

Optimization of Canister Loading Patterns in Dual Purpose Canisters for Criticality Suppression

Spent Fuel and Waste Disposition

***Prepared for
US Department of Energy
Spent Fuel and Waste Science
and Technology***

***Pacific Northwest National Laboratory:
Justin Clarity, Harish Gadey, Peter
Stefanovic, Paul Miller, and Kaushik
Banerjee***

November 29, 2022

M3SF-22PN010305094

PNNL-33974

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

This is a technical report that does not take into account contractual limitations or obligations under the Standard Contract for Disposal of Spent Nuclear Fuel and/or High-Level Radioactive Waste (Standard Contract) (10 CFR Part 961). For example, under the provisions of the Standard Contract, spent nuclear fuel in multi-assembly canisters is not an acceptable waste form, absent a mutually agreed to contract amendment.

To the extent discussions or recommendations in this report conflict with the provisions of the Standard Contract, the Standard Contract governs the obligations of the parties, and this report in no manner supersedes, overrides, or amends the Standard Contract.

This report reflects technical work which could support future decision making by DOE. No inferences should be drawn from this report regarding future actions by DOE, which are limited both by the terms of the Standard Contract and Congressional appropriations for the Department to fulfill its obligations under the Nuclear Waste Policy Act including licensing and construction of a spent nuclear fuel repository.

SUMMARY

This report documents work performed in support of the US Department of Energy Office of Nuclear Energy (NE) Spent Fuel and Waste Disposition, Spent Fuel and Waste Science and Technology, under work breakdown structure element 1.08.01.03.05, “Direct Disposal of Dual Purpose Canisters.” In particular, this report fulfills milestone M3SF-22PN010305094, “Application of AI/ML techniques to optimize DPC loading,” within work package SF-22PN01030509, “Direct Disposal of Dual Purpose Canisters - PNNL.”

This report continues the process of examining the potential for using loading optimization as a disposal criticality suppression technique. This latest update to the report added:

1. A validation of the artificial neural network (ANN) reactivity prediction tool against as-loaded dual-purpose canisters from the UNF-ST&DARDS database
2. Refinement of the algorithm to incorporate an additional DPC design to improve the predictions
3. Investigation of a storage and transportation focused loading optimization algorithm.

A key result from this year’s work is that the ANN performance was significantly improved by including the additional model for the NUHOMS canisters and it is apparent that fuel type specific modeling considerations should be incorporated into the ANN in future work. The GRASP-enabled adaptive multi-objective memetic algorithm with partial clustering (GAMMA-PC) algorithm was evaluated as a tool for storage and transportation oriented optimization. The GAMMA-PC optimization routine considers decay heat at loading, the time between loading and when the canister is eligible for transportation and minimizes the number of casks loaded. Future work will develop a dose minimization optimization routine.

CONTENTS

SUMMARY	iii
LIST OF FIGURES	v
REVISION HISTORY.....	vii
ACRONYMS.....	viii
1. INTRODUCTION.....	1
2. METHODOLOGY	3
2.1 Canister Criticality Calculations	3
2.2 Assembly Reactivity Calculations	3
2.3 Canister Reloading of Fuel at a Single Site	3
2.4 Development of a full life-cycle DPC loading optimization tool	6
2.4.1 Development of ANN for Rapid Reactivity Prediction	7
2.5 Storage and Transportation Fuel Loading Optimization in UNF-ST&DARDS	10
2.5.1 Typical Acceptance Criteria for Loading into Dry Storage and Transport.....	11
2.5.2 Fuel Loading Optimization Methodology in UNF-ST&DARDS	12
2.5.3 Data Input from UNF-ST&DARDS to Fuel Loading Optimization.....	12
3. RESULTS AND DISCUSSION.....	13
3.1 Single Site Reloading Results.....	13
3.2 Development of ANN for canister reactivity prediction	14
3.3 Impacts of training data on k_{eff} predictions	20
3.4 Validation of the ANN against as-loaded casks.....	21
3.5 Damaged Fuel Predictions	27
3.6 Fuel Loading Optimization Case Study for a PWR Nuclear Plant	28
4. CONCLUSIONS AND FUTURE WORK.....	33
5. REFERENCES	34

