then performed using this source. The resulting adjoint flux provides an estimate of the importance of phase space regions to the detector response. The detector response is defined as the volume-integrated product between the forward flux and the detector response function

$$R = \sum_c \sum_g \phi_c^g \sigma_c^g V_c, \quad (2)$$

where  $V_c$  is the volume of mesh cell  $c$ . By reciprocity, an estimate of the detector response can be formulated in terms of the estimated adjoint flux and the forward source [23]

$$R = \sum_c \sum_g \phi_c^{\dagger g} q_c^g V_c. \quad (3)$$

The optimal biased source is then given by

$$\hat{q}_c^g = \frac{\phi_c^{\dagger g} q_c^g}{R}. \quad (4)$$

To play a fair game, particles born from this source must have the modified weight ( $w$ ) given by

$$w_c^g = \frac{R}{\phi_c^{\dagger g}}. \quad (5)$$

The corresponding WW lower bounds ( $w_l$ ) are

$$w_{l,c}^g = \frac{R}{\left(\frac{1+\beta}{2}\right) \phi_c^{\dagger g}}, \quad (6)$$

where  $\beta$  is the ratio of WW upper bound to lower bound (typically 5). The factor of  $(1 + \beta)/2$  ensures that the center of each WW is  $w_c^g$  as given in Eq. (5). The process of discretizing the MC geometry and source, carrying out deterministic transport, calculating the CADIS VR parameters, then carrying out MC transport with VR has been automated in the industry-standard software suites mentioned in Section 1. Geometry discretization

is generally performed using a ray tracing algorithm, and point sampling is used to discretize sources in space and energy [8].

The performance of CADIS is affected by how well the deterministic estimate of the adjoint flux matches the true importance function and the degree to which the true adjoint is angularly dependent. The latter is significant because industry-standard MC codes do not typically support angularly dependent WWs. As a result, biased sampling is only achieved in space and energy, although methods have been proposed to account for the angular dependence of the adjoint flux without explicitly performing angular biasing [24]. The performance of CADIS can be quantified by the speedup,

$$\text{speedup} = \frac{\text{FOM}_{\text{CADIS}}}{\text{FOM}_{\text{analog}}}, \quad (7)$$

where FOM is standard MC figure of merit in terms of the relative error in the tally result ( $R$ ) and the computer processor time ( $t_{\text{proc}}$ ),

$$\text{FOM} = \frac{1}{R^2 t_{\text{proc}}}. \quad (8)$$

For the purpose of this work, *analog* is taken to mean MC transport in the absence of VR, aside from implicit capture.

### 3. Singular Value Decomposition

SVD is a prominent matrix factorization technique closely related to principle component analysis and proper orthogonal decomposition. SVD has a wide variety of applications, including image compression [25], bioinformatics [26], and signal processing [27]. In the nuclear engineering field, SVD has previously been applied to reactor noise analysis [28], and also the approximation of Jacobian matrices of sensitivity coefficients for uncertainty quantification [29].

SVD can be applied to any arbitrary two-dimensional matrix  $\mathbf{A}$ . Two assumptions are made without loss of generality for the purpose of simplifying this discussion and are not necessary for the implementation of this method as described in Section 5. It is first assumed that  $\mathbf{A}$  is a real-valued  $m \times n$  matrix with  $m > n$ . The second assumption pertains to the matrix rank of  $\mathbf{A}$ . The rank  $r$  of a matrix refers to the size of the vector space that is spanned by the rows or columns of the matrix. A matrix is said to be full rank if

$r = \min(m, n)$ . Here, matrix  $\mathbf{A}$  is assumed to be full rank with  $r = n$ . With these assumptions, the thin SVD<sup>2</sup> of  $\mathbf{A}$  is given by the following [30]:

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = [\mathbf{u}_1 \quad \mathbf{u}_2 \quad \dots \quad \mathbf{u}_n] \begin{bmatrix} \sigma_1 & 0 & \dots & 0 \\ 0 & \sigma_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_n \end{bmatrix} \begin{bmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \vdots \\ \mathbf{v}_n^T \end{bmatrix}. \quad (9)$$

In this equation,  $\mathbf{\Sigma}$  is a diagonal matrix of monotonically decreasing values, known as singular values.  $\mathbf{U}$  is a matrix comprised of the orthogonal left-singular vectors  $\{\mathbf{u}\}$ , each of length  $m$ .  $\mathbf{V}$  is a unitary matrix comprised of the right-singular vectors  $\{\mathbf{v}\}$ , each of length  $n$ . Here,  $T$  denotes the matrix transpose.

One application of this decomposition is the ability to form low-rank approximations of  $\mathbf{A}$ , i.e., approximations that have some rank  $k < n$ . If an approximation  $\mathbf{A}_k$  can be found with some significantly reduced rank  $k$ , then the storage requirements of  $\mathbf{A}$  can be diminished by only storing the linearly independent components of  $\mathbf{A}$ . The Eckart-Young theorem [31] states that the optimal rank  $k$  approximation of  $\mathbf{A}$  can be formed via a truncated SVD, from the singular values and left- and right-singular vectors of  $\mathbf{A}$ ,

$$\mathbf{A}_k = \sum_{l=1}^k \sigma_l \mathbf{u}_l \mathbf{v}_l^T, \quad (10)$$

where individual values of the approximation can be calculated via

$$a_{k,ij} = \sum_{l=1}^k \sigma_l u_{il} v_{jl}. \quad (11)$$

With this notation,  $a_{k,ij}$  denotes the value in the  $i^{\text{th}}$  row and  $j^{\text{th}}$  column in the  $\mathbf{A}_k$  matrix. Note that when  $k = r = n$ , the approximation is exact. This approximation minimizes the Frobenius norm of the difference between  $\mathbf{A}$

---

<sup>2</sup>This work uses the so-called *thin SVD* as opposed to the *full SVD*, in which constituent matrices have larger storage requirements but do not contain any additional information.

and  $\mathbf{A}_k$ , which is defined as

$$\|\mathbf{A} - \mathbf{A}_k\|_F = \sqrt{\sum_i \sum_j |a_{ij} - a_{k,ij}|^2}. \quad (12)$$

This quantity can be related to the singular values of  $\mathbf{A}$  via,

$$\|\mathbf{A} - \mathbf{A}_k\|_F = \sqrt{\sum_{i=k+1}^n \sigma_i^2}. \quad (13)$$

From this, it follows that the relative error in a rank  $k$  approximation is given by

$$R_k = \frac{\sqrt{\sum_{i=k+1}^n \sigma_i^2}}{\sqrt{\sum_{i=1}^k \sigma_i^2}}. \quad (14)$$

This can be used to discern the minimum rank necessary to achieve a target relative error, which may be predictive of some performance metric of  $\mathbf{A}_k$ .

Based on Eq. (10), it is clear that only the first  $k$  singular values and the first  $k$  right and left singular vectors are necessary to calculate any value in  $\mathbf{A}_k$ . Computer memory can be saved by storing only this data, and individual values of  $\mathbf{A}_k$  can be recalculated as needed via Eq. (11). To quantify the amount of memory saved, the memory saving factor is defined as follows:

$$M = \frac{\text{memory with full rank}}{\text{memory with rank } k} = \frac{mn}{mk + k + nk}. \quad (15)$$

An example of SVD is shown in Fig. 1. Figure 1d shows a  $1,333 \times 1,000$  pixel grayscale image, that can be represented as a matrix of floating point values over the range  $[0, 1]$ . Figures 1a to 1c show ranks 1, 10, and 100 approximations of this matrix. Figure 2 shows the relative error of the approximations as a function of rank. The plot shows that the rank 100 approximation has a relative error of 12.3%, and Eq. (15) indicates that this approximation requires 5.71 times less memory. Visually, the rank 100 image closely matches the full rank image, so this memory savings could be useful for some applications.




Figure 1: SVD reconstructions of an image represented as a  $1,333 \times 1,000$  matrix of grayscale values between 0 and 1. The photograph shows a waterfall in Stephens Gap Cave in Jackson County, Alabama, and is an original work by the author.




Figure 2: Relative error as a function of rank for the matrix shown in Fig. 1d.

## 4. Application of SVD to CADIS

In the example in Fig. 1, SVD was used to approximate a 2D matrix, with the rows and columns both representing spatial dimensions. There are a multitude of ways this same principle can be used to approximate CADIS/FW-CADIS adjoint flux distributions. This section introduces two possible methods: one which performs SVD in linear space, and another which performs SVD in log space. The application of these two methods to WWs and biased sources is then discussed.

### 4.1. SVD in Linear Space

With the first method, the adjoint flux is converted into a 2D matrix with  $m$  rows representing mesh cells and  $n$  columns representing energy groups, as follows:

$$\mathbf{A} = \begin{bmatrix} \phi_{c=0}^{\dagger g=0} & \phi_{c=0}^{\dagger g=1} & \cdots & \phi_{c=0}^{\dagger g=N_g-1} \\ \phi_{c=1}^{\dagger g=0} & \phi_{c=1}^{\dagger g=1} & \cdots & \phi_{c=1}^{\dagger g=N_g-1} \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{c=N_c-1}^{\dagger g=0} & \phi_{c=N_c-1}^{\dagger g=1} & \cdots & \phi_{c=N_c-1}^{\dagger g=N_g-1} \end{bmatrix}. \quad (16)$$

With this configuration, the resulting left- and right- singular vectors of the SVD can be thought of as characteristic spatial distributions and energy spectra, respectively.

Though the adjoint flux is known to be the optimal biasing function, it is not clear which norm should be used to judge the quality of low rank approximations relative to the full rank adjoint flux, which in this case is a deterministic estimate. By using the matrix construction shown in Eq. (16), the SVD optimizes rank  $k$  approximations with respect to the Frobenius norm in linear space, as shown in Eq. (12). Because this norm uses the arithmetic difference between the adjoint flux and the rank  $k$  approximations, phase space regions in which the adjoint flux is higher—near the detector—are weighted more heavily than regions of low adjoint flux (i.e., near the source). This is significant because for typical shielding problems, adjoint flux distributions may span 5–10 orders of magnitude or more through important phase space regions. Optimizing low-rank approximations of the adjoint flux in a manner

that provides more equal weighting across all phase space regions may result in better performance.

#### 4.2. SVD in Logarithmic Space

This second method involves forming an  $\mathbf{A}$  matrix by first taking the log of each value,

$$\mathbf{A} = \begin{bmatrix} \log(\phi_{c=0}^{\dagger g=0}) & \log(\phi_{c=0}^{\dagger g=1}) & \dots & \log(\phi_{c=0}^{\dagger g=N_g-1}) \\ \log(\phi_{c=1}^{\dagger g=0}) & \log(\phi_{c=1}^{\dagger g=1}) & \dots & \log(\phi_{c=1}^{\dagger g=N_g-1}) \\ \vdots & \vdots & \ddots & \vdots \\ \log(\phi_{c=N_c-1}^{\dagger g=0}) & \log(\phi_{c=N_c-1}^{\dagger g=1}) & \dots & \log(\phi_{c=N_c-1}^{\dagger g=N_g-1}) \end{bmatrix}. \quad (17)$$

The resultant SVD will minimize the Frobenius norm in log space, as given by

$$\|\mathbf{A} - \mathbf{A}_k\|_F = \sqrt{\sum_c \sum_g \left| \log(\phi_c^{\dagger g}) - \log(\phi_{k,c}^{\dagger g}) \right|^2}. \quad (18)$$

With this method, the relative difference between the adjoint flux and the rank  $k$  approximation is weighted more uniformly across phase space. This method has the same memory requirements as the linear method.

#### 4.3. Weight Windows

With the WW game, the WW lower bounds (and corresponding upper bounds) within a region of phase space must be determined many times per particle history, typically at some combination of boundary crossings, collisions, WW mesh crossings, or traversals of a certain number of particle mean free paths. WW lower bounds can be calculated via Eq. (6), which requires the adjoint flux within a specific region of phase space.

When a reduced rank estimate of the adjoint flux is being used, individual values of the adjoint flux must be calculated on demand. When SVD in linear space is being used, this is done using Eq. (11), duplicated here:

$$\phi_{k,c}^{\dagger g} = a_{k,ij} = \sum_{l=1}^k \sigma_l u_{il} v_{jl}. \quad (19)$$

For SVD in log space, the adjoint flux is recovered via an exponential evaluation:

$$\phi_{k,c}^{\dagger g} = \exp(a_{k,ij}) = \exp\left(\sum_{l=1}^k \sigma_l u_{il} v_{jl}\right). \quad (20)$$

In both of these cases, additional floating point operations must be performed to recover the adjoint flux, which is normally simply retrieved from memory. This potential performance penalty is accounted for in the MC FOM. Since it is expected that the number of spatial regions greatly exceeds the number of energy groups, the memory reduction factor for the calculation of WWs can be simplified to

$$M = \frac{mn}{mk + k + nk} \approx \frac{n}{k}. \quad (21)$$

#### 4.4. Biased Sources

For source biasing, using a reduced-rank estimate of the adjoint flux is not as straightforward. Sources can be subdivided into two categories: separable in space/energy, and non-separable in space/energy. For separable sources, the source in any mesh cell and energy group can be represented as

$$q_c^g = q_c q^g. \quad (22)$$

In other words, the source spectrum is the same throughout all space. Separable sources are used in a wide variety of nuclear analysis scenarios, including fission sources for standard light water reactor excore problems, and simple models of fusion plasma sources for tokamak problems. In contrast, with non-separable sources the source spectrum varies spatially, as is the case for scenarios such as shutdown dose rate analysis for fusion energy systems [32].

To sample a particle source probability density function (PDF), the PDF must be converted to a cumulative distribution function (CDF), which must remain in memory for the entirety of a simulation. For non-separable sources, the size of this CDF will be  $N_c \times N_g$ , where  $N_c$  is the number of mesh cells, and  $N_g$  is the number of energy groups. When using source biasing, the biased version of a non-separable source also must be of size  $N_c \times N_g$ , so using a low-rank approximation of the adjoint flux does not result in memory reduction.

For separable sources, separate CDFs can be created in space and in energy. The sizes of these CDFs are  $N_c$  and  $N_g$  for space and energy, respectively.

With the standard CADIS method, the adjoint flux is of size  $N_c \times N_g$ , so the biased CDF will also be of size  $N_c \times N_g$ , even if the source is separable. Using a reduced rank adjoint provides one way of applying the CADIS method to separable sources without creating an  $N_c \times N_g$  CDF, resulting in a factor of  $N_g$  memory savings.

Consider the case in which the adjoint flux can be decomposed in space and energy, as follows:

$$\phi_c^\dagger{}^g = \phi_c^\dagger \phi^\dagger{}^g. \quad (23)$$

Using this adjoint flux, the response can also be decomposed in space and energy:

$$\begin{aligned} R &= \sum_c \sum_g \phi_c^\dagger \phi^\dagger{}^g q_c q^g V_c \\ &= \left( \sum_c \phi_c^\dagger q_c V_c \right) \left( \sum_g \phi^\dagger{}^g q^g \right) \\ &= R_C R_G, \end{aligned} \quad (24)$$

which is also true of the biased source:

$$\begin{aligned} \hat{q}_c^g &= \frac{\phi_c^\dagger \phi^\dagger{}^g q_c q^g}{R_C R_G} \\ &= \left( \frac{\phi_c^\dagger q_c}{R_C} \right) \left( \frac{\phi^\dagger{}^g q^g}{R_G} \right) \\ &= \hat{q}_c \hat{q}_g. \end{aligned} \quad (25)$$

Because this biased source is separable in space and energy, it can be sampled with two separate CDFs as is done for the analog separable source.