LIST OF FIGURES

Figure 1. Canister loaded with highly reactive assembly in low-leakage locations (left) and assembly reactivity distribution for all assemblies from shutdown sites vs. those discharged from the last cycle of operation (right).....	2
Figure 2. Histogram of Zion intact fuel assembly reactivities, with the single bin approach (left) and the four-bin approach (right).....	5
Figure 3. Assembly reactivity rank loading scheme for the TSC-37 canisters used at Zion.	5
Figure 4. Flow chart showing the components of a wholistic optimization approach that incorporates post-closure criticality.....	6
Figure 5. Diagram of the neural network used with the input, output and hidden layers labeled.....	8
Figure 6. Burnup and enrichment probability distribution of the 17×17 assemblies in the UDB.....	9
Figure 7. Loading scheme used to generate the “maximum” case canister k_{eff} calculations.	9
Figure 8. Loading scheme used to generate the “minimum” case canister k_{eff} calculations.....	10
Figure 9. UNF-ST&DARDS data input to Loading Optimization (GAMMA-PC).....	13
Figure 10. Box plots for the single site reloading analysis.	14
Figure 11. Histogram of assembly reactivity values used to train the ANN.....	15
Figure 12. Histogram of MPC-32 canister k_{eff} values used to train the ANN.....	16
Figure 13. Histogram of the NUHOMS canister k_{eff} values used to train the ANN.	17
Figure 14. Training and validation errors for the neural network.....	18
Figure 15. Predicted k_{eff} as a function of testing k_{eff} (left) and frequency of error (right) for the testing data set.	19
Figure 16. (a) Prediction error for MPC-32 canisters while using random training data; (b) error in the k_{eff} prediction while using worse, optimal, and random assembly loading scenarios for MPC-32 canisters.....	20
Figure 17. a) Prediction error for NUHOMS-32 canisters while using random training data; b) error in the k_{eff} prediction while using worse, optimal, and random assembly loading scenarios for NUHOMS-32 canisters.	21
Figure 18. The true and predicted k_{eff} values of a wide variety of fuel assemblies and canister designs using the MPC-32 ANN model	22
Figure 19. Comparison of the UNF-ST&DARDS calculated and Randomized MPC-32 ANN predicted canister k_{eff} values for 432 DPC validation set.	24
Figure 20. Comparison of the UNF-ST&DARDS calculated and Randomized MPC-32 and NUHOMS-32 ANN predicted canister k_{eff} values for 432 DPC validation set.	25
Figure 21. Comparison of the UNF-ST&DARDS calculated and Expanded MPC-32 and NUHOMS-32 ANN predicted canister k_{eff} values for 432 DPC validation set.	26
Figure 22. True and predicted canister k_{eff} values for damaged assemblies while using predefined k_{inf} assembly reactivity and a k_{inf} of 1.4.....	27
Figure 23. Loading map for MPC-32 Cask 1.....	31
Figure 24. Number of Casks in the Loading Plan by Decay Heat	32

Figure 25. Number of Casks Loaded per Decade.	32
--	----

REVISION HISTORY

Revision	Effective Date	Changes Made
0	June 25, 2021	Initial issue
1	November 29, 2022	The following items were added in Revision 1 of the report: (1) This work expands the ANN model to include the NUHOMS 32-assembly cask models, (2) a wider range of loadings with the same fuel types was included in the training data set, (3) the storage and transportation oriented loading optimization algorithm was used to generate a baseline loading dataset for future reactivity analysis.