Equation (19) shows that a rank 1 approximation can be decomposed in the form of Eq. (23), with

$$\phi_c^\dagger = \sigma_1 u_{i1}, \quad (26)$$

and

$$\phi^\dagger{}^g = v_{j1}, \quad (27)$$

noting that the  $\sigma_1$  could be included as part of either the spatial or energy term. It is also apparent from Eqs. (19) and (20) that approximations using the linear form with  $k > 1$ , as well as all approximations using the log form, cannot be represented in the separable form of Eq. (23). As a result, the memory savings factor for biased separable sources is given by

$$M \approx \begin{cases} n, & \text{linear, } k = 1 \\ 1, & \text{otherwise.} \end{cases} \quad (28)$$

#### 4.5. Biased Separable Sources with $k > 1$

Though the linear SVD cases with  $k > 1$  cannot be used to form a single separable source in the form of Eq. 22, one method for reducing memory requirements in these cases is proposed here. This method is left for future work, as it requires more significant modifications to the MC source sampling procedure.

For a rank  $k$  linear SVD case, independent separable sources can be created from each of the  $k$  sets of singular values and left- and right-singular vectors. Each one of these sources can then be biased in the same fashion as the linear rank 1 case (as described in Section 4.4), resulting in  $k$  biased separable sources. To sample a source particle, the rank is first sampled using a PDF formed from the  $k$  singular values. The biased separable source corresponding to this rank is then sampled. This strategy yields a memory saving factor of  $M = n/k$ .

## 5. Implementation in Exnihilo

To test the performance of the methods described in Section 4, modifications were made to the Exnihilo code base. The Omnibus front end application allows for CADIS to be run automatically using the Shift MC code and the Denovo 3D  $S_N/SP_N$  code. This is done using a single Omnibus input file that specifies the parameters for both Shift and Denovo. Omnibus discretizes the geometry and detector response function onto a supplied superimposed mesh to set up a Denovo problem. Omnibus then runs an adjoint transport problem in Denovo and passes the resulting adjoint flux distribution to Shift. Shift then runs MC transport with CADIS WWs and source biasing.

An additional step is added to this process in which Denovo performs an SVD of the adjoint flux in linear or log space using the matrix constructions described in Sections 4.1 and 4.2, respectively. In both cases, only rank

1 approximations are currently supported, although in principle, Exnihilo could be further modified to support approximations of any rank.<sup>3</sup> As a result, only the first singular value and the corresponding left- and right-singular vectors are required. These are obtained using an iterative method as described in the following subsection. Shift was also modified to use these rank 1 approximations of the adjoint flux to create biased sources (which are separable when the unbiased source is separable) and also to calculate WWs on demand.

### 5.1. Calculating the SVD

For typical CADIS shielding applications, the deterministic transport problem may have  $10^6 - 10^8$  spatial elements and 10–100 energy groups. The size of the resulting  $\mathbf{A}$  matrix makes direct methods for computing the SVD impractical. In addition, since Denovo uses domain decomposition, no single process has a full copy of the  $\mathbf{A}$  matrix, which would further complicate the use of direct methods. Instead, the first singular value and corresponding left- and right-singular vectors are obtained by reformulating the SVD in terms of a standard eigenvalue problem. Because  $\mathbf{V}$  is unitary, Eq. (9) can be rewritten as

$$\mathbf{A}\mathbf{V} = \mathbf{U}\mathbf{\Sigma}, \quad (29)$$

which implies

$$\mathbf{A}\mathbf{v}_i = \sigma_i \mathbf{u}_i. \quad (30)$$

Next, consider the transpose of the SVD,

$$\mathbf{A}^T = (\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T)^T = \mathbf{V}\mathbf{\Sigma}\mathbf{U}^T. \quad (31)$$

This yields

$$\mathbf{A}^T\mathbf{U} = \mathbf{V}\mathbf{\Sigma}, \quad (32)$$

which implies

$$\mathbf{A}^T\mathbf{u}_i = \sigma_i \mathbf{v}_i. \quad (33)$$

---

<sup>3</sup>Note that in this work, the performance of approximations with ranks greater than 1 were also tested; the method for doing so is described in Section 6.

Equations (30) and (33) can then be combined to form the following standard eigenvalue problem

$$\begin{pmatrix} 0 & \mathbf{A} \\ \mathbf{A}^T & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u}_i \\ \mathbf{v}_i \end{pmatrix} = \sigma_i \begin{pmatrix} \mathbf{u}_i \\ \mathbf{v}_i \end{pmatrix}, \quad (34)$$

which is in the form

$$\mathbf{L}\mathbf{x} = \lambda\mathbf{x}, \quad (35)$$

with

$$\begin{aligned} \mathbf{L} &= \begin{pmatrix} 0 & \mathbf{A} \\ \mathbf{A}^T & 0 \end{pmatrix} \\ \mathbf{x} &= (\mathbf{u}_i \quad \mathbf{v}_i)^T \\ \lambda &= \sigma_i. \end{aligned} \quad (36)$$

A standard eigensolver can then be used to solve for the largest eigenvalue and corresponding eigenvector. Since  $\sigma_i$  values decrease monotonically in  $\Sigma$ , this pair will provide  $\sigma_1$ ,  $\mathbf{u}_1$ , and  $\mathbf{v}_1$ , as desired. In principle, this could be accomplished using a symmetric eigensolver such as the Lanczos algorithm. However, in this case, Arnoldi iteration was used due to its availability within the Anasazi package within Trilinos [33].

Arnoldi iteration requires successive evaluation of  $\mathbf{L}\mathbf{x} = \mathbf{x}'$ , where  $\mathbf{x}$  and  $\mathbf{x}'$  are iteration vectors. This is done in parallel with  $\mathbf{L}$ ,  $\mathbf{x}$ , and  $\mathbf{x}'$  decomposed across multiple processor domains, as shown in Fig. 3. When multiplying the first  $m$  rows of  $\mathbf{L}$  by  $\mathbf{x}$ , only the  $\mathbf{v}_i$  portion of  $\mathbf{x}$  is required. This portion is stored on domain 0 by convention and is communicated to other processes via an MPI broadcast. When multiplying the last  $n$  rows of  $\mathbf{L}$  by  $\mathbf{x}$ , each process only operates on the local portion of the  $\mathbf{u}_i$  vector, and values of  $\mathbf{x}'$  are then calculated using an MPI global sum across all processes.

Though this procedure only calculates the first singular value and corresponding left- and right-singular vectors, it could be modified to calculate these qualities for SVD reconstructions of any rank. This will be necessary to add support within Exnihilo for linear and log SVD approximations with  $k > 1$ .




Figure 3: Parallel decomposition of the  $L$  matrix given in Eq. (36). Colors represent how data is divided and stored by compute processes, and gray regions represent the block zeros. The  $\mathbf{u}_i$  portion of  $\mathbf{x}$  is divided evenly across all processes, and the smaller  $\mathbf{v}_i$  portion of  $\mathbf{x}$  is assigned to domain 0 by convention.

## 6. Performance

In Section 4, two methods were proposed for reducing the memory requirements of WWs and biased sources using SVD in linear and log space. In this section, a description is presented of how the performance of these methods was tested for streaming and diffusion base cases, a typical fission problem, and a typical fusion problem. This was done by finding the speedup (as defined in Eq. (7)) as a function of the rank of the approximation. For the rank 1 approximation, the SVD and subsequent source biasing and runtime calculation of WWs were performed using the newly implemented code described in Section 5. This method is referred to as *runtime reconstruction*.

Approximations with ranks greater than 1 could in principle be implemented in Exnihilo, but they are not currently supported. To test these approximations, an adjoint transport calculation was first run with Denovo, and the resultant adjoint flux distribution was written to disk. A Python script was then used to perform the SVD and generate approximations over all ranks. These approximations were then used to form reconstructions of the original adjoint flux. With these reconstructions, the rank  $k$  approximation for each value of the adjoint flux distribution is stored explicitly. The result

is an approximate adjoint flux distribution with the same memory storage requirements as the original distribution. These distributions are written to disk and then read into Shift to form CADIS WWs and biased sources. This method will be referred to as *preprocessed reconstruction*. Although this method is useful for finding the relationship between speedup and rank, it does not have any practical application, as it decreases the quality of the VR without any memory savings since the WW storage is not exploiting the low rank approximation.