ACRONYMS

ANN	artificial neural network
CoC	certificate of compliance
DPC	dual-purpose canister
FEPs	features, events, and processes
GAMMA-PC	GRASP-enabled adaptive multi-objective memetic algorithm with partial clustering
GELU	Gaussian exponential linear unit
KDE	kernel density estimator
NA	no absorber (loss of neutron absorber case)
NE	Office of Nuclear Energy
ORNL	Oak Ridge National Laboratory
PWR	pressurized water reactor
SELU	scaled exponential linear unit
UDB	Unified Database
UNF-ST&DARDS	Used Nuclear Fuel Storage, Transportation and Disposal Analysis Resource and Data System

OPTIMIZATION OF CANISTER LOADING PATTERNS IN DUAL PURPOSE CANISTERS FOR CRITICALITY SUPPRESSION

1. INTRODUCTION

The current spent nuclear fuel management strategy includes reliance on dry storage. Utilities are meeting their interim storage needs on an individual basis with the use of large-capacity dry storage casks. The focus is currently on meeting existing storage and transportation requirements, as disposal requirements are not currently available. The casks are commonly known as *dual-purpose canisters* (DPCs), serving the dual purposes of storage and transportation. However, a small percentage of single-purpose (storage-only) systems is also being used to meet storage needs. This report investigates the post-closure criticality safety aspects of DPCs.

Placement of large, heavy waste packages containing DPCs into a repository for direct disposal has not yet been implemented domestically or internationally. Therefore, the prospect of direct disposal of DPCs presents new engineering and scientific challenges. Some of these challenges have already been investigated, including handling and placement, use of ramps vs. shafts, use of hoists, use of transport equipment, and thermal management. Additionally, some studies have been conducted to determine the feasibility of direct disposal from a criticality analysis perspective. These studies concluded that direct disposal is possible but showed that it is challenging to demonstrate subcriticality over the disposal time period [1,2]. The alternative to directly disposing of DPCs in a repository is to repackage the spent nuclear fuel into smaller canisters for disposal. Direct disposal of DPCs without cutting them open and repackaging is appealing because it could be more cost effective, could reduce the complexity of fuel management operations inside and outside reactor facilities, and could result in lower cumulative worker doses during interim storage and handling before eventual disposal in a deep geologic repository.

The performance of the basket neutron absorber material as a function of time inside the canister is a key factor for demonstrating subcriticality. The neutron absorber panel material used in the majority of currently loaded DPCs is Boral®. Boral is composed of B₄C particles and alloy 1100 aluminum that are hot-rolled together to form a neutron-absorbing core. The neutron-absorbing core is then bonded to two outer layers of alloy 1100 aluminum. Various corrosion tests have been performed on this material because it is used in existing casks and in spent fuel pools. Corrosion tests that were conducted under pool chemistry conditions showed that Boral cladding material was lost at a rate of 0.28 mil per year, which equates to a service life of approximately 40 years (in the presence of water) before degradation of the neutron-absorbing core [1]. Additionally, some tests of Boral under simulated vacuum drying processes have shown that the drying process induces blisters within the Boral [3]. Since the repository periods of interest are expected to be at least 10,000 years, the Boral neutron absorber will not likely maintain its criticality control function that long if the package cavity is exposed to an aqueous environment.

A repository performance assessment examines the features, events, processes (FEPs), and sequences of FEPs that might affect the repository. Criticality is considered such an event within the FEPs nomenclature, with the potential to affect repository performance. Before a performance assessment is conducted, any FEPs that can affect repository performance are screened for inclusion or exclusion. Based on previous screening criteria, options available for excluding a feature, event, or process consisted of a low-probability criterion, a low-consequence criterion, and regulation. Previous work by Clarity et al. using The Used Nuclear Fuel Storage, Transportation and Disposal Analysis Resource and Data System (UNF-ST&DARDS) [4] showed that, in many cases, considering the as-loaded contents of DPCs results in enough evidence to offset the potential degradation of the aluminum-based neutron absorbers.

However, there are a substantial number of DPCs for which subcriticality cannot be demonstrated based on crediting the as-loaded contents alone.