All experimentation was performed using the Summit supercomputer [34] at Oak Ridge National Laboratory. Each node on this machine has two 22-core IBM Power9 processors and 512 GB of RAM, as well as 6 NVIDIA Tesla GPUs (which were not used for this work).

### 6.1. Base Case

As a proof of principle, two base case problems with nearly the same detector response were first tested: one dominated by neutron streaming, and one dominated by diffusion. This was done to determine if the dominant transport mode affects the performance of VR parameters compressed with SVD, as streaming and diffusion problems generally have different spectral shapes through important spatial regions. The geometries for these problems are shown in Figure 4. The streaming geometry is based on the Kobayashi benchmark [35] and consists of a  $60 \times 100 \times 60$  cm paraffin box of density  $0.93 \text{ g/cm}^3$  with a dog-leg streaming channel containing dry air that is  $10 \times 10$  cm in cross section. At the beginning and ending of the streaming channels,  $10 \times 10 \times 10$  cm dry air regions represent the source and tally regions, respectively. The source has a uniform probability over the energy range  $[10^{-3} \text{ eV}, 10^5 \text{ eV}]$ .

The diffusion case is exactly the same as the streaming case, except the streaming channel, which does not include the source and tally regions, is filled with paraffin, and all of the paraffin in the problem has a density that is reduced by a factor of 6. As a result, the tally responses for the two cases were equal within 5%, so both problems should be nearly equal in difficulty and should converge in analog at similar rates.

For each geometry, an adjoint Denovo calculation was performed using the detector as a source with a uniform spatial mesh consisting of  $5 \times 5 \times 5$  cm mesh cells, 28-group cross sections [11], and an  $S_{16}$  level symmetric quadruple range quadrature set. These parameters, as well as those used in subsequent problems, are representative of the typical resolution used for MC VR. The

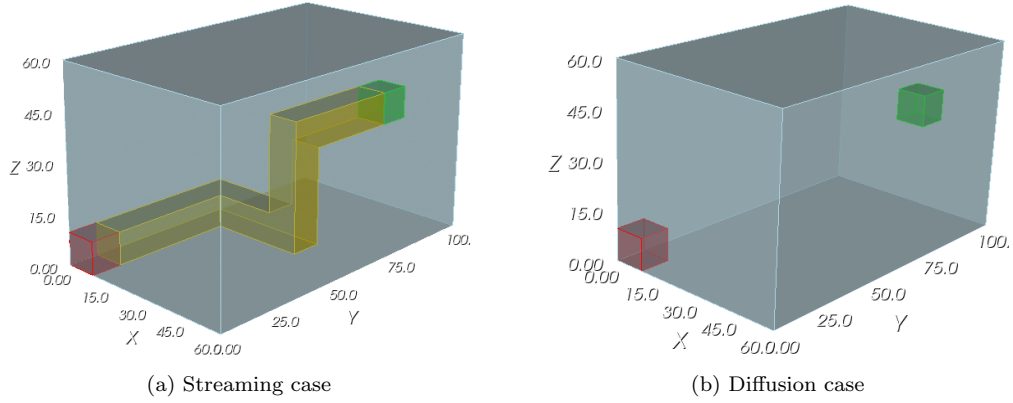


Figure 4: Geometries used for preliminary analysis.

resulting adjoint flux was then decomposed with SVD in both linear and log space and then was reconstructed with ranks 1–28. The relative error from these reconstructions is plotted in Fig. 5. It appears that the log SVD provides a better fit for the lowest rank reconstructions. This is apparent in the adjoint flux distributions for the streaming case shown in Fig. 6a for the lowest energy group,  $10^{-5} - 0.01$  eV. In these plots, the log version of the rank 1 reconstruction visually appears to match the full rank adjoint flux, whereas the linear version overestimates the adjoint flux in the streaming channel by an order of magnitude.

The columns in the  $\mathbf{V}$  matrix can be thought of as characteristic energy spectra, the linear combination of which can be used to exactly describe the spectra at any spatial location. Figure 7 shows the first 4 columns in  $\mathbf{V}$  for the linear version of the SVD. In both the streaming and diffusion cases, the first column provides a relatively flat energy spectrum, indicating that the best rank 1 approximation to the adjoint flux essentially eliminates energy biasing. The second column in  $\mathbf{V}$  is similar for both problems, but the third and fourth columns differ significantly. These columns appear to apply corrections only for high and low energies in the streaming cases, but more uniformly across all energies in the diffusion case. This may be due to the fact that the diffusion case is expected to provide relatively smooth moderation as neutrons emanate away from the source. The columns of  $\mathbf{V}$  for the log version of the SVD have a similar trend, but do not have an obvious physical interpretation.

Shift was then run using the preprocessed reconstructions, the rank 1




Figure 5: Relative error as a function of rank for SVD reconstruction of the adjoint flux.

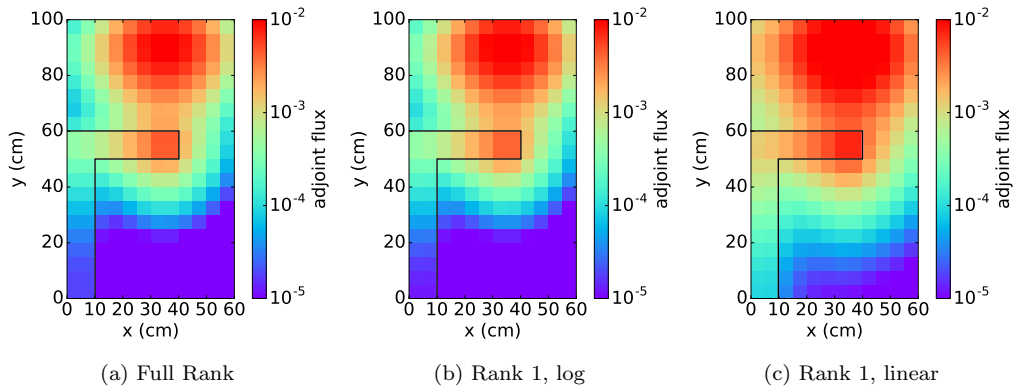


Figure 6: Adjoint flux distribution for the streaming case at  $z = 0$  for the  $10^{-5} - 0.01$  eV energy group.




Figure 7: Columns in  $\mathbf{V}$  for the linear SVD, which represent characteristic spectra.

runtime reconstructions, and the normal CADIS procedure. Ten trials with different random number seeds were run for each case, with  $5 \cdot 10^7$  histories per trial in the streaming case and  $10^8$  histories per trial in the diffusion case. Additionally, trials were run in analog to calculate the speedup, with  $2 \cdot 10^8$  histories per trial in both the streaming and diffusion cases. Four Summit nodes were used for all trials, with 40 processes per node.

Figure 8 shows speedups as a function of rank for streaming and diffusion cases. In these plots, error bars represent the standard deviation of the speedup from the 10 trials executed with VR (i.e., the uncertainty in the FOM for the analog run is not taken into account). A much larger speedup is attained in the diffusion case. This is because there is a strong angular dependence in the importance function in the streaming case, and no angular biasing is being performed. In both the streaming and diffusion cases, the log version of the SVD provides a nearly constant speedup across all ranks. The rank 28 preprocessed reconstruction has the same speedup as the CADIS base case because these two methods use the exact same adjoint flux. The log version of the rank 1 runtime reconstruction has nearly the same speedup as the preprocessed reconstruction in the streaming case, but the runtime reconstruction is slightly worse in the diffusion case. This discrepancy is accounted for by the extra computational cost associated with the runtime evaluation of the exponential function required to calculate the adjoint flux.