The as-loaded criticality assessment described in Clarity et al. [4] showed that many of the canisters with k_{eff} values greater than the hypothetical limit of 0.98 were placed into configurations so that the most reactive individual fuel assemblies were loaded in the lowest neutron leakage locations in the center of the basket. This configuration resulted in the highest canister reactivity for a given set of fuel assemblies loaded in a cask. This loading scenario is illustrated in the left portion of Figure 1 using a TSC-37 canister from Zion. In Figure 1, the canister basket cells are shaded according to the reactivity of the individual assembly (assembly reactivity calculations are discussed in Section 2.2) that they contain. The k_{eff} for the canister shown in the left portion of Figure 1 is approximately 1.09, because the most reactive assemblies are placed in the center of the DPC. Assembly loading within a canister is a particularly important issue for shutdown sites for which the fuel assemblies discharged for the final cycle of operation have not had the ability to achieve the same burnup as an equivalent enrichment assembly would earlier in a plant's operation. The right-hand image in Figure 1 shows a histogram of assembly reactivities for the pressurized water reactor (PWR) assemblies from many shutdown sites in the Unified Database (UDB) shaded in blue, and the reactivities of those assemblies discharged from the last cycle of operation of the shutdown sites are shown in red. It is apparent that the most reactive assemblies in the inventory are discharged at the end of operation.




Figure 1. Canister loaded with highly reactive assembly in low-leakage locations (left) and assembly reactivity distribution for all assemblies from shutdown sites vs. those discharged from the last cycle of operation (right).

This report documents an investigation of potential reactivity management strategies that can be used to reduce the number of canisters susceptible to criticality over the repository performance assessment period. The investigations include reloading the assemblies in the canisters at the Zion site using the individual assembly reactivity and the development of an artificial neural network (ANN) for canister reactivity prediction. The goal of the canister reactivity ANN is to be able to rapidly predicted k_{eff} so that it can be used for optimum loading searches. These approaches will facilitate criticality analysis to be incorporated in the decision-making process that will be used for canister loading.

The remainder of this report is organized as follows: Section 2 discusses the methods and inputs used to reload canisters at Zion, to develop the ANN for rapid canister reactivity estimation and for the storage and transportation oriented loading algorithm; Section 3 presents the canister reloading results, documents the performance of the ANN, and describes the use of the loading algorithm for a sample site; and Section 4 discusses the conclusions derived from this work.

2. METHODOLOGY

2.1 Canister Criticality Calculations

The burnup credit and criticality computational methods used to perform this work are discussed at length in Clarity et al. [4]. The same computational pipeline, ORIGEN reactor libraries, and DPC computational templates are used here as in the previous work [4]. All criticality calculations performed to support this work used the SCALE 6.2.3 version of the KENO-VI Monte Carlo code with ENDF/B-VII.1 nuclear data. The Monte Carlo parameters used herein are 550 generations of 50,000 neutron histories, skipping the first 50 generations to allow for source convergence. These parameters are sufficient to produce a Monte Carlo uncertainty on k_{eff} of 25 pcm (one sigma) or less for all calculations. All canister criticality calculations performed here are analyzed for the date July 1, 22,000, since this typically corresponds to the most reactive time step analyzed under disposal conditions [4].

For the single site reloading calculations, new loading maps were constructed and uploaded to a developmental copy of the UDB using the same method that is implemented for as-loaded analysis of real canisters existing at current sites.

2.2 Assembly Reactivity Calculations

Assembly reactivity calculations are single assembly criticality calculations in which the assembly is modeled as being infinitely reflected on all four radial surfaces. Because the models are relatively simple, the Monte Carlo parameters are relaxed to 10,000 particles per generation, skipping the first 25 generations. The calculations are run with enough generations to obtain a Monte Carlo uncertainty on k_{eff} of 75 pcm (one sigma). Assembly reactivities have traditionally been calculated at 9,500 years of cooling time in UNF-ST&DARDS, and that cooling time is used for all calculations in this report.

The assembly reactivity values are used as a single parameter that estimates how much a given assembly would contribute to the k_{eff} of a canister in comparison to other assemblies. The assembly reactivity allows for comparison of assemblies with different enrichment, burnup, and cooling time combinations. For the canister reloading analysis, assembly reactivity is used to determine the order in which the assemblies are loaded into canisters. Assembly reactivity is also used to bin the assemblies for distribution amongst the canisters. For the ANN, assembly reactivities are used as predictor variables in the training and prediction routines.