For the linear SVD trials, the speedup does not seem to stabilize to the full rank value until rank 13 for the streaming case and rank 8 for the diffusion case. Though the speedups in the diffusion case are all less than or equal to the full rank speedup, in the streaming case, several ranks provide a higher speedup than the full rank speedup. This is only possible because the the full rank cases are using an approximation of the adjoint flux, and angular biasing is not employed. It makes sense that this effect occurs in the streaming problem, as the true importance function should have a strong angular dependence in this case. By chance, these low rank approximations outperformed the full rank adjoint, but this is not systematic, as several low rank approximations (e.g., ranks 2 and 3) yield decidedly poor performance.

For these problems it appears that the rank 1 approximation of the log version of the SVD provides the best option. In both cases it provides a factor of  $M = 28$  reduction in memory for the weight windows, with no appreciable decrease in performance. Because the log SVD is used, the biased source is the same size as the CADIS base case. Since the source occupies a small spatial region, this does not present a problem. The ensuing problems consider cases in which source biasing is more important, with sources spanning significant spatial regions.

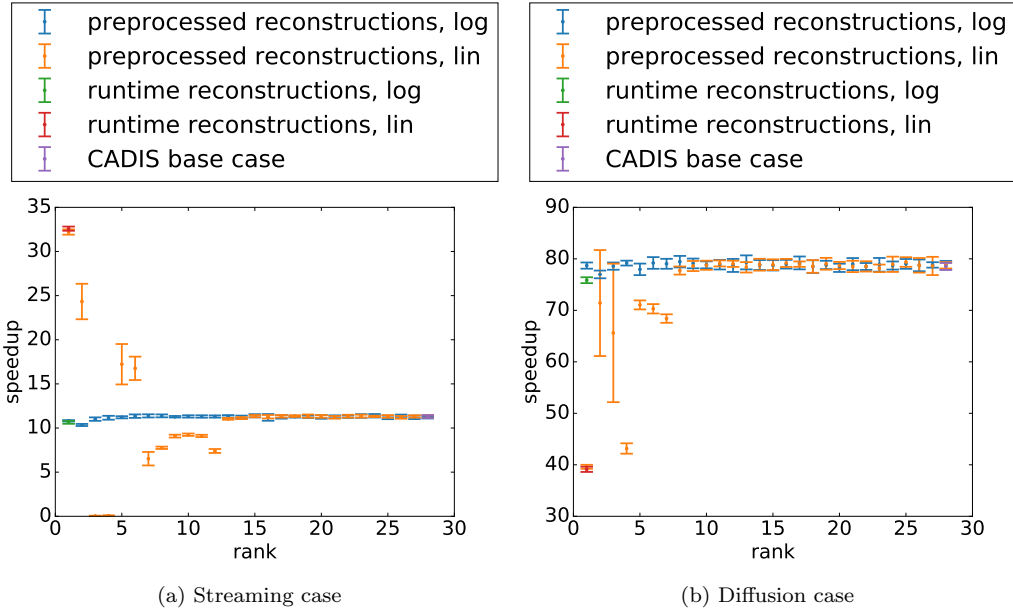


Figure 8: Speedup as a function of rank. Error bars show the standard deviation of the speedup over 10 trials using different random number generator seeds.

## 6.2. Westinghouse AP1000

After confirming that SVD can be used to compress VR parameters for the streaming and diffusion base case problems as described in Section 6.1, the same analysis was applied to production-level problems, the first being a Westinghouse AP1000 excore problem. The geometry for this problem, which is a quarter core with reflecting boundaries, is shown in Fig. 9. The reactor core portion of this model (i.e., the fuel assemblies and surrounding core baffle) was originally created as a test stand for the VERA software suite [21]. The core barrel, downcomer, and pressure vessel were added according to the dimensions and materials specified in Zheng et al. [36]. Neutron pads were not modeled. A 5 cm radius spherical tally region containing air was added to the excore region, as shown in green in Fig. 9. This tally was modified with  $^{10}\text{B}$  total reaction cross section to simulate a  $\text{BF}_3$  neutron detector [37]. Since there are no streaming channels in this problem, it more closely resembles the diffusion base case.

This problem was run as a fixed source problem, and CADIS was used to optimize the convergence of the aforementioned tally. This required a converged mesh-based fission source which was obtained using a standard MC




Figure 9: Midplane slice of the Westinghouse AP1000 geometry with an excore  $^{10}\text{B}$  reaction rate tally, shown in green.




Figure 10: Converged fission source distribution used as a fixed source for calculating the  $^{10}\text{B}$  reaction rate in the tally region.

eigenvalue calculation in Shift. This calculation ran 750 inactive cycles and 750 active cycles with  $5 \cdot 10^8$  histories per cycle. The fission source distribution, which appears in Fig. 10, was tallied on a mesh with 1-pin resolution in the  $x$  and  $y$  direction and 2 cm resolution in  $z$ . This mesh was saved to disk and was later read into Shift to create a separable source—assuming a  $^{235}\text{U}$  Watt fission spectrum in energy—for subsequent fixed-source transport.

For adjoint transport in Denovo, a mesh with 4-pin resolution in  $x$ , 4-pin resolution in  $y$ , and 10 cm resolution in  $z$  was used, along with a 28-group cross section library and a quadruple range quadrature set with 4 polar and 4 azimuthal angles per octant. The adjoint flux for one of the dominant source energy groups (1.4227–1.8268 MeV) is shown in Fig. 11. This plot shows that the adjoint flux drops by 6 orders of magnitude as it reaches the




Figure 11: Contour plot of the adjoint flux distribution at the midplane for the energy group spanning 1.4227–1.8268 MeV.

fuel region, and it continues to drop steeply as it propagates deeper into the core. As a result, only a small fraction of the fuel region is important to the detector of interest, with the most important pins lying to the right of the  $10^{-6}$  contour line in Fig. 11. This indicates that source biasing will be important for improving the convergence of this problem.

Fixed source problems with preprocessed/runtime reconstructions of the linear/log versions of the SVD, as well as the analog case and CADIS base case, were then run on Summit. Ten trials with different random number seeds were run for each configuration. Each trial was run on 264 nodes, with 7 processes per node (due to memory constraints). For each trial,  $8 \cdot 10^8$  histories were simulated, except for analog, in which  $3 \cdot 10^{10}$  histories were simulated. Figure 12 shows the speedup results as a function of rank for each configuration. Unlike the base case, the linear SVD seemed to perform

approximately as well as the log version. For both versions, the rank 1 approximation incurs a factor of  $\sim 3$  performance penalty. However, for the rank 1 linear case, the biased source is separable in space/energy, and therefore its memory requirements can be reduced via Eq. (28). This is significant, as the source is spatially distributed throughout a large fraction of the geometry. As a result, the two best options for this problem are as follows:

1. Use a rank 1 approximation with linear SVD, and reduce the memory required to store both the WWs and biased source by a factor of 28 while incurring a  $\sim 3\times$  performance penalty,
2. Use a rank 2 approximation with a log SVD, which only reduces the memory required to store the WWs by a factor of 14, but which incurs no performance penalty.

In practice, the second option may be more useful, as the log SVD appears to be more reliable.