2.3 Canister Reloading of Fuel at a Single Site

To gain an understanding of the potential reactivity reduction obtainable by optimizing canister loadings, the assemblies at Zion were reorganized into new canisters. Zion was selected for this analysis because it was identified [4] as having several canisters with k_{eff} values in excess of the hypothetical 0.98 subcritical limit. The Zion site has 61 canisters that contain 2,226 assemblies. Of these assemblies, 2,166 are intact, and 60 have been declared as damaged fuel. The damaged fuel assemblies are eligible for storage in the four peripheral locations along the minor axes of the canister. The damaged fuel assemblies are modeled with the fresh fuel assumption, with an expanded pitch to conservatively account for uncertainties in assembly geometry.

Two types of decisions must be made when determining the method to be used to load the fuel assemblies into hypothetical canisters:

1. Selection of assemblies to be loaded into a specific canister, and
2. Selection of the location in the canister that each assembly will occupy.

The algorithm used to assign assemblies to canisters is as follows:

1. Divide the assembly inventory into damaged and intact fuel categories.
2. Determine the number of canisters needed to accommodate the damaged fuel assemblies. This was done by dividing the number of damaged fuel assemblies by 4, which is the number of locations that can accommodate damaged fuel assemblies in a TSC-37. The number of canisters needed to accommodate the damaged fuel is 15.
3. Determine the overall number of canisters to be filled with only intact fuel. This was done by subtracting the number of intact assemblies needed to completely fill the canister with damaged fuel canisters (495) from the total inventory of intact fuel and dividing that number by the canister capacity (37). Since the number of canisters must be an integer, an extra canister was added. The number of canisters necessary to accommodate all of the intact fuel is 46, for a total of 61 canisters.
 - a. It was decided that short-loaded canisters should be loaded with 36 intact assemblies. A total of 30 canisters are filled with 36 assemblies. These canisters were filled first.
 - b. The remainder of the canisters containing only intact fuel were filled with 37 assemblies. These canisters were filled second.
 - c. Finally, the damaged fuel canisters were filled with 4 damaged fuel assemblies and 33 intact fuel assemblies.
4. When the intact fuel assemblies were loaded into each canister, the assemblies were binned into groups with equal numbers of assemblies according to assembly reactivity, and the top and bottom assemblies were chosen sequentially from each group, starting with the highest reactivity group. The process was then repeated for the next group. To investigate the impact of the uniformity of selecting assemblies from the assembly reactivity distribution, the number of groups used was varied from one to four. The histogram of the one-bin and four-bin assembly reactivity group structures is shown in Figure 2. Each canister was filled in this manner before moving on to the next canister.
5. Damaged fuel assemblies do not have assembly reactivity values in the UDB. Because they are modeled as being fresh fuel with the manufactured enrichment, the enrichment was used as a proxy for assembly reactivity because reactivity scales with enrichment for fresh fuel. Using the same process of sequentially selecting the highest and lowest enrichments, only a single group of damaged fuel was used.

Once the assemblies were selected for each canister, they were ranked according to assembly reactivity and distributed according to the map shown in Figure 3, where the highest reactivity is one, and the lowest reactivity is 37.




Figure 2. Histogram of Zion intact fuel assembly reactivities, with the single bin approach (left) and the four-bin approach (right).




Figure 3. Assembly reactivity rank loading scheme for the TSC-37 canisters used at Zion.

2.4 Development of a full life-cycle DPC loading optimization tool

The criticality-oriented optimization techniques presented in Section 2.3 represent an idealized case in which criticality is the only parameter considered. Ultimately, the goal is to develop a tool that will allow for the consideration of disposal criticality in a more wholistic optimization scheme that would also meet loading constraints associated with on-site storage, such as decay heat and dose. It is envisioned that there would be at least three components of an optimization routine that would be able to meet the integrated life-cycle needs of storage, transportation and disposal. The components that will be necessary are 1.) a storage and transportation focused optimization tool, 2.) the ability to rapidly predict canister reactivity, and 3.) the ability to adapt the loadings of highly reactive canisters. These techniques will need to be employed iteratively to arrive at loadings that are acceptable for all storage, transportation, and disposal as is depicted in Figure 4. The initial implementation of a tool to optimize canister loadings for storage and transportation is the GRASP-enabled adaptive multi-objective memetic algorithm with partial clustering (GAMMA-PC) [5] which has been included in UNF-ST&DARDS and is further discussed in Section 2.5.2 and the development of the ANN for rapid reactivity predictions is discussed in Section 2.4.1. Algorithms for modifying loadings will be explored in future work.