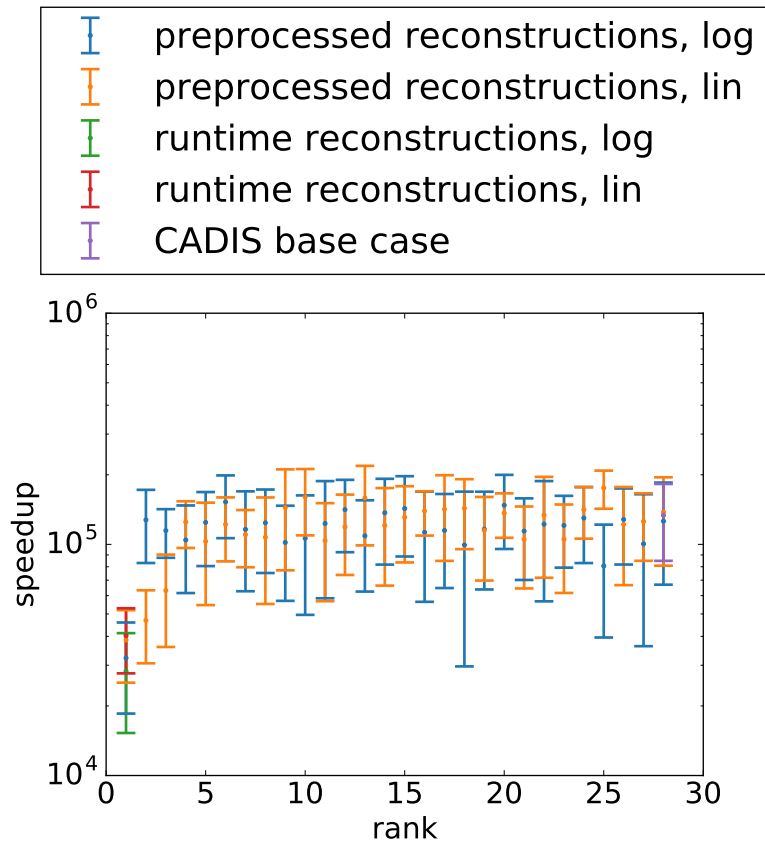


Figure 12: Speedup as a function of rank. Error bars show the standard deviation of the speedup over 10 trials using different random number generator seeds.

### 6.3. Joint European Torus

As was mentioned in Section 6.2, the AP1000 problem more closely resembles the diffusion base case. In this section, analysis is done using the Joint European Torus, which more closely resembles the streaming base case due to the presence of streaming channels.

The geometry for this problem is from an MCNP model [38] that was originally created for neutron flux and shutdown dose rate benchmarking activities [22]. This geometry is shown in Figure 13. The tally of interest for this problem is the spherical detector indicated on the geometry by an arrow, modified with flux-to-kerma conversion factors. This detector (labeled A8 in the original model) is aligned with a penetration and is located at  $z = -8.3$  cm, just below the tokamak midplane at  $z = 0$  cm. For the purposes of this problem, the radius of this detector has been expanded from 1.5 to 15 cm. The plasma neutron source for this model is separable: it is toroidal in shape, and it uses a single energy distribution, which is a Gaussian fusion spectrum for DD fusion at 10 keV.

Exnihilo has the ability to use MCNP geometries and sources for transport. However, sources defined by MCNP `SDEF` cards cannot be biased in Shift. For this reason, the plasma source in the MCNP input file was manually discretized onto a Cartesian mesh to use CADIS source biasing. This was done by extracting the `SDEF` cards from the original MCNP input file and placing them in a new file containing a dummy geometry (a 1,000 cm radius sphere). This new input file was then run with the `PTRAC` option enabled, filtering for source events, with  $1 \cdot 10^7$  histories simulated. A Python script was then used to read the position and energies of the source events and create a discrete source using a mesh with  $10 \times 10 \times 10$  cm elements spanning  $[-400$  cm,  $400$  cm] in  $x$  and  $y$  and  $[-200$  cm,  $200$  cm] in  $z$ , and an energy group structure with 100 groups between 1.7–3.2 MeV. This source is shown in Fig. 14.

For deterministic adjoint transport, a  $102 \times 92 \times 78$  nonuniform mesh was used. This mesh is finer along the path between the source, through the penetration, and to the detector, and it is also coincident with the source mesh in the source region. As in the analysis described above, the 28-group energy group structure was used. It is noted that with this group structure, the entire source energy range falls within a single energy group. Although this means that no source biasing can be done in energy, this is not important, as the source spectrum is so narrow that source biasing in energy would not have a significant effect. The deterministic adjoint transport step also used

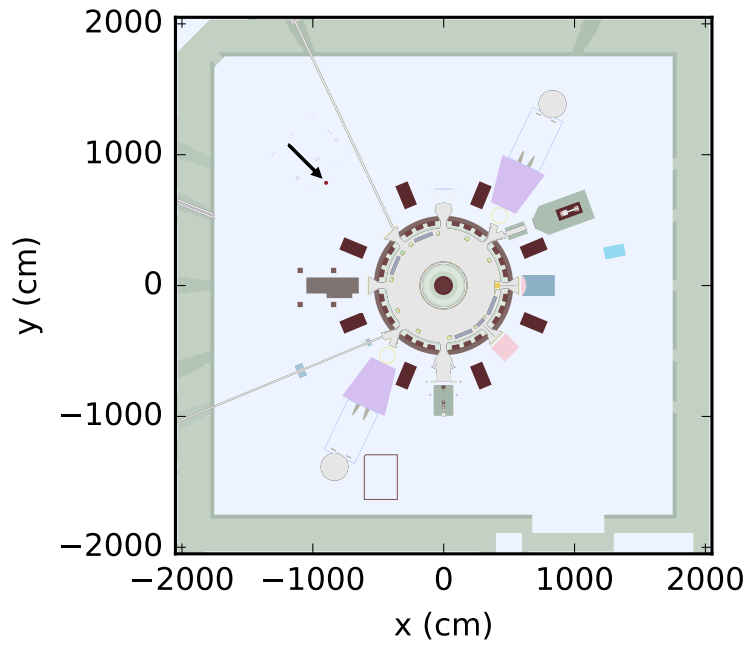


Figure 13: Slice through the JET geometry at  $z = -8.3$ . The arrow indicates the location of the detector.

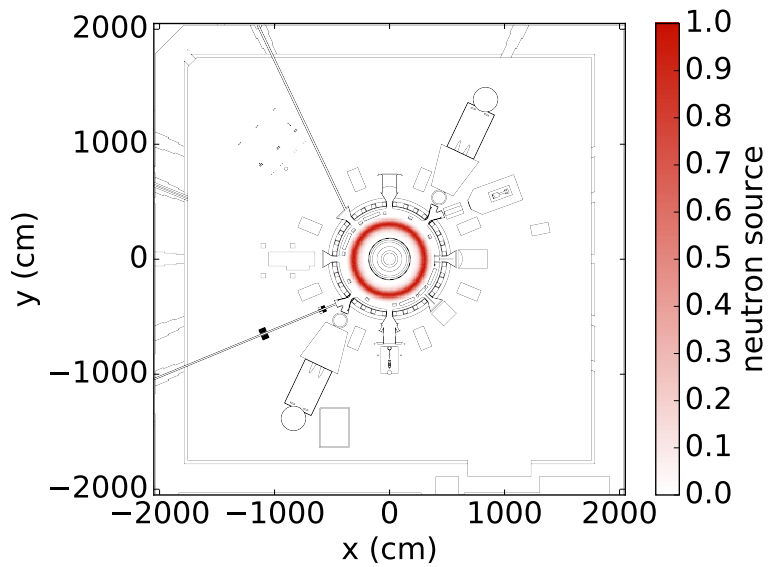


Figure 14: Slice of the discretized neutron source at  $z = -8.3$  cm.




Figure 15: Slice of the adjoint flux distribution at  $z = -8.3$  cm for the 1.83–3.01 MeV energy group.

a quadruple range quadrature set with 8 polar and 8 azimuthal angles per octant to mitigate ray effects, as the adjoint source is surrounded by void. The adjoint flux for the 1.83–3.01 MeV energy group is shown in Fig. 15. This plot clearly shows the streaming path through the penetration, which underscores the importance of source biasing in this problem.

Ten trials with different random number seeds were conducted for each SVD configuration. These trials each ran  $5 \cdot 10^8$  histories. For analog transport,  $10^9$  histories were run. Each trial was run on 20 nodes, with 20 processes per node. Speedup results for this problem are shown in Fig. 16. As in the base case problems, the speedups for this streaming problem were more modest when compared to the diffusion-dominated AP1000 problem. Due to the small solid angle subtended by the streaming channel between the source and detector, angular source biasing would be necessary to achieve speedups commensurate with the AP1000 problem. There does not appear to be a practical difference in the performance of the rank 1 reconstructions for both the linear and log SVD when compared to the full rank speedups. This means that the best option for this problem is the rank 1 linear SVD, which allows for a factor of 28 reduction in memory for both the weight windows and the biased source.


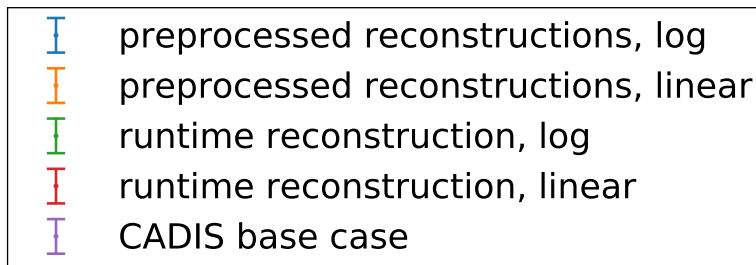


Figure 16: Speedup as a function of rank. Error bars show the standard deviation of the speedup over 10 trials using different random number generator seeds.

## 7. Conclusion

Taking the SVD of the adjoint flux is an effective method for reducing the memory requirements of VR distributions without a significant sacrifice in performance. For the log SVD, rank 1 approximations—which reduced the size of the WWs by a factor of 28—incur no performance penalty for the two base case problems and JET. A modest factor of  $\sim 3$  penalty was observed for AP1000, but this can be remedied with a rank 2 approximation, which reduces the size of WWs by a factor of 14. Low rank approximations with the linear SVD were ineffective for the base case problems, but they were similarly effective to the log version for AP1000 and JET. More work is necessary to determine the reason for this result, and is motivated by the fact that rank 1 approximations of the linear SVD have the added benefit of reducing the size of the biased source when the source is separable. For AP1000 and JET, biased sources were reduced by a factor of 28, with the same performance as the rank 1 approximations of the log SVD.

Future work will involve implementing the ability to form higher rank ( $k > 1$ ) approximations of the linear and log SVDs in Exnihilo, for cases such as AP1000, in which rank 1 approximation results in a performance penalty. The method described in Section 4.5 will be implemented to reduce the memory requirements of the biased source for the linear SVD case with  $k > 1$ . It is also worth exploring whether the relative error of the SVD or some other metric can be used to predict the effectiveness of approximations as a function of rank. Finally, the Frobenius norm in linear or log space is not mathematically optimal for approximating the adjoint flux with respect to the performance of the resulting VR parameters. New theory is necessary to determine a more appropriate norm.

## Acknowledgments

The authors would like to thank Eva Davidson for her advice on setting up the AP1000 excore problem and Katherine Royston for her help with JET. Thanks also go to Jonathan Naish for creating the MCNP model of JET and to Paola Batistoni for authorizing its use for this work. Special thanks go to Aaron Bevill for his insight in regards to performing SVD in log space. Work for this paper was supported by Oak Ridge National Laboratory, which is managed and operated by UT-Battelle, LLC, for the US Department of Energy under Contract No. DEAC05-00OR22725. This research was supported by

the Exascale Computing Project (ECP), project number 17-SC-20-SC. The ECP is a collaborative effort of two DOE organizations, the Office of Science and the National Nuclear Security Administration, that are responsible for the planning and preparation of a capable exascale ecosystem—including software, applications, hardware, advanced system engineering, and early testbed platforms—to support the nation’s exascale computing imperative. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the US Department of Energy under Contract No. DE-AC05-00OR22725.

## References

## References

- [1] G. Goertzel, M. H. Kalos, Monte Carlo methods in transport problems, *Progress in Nuclear Energy. Series 1, Physics and Mathematics* 2 (1958) 315–369.
- [2] T. E. Booth, J. S. Hendricks, Importance estimation in forward Monte Carlo calculations, *Fusion Science and Technology* 5 (1) (1984) 90–100. doi:10.13182/FST84-A23082.  
URL <https://www.tandfonline.com/doi/abs/10.13182/FST84-A23082>
- [3] M. A. Cooper, E. W. Larsen, Automated weight windows for global Monte Carlo particle transport calculations, *Nuclear Science and Engineering* 137 (2001) 1–13. doi:10.13182/NSE00-34.  
URL <https://www.tandfonline.com/doi/abs/10.13182/NSE00-34>
- [4] A. Davis, A. Turner, Comparison of global variance reduction techniques for Monte Carlo radiation transport simulations of ITER, *Fusion Engineering and Design* 86 (2011) 2698–2700. doi:10.1016/j.fusengdes.2011.01.059.  
URL <https://www.sciencedirect.com/science/article/pii/S0920379611000718>
- [5] A. Haghghat, J. C. Wagner, Monte Carlo variance reduction with deterministic importance functions, *Progress in Nuclear Energy* 42 (1) (2003) 25–53. doi:10.1016/S0149-1970(02)00002-1.

- URL <https://www.sciencedirect.com/science/article/pii/S0149197002000021>
- [6] T. E. Booth, A sample problem for variance reduction in MCNP, Tech. Rep. LA-10363-MS, Los Alamos National Laboratory (1985).
- [7] J. C. Wagner, D. E. Peplow, S. W. Mosher, FW-CADIS method for global and regional variance reduction of Monte Carlo radiation transport calculations, Nuclear Science and Engineering 176 (1) (2014) 37–57. doi:10.13182/NSE12-33.  
URL <https://www.tandfonline.com/doi/abs/10.13182/NSE12-33>
- [8] T. M. Pandya, S. R. Johnson, T. M. Evans, G. G. Davidson, S. P. Hamilton, A. T. Godfrey, Implementation, capabilities, and benchmarking of Shift, a massively parallel Monte Carlo radiation transport code, Journal of Computational Physics 308 (2016) 239–272. doi:10.1016/j.jcp.2015.12.037.  
URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999115008566>
- [9] S. W. Mosher, A. M. Bevill, S. R. Johnson, A. M. Ibrahim, C. R. Daily, T. M. Evans, J. C. Wagner, J. O. Johnson, R. E. Grove, ADVANTG—An Automated Variance Reduction Parameter Generator, Tech. Rep. ORNL/TM-2013/416, Oak Ridge National Laboratory, Oak Ridge, TN (Nov. 2013).
- [10] D. Griesheimer, D. Gill, B. Nease, T. Sutton, M. Stedry, P. Dobreff, D. Carpenter, T. Trumbull, E. Caro, H. Joo, D. Millman, MC21 v.6.0 – a continuous-energy Monte Carlo particle transport code with integrated reactor feedback capabilities, Annals of Nuclear Energy 82 (2015) 29 – 40. doi:10.1016/j.anucene.2014.08.020.  
URL <http://www.sciencedirect.com/science/article/pii/S0306454914004058>
- [11] B. Rearden, M. A. Jessee, Eds., SCALE code system, version 6.2, Tech. Rep. ORNL/TM-2005/39, Oak Ridge National Laboratory, Oak Ridge, TN (2016).
- [12] K. Royston, S. Wilson, J. Risner, A. Ibrahim, M. Loughlin, Analysis of radiation transport due to activated coolant in the ITER neutral beam