Figure 4. Flow chart showing the components of a wholistic optimization approach that incorporates post-closure criticality.

2.4.1 Development of ANN for Rapid Reactivity Prediction

Optimization techniques generally require that multiple loading patterns be considered before a solution is chosen that meets all constraints and generates an optimum value for parameters of interest. Criticality assessment is typically performed using Monte Carlo eigenvalue calculations. The durations of these calculations for the highly heterogeneous models considered in as-loaded analysis typically spans several hours to a few days per configuration. The time required to run these calculations makes them ill-suited for the rapid criticality assessment desired when considering multiple design constraints. To allow for approximate reactivity assessment of as-loaded canisters, an ANN is developed for this purpose.

This work aims to develop an ANN to predict the k_{eff} of a canister given its contents. ANNs are mathematical constructs loosely modeled after neurological pathways in the brain that are composed of a series of neurons or units that perform mathematical functions. Units within a neural network are organized in layers that fall into three broad classifications: 1) the input layer, 2) the hidden layers, and 3) the output layer. The input layer accepts input variable values. The first hidden layer receives a weighted sum of the inputs and applies an activation function to the sum, and each subsequent layer takes the weighted sum of the outputs of the units in the preceding layer and applies an activation function to the sum. The output layer applies an activation function to the weighted sum of the last hidden layer. Tunable parameters that are used in developing a neural network include the activation functions used, the number of hidden layers, and the number of units per hidden layer.

ANNs are developed or trained by feeding a number of known input and output sets so that the appropriate weights for each variable at each unit can be determined. For the development of canister k_{eff} ANN, the values input into the ANN are the assembly reactivities, and the output values are the Monte Carlo-calculated canister k_{eff} values. A diagram showing five input nodes, three hidden layers of ten units, and an output unit is provided in Figure 5.

To generate training data for the preliminary version of the ANN, 1,000 randomly loaded canisters were analyzed. The fuel assembly used for this purpose was the Westinghouse 17×17 Standard fuel assembly (GC-859 designation W1717WL). The canister used for this analysis was the MPC-32. The first step in this process was to generate a random inventory of fuel assemblies representative of the assemblies currently being stored. A two-dimensional kernel density estimator (KDE) was developed using the enrichments and burnups of 105,139 PWR fuel assemblies in the UDB. A plot of the KDE's probability density is provided in Figure 6. Using the KDE, 32,000 burnup and enrichment combinations were generated to represent the assemblies loaded into 1,000 canisters. Assembly reactivity calculations were run for each of the 32,000 assemblies at 9,500 years of cooling time, and 1,000 canister criticality calculations were run in the year 22,000—which represents the peak reactivity time for most canisters—with the loss of neutron absorber (NA) hypothetical disposal scenario.

Based on the work performed in Revision 0 of this document it was decided that it may be valuable to try to expand the range of the canister loadings the ANN is trained with to include high and low reactivity loading patterns as well as random loading patterns. The low reactivity patterns rank the assemblies assigned to the canisters in a manner similar to that described for the single site reloading exercise in Section 2.3 and were assigned to positions in the canister according to the loading scheme in Figure 7. For the minimum reactivity configurations the same process is followed, except the assemblies are assigned to positions according to the reactivity map in Figure 8.

The ANN was trained using the assembly reactivity values as inputs and the canister k_{eff} values as outputs. The data were initially divided so that 70% of the canisters were used as the training data, and 20% canisters were reserved as testing data. Reserving the testing data allowed for independent quantification of algorithm performance once the ANN has been fully trained. Furthermore, 10% of canisters were withheld to serve as a validation data set. The validation data set is used to perform preliminary assessments of the algorithm's accuracy and to decide the number of passes to be made through the training process, or the epochs to be used. Once the validation error has converged to a stable

value, additional epochs are no longer needed. Canisters were randomly assigned to the training, testing, and validation data sets. The training process involves