- injection cell, *Fusion Science and Technology* 72 (3) (2017) 368–373.  
doi:10.1080/15361055.2017.1333867.  
URL <https://www.tandfonline.com/doi/full/10.1080/15361055.2017.1333867>
- [13] J. Yang, S. C. Wilson, S. W. Mosher, G. Radulescu, Integration of the full tokamak reference model with the complex model for ITER neutronic analysis, *Fusion Science and Technology* 74 (4) (2018) 277–287.  
doi:10.1080/15361055.2018.1493325.  
URL <https://www.tandfonline.com/doi/full/10.1080/15361055.2018.1493325>
- [14] G. G. Davidson, T. M. Pandya, S. R. Johnson, T. M. Evans, A. E. Isotalo, C. A. Gentry, W. A. Wieselquist, Nuclide depletion capabilities in the Shift Monte Carlo code, *Annals of Nuclear Energy* 114 (2018) 259–276. doi:<https://doi.org/10.1016/j.anucene.2017.11.042>.  
URL <http://www.sciencedirect.com/science/article/pii/S0306454917304322>
- [15] S. P. Hamilton, T. M. Evans, Continuous-energy Monte Carlo neutron transport on GPUs in the Shift code, *Annals of Nuclear Energy* 128 (2019) 236–247. doi:10.1016/j.anucene.2019.01.012.  
URL <http://www.sciencedirect.com/science/article/pii/S0306454919300167>
- [16] D. E. Peplow, Comparison of hybrid methods for global variance reduction in shielding calculations, in: *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, Sun Valley, ID, 2013.
- [17] J. A. Ellis, T. M. Evans, S. P. Hamilton, C. Kelley, T. M. Pandya, Optimization of processor allocation for domain decomposed Monte Carlo calculations, *Parallel Computing* 87 (2019) 77–86. doi:<https://doi.org/10.1016/j.parco.2019.06.001>.  
URL <http://www.sciencedirect.com/science/article/pii/S0167819119301243>
- [18] T. M. Evans, A. S. Stafford, R. N. Slaybaugh, K. T. Clarno, Denovo: A new three-dimensional parallel discrete ordinates code in SCALE,

- Nuclear Technology 171 (2010) 171–200. doi:10.13182/NT171-171.  
URL <https://www.tandfonline.com/doi/abs/10.13182/NT171-171>
- [19] S. P. Hamilton, T. M. Evans, Efficient solution of the simplified  $P_N$  equations, *Journal of Computational Physics* 284 (2015) 155–170. doi:10.1016/j.jcp.2014.12.014.  
URL <https://www.sciencedirect.com/science/article/pii/S0021999114008225>
- [20] S. R. Johnson, Omnibus: a new front end to Denovo and Shift, *Transactions of the American Nuclear Society* 117 (2017).
- [21] F. Franceschini, A. Godfrey, J. Kulesza, R. Oelrich, Westinghouse VERA test stand: zero power physics test simulations for the AP1000 PWR, Tech. Rep. CASL-U-2014-0012-000, CASL (2014).
- [22] R. Villari, P. Batistoni, J. Catalan, B. Colling, D. Croft, U. Fischer, D. Flammini, N. Fomesu, L. Jones, A. Klix, B. Kos, M. Kłosowski, I. Kodeli, S. Loreti, F. Moro, J. Naish, B. Obryk, L. Packer, P. Pereslavitsev, R. Pilotti, S. Popovichev, P. Sauvan, I. Stamatelatos, T. Vasilopoulou, ITER oriented neutronics benchmark experiments on neutron streaming and shutdown dose rate at JET, *Fusion Engineering and Design* 123 (2017) 171 – 176, proceedings of the 29th Symposium on Fusion Technology (SOFT-29) Prague, Czech Republic, September 5-9, 2016. doi:10.1016/j.fusengdes.2017.03.037.  
URL <http://www.sciencedirect.com/science/article/pii/S0920379617302405>
- [23] G. Bell, S. Glasstone, *Nuclear Reactor Theory*, New York (1971).
- [24] M. Munk, R. N. Slaybaugh, Review of hybrid methods for deep-penetration neutron transport, *Nuclear Science and Engineering* 193 (10) (2019) 1055–1089. doi:10.1080/00295639.2019.1586273.  
URL <https://www.tandfonline.com/doi/full/10.1080/00295639.2019.1586273>
- [25] H. Andrews, C. Patterson, Singular value decomposition (SVD) image coding, *IEEE Transactions on Communications* 24 (4) (1976) 425–432. doi:10.1109/TCOM.1976.1093309.  
URL <https://ieeexplore.ieee.org/document/1093309>

- [26] O. Alter, P. O. Brown, D. Botstein, Singular value decomposition for genome-wide expression data processing and modeling, *Proceedings of the National Academy of Sciences* 97 (18) (2000) 10101–10106. doi:10.1073/pnas.97.18.10101.  
URL <https://www.pnas.org/content/97/18/10101>
- [27] R. Kumaresan, D. Tufts, Estimating the parameters of exponentially damped sinusoids and pole-zero modeling in noise, *IEEE Transactions on Acoustics, Speech, and Signal Processing* 30 (6) (1982) 833–840. doi:10.1109/TASSP.1982.1163974.  
URL <https://ieeexplore.ieee.org/document/1163974>
- [28] J. Navarro-Esbrí, G. Verdú, D. Ginestar, J. Muñoz-Cobo, Reactor noise analysis based on the singular value decomposition (SVD), *Annals of Nuclear Energy* 25 (12) (1998) 907 – 921. doi:10.1016/S0306-4549(97)00090-X.  
URL <http://www.sciencedirect.com/science/article/pii/S030645499700090X>
- [29] H. S. Abdel-Khalik, P. J. Turinsky, M. A. Jessee, Efficient subspace methods-based algorithms for performing sensitivity, uncertainty, and adaptive simulation of large-scale computational models, *Nuclear Science and Engineering* 159 (3) (2008) 256–272. doi:10.13182/NSE159-256.  
URL <https://www.tandfonline.com/doi/abs/10.13182/NSE159-256>
- [30] G. Golub, C. Van Loan, *Matrix Computations*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, 2013.
- [31] C. Eckart, G. Young, The approximation of one matrix by another of lower rank, *Psychometrika* 1 (3) (1936) 211–218. doi:10.1007/BF02288367.  
URL <https://ideas.repec.org/a/spr/psycho/v1y1936i3p211-218.html>
- [32] E. D. Biondo, A. Davis, P. P. Wilson, Shutdown dose rate analysis with CAD geometry, Cartesian/tetrahedral mesh, and advanced variance reduction, *Fusion Engineering and Design* 106 (2016) 77–84. doi:<https://doi.org/10.1016/j.fusengdes.2016.03.004>.  
URL <http://www.sciencedirect.com/science/article/pii/S0920379616302009>

- [33] M. Heroux, R. Bartlett, V. H. R. Hoekstra, J. Hu, T. Kolda, R. Lehoucq, K. Long, R. Pawlowski, E. Phipps, A. Salinger, H. Thornquist, R. Tuminaro, J. Willenbring, A. Williams, An Overview of Trilinos, Tech. Rep. SAND2003-2927, Sandia National Laboratories (2003).
- [34] Oak Ridge Leadership Computing Facility, Summit: Oak Ridge National Laboratory's next high performance supercomputer, <https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit> (April 2018).
- [35] K. Kobayashi, N. Sugimura, Y. Nagaya, 3D radiation transport benchmark problems and results for simple geometries with void region, *Progress in Nuclear Energy* 39 (2) (2001) 119 – 144. doi:10.1016/S0149-1970(01)00007-5.  
URL <http://www.sciencedirect.com/science/article/pii/S0149197001000075>
- [36] J. Zheng, B. Zhang, S. Shi, Y. Chen, Analysis for the effect of spatial discretization method on AP1000 reactor pressure vessel fluence calculation, *Science and Technology of Nuclear Installations* 2016 (Article ID 3461290) (2016). doi:10.1155/2016/3461290.  
URL <https://www.hindawi.com/journals/stni/2016/3461290/>
- [37] H. P. Smith, E. E. Davidson, A. T. Godfrey, T. Pandya, An analysis of various solution strategies and perturbations on inputs of the reactor shielding problem, in: ANS RPSD 2018 - 20th Topical Meeting of the Radiation Protection & Shielding Division of ANS, Santa Fe, NM, 2018.
- [38] J. Naish, Model of JET — DT Source, WPJET3-NEXP STRE subproject within EUROfusion Consortium (2014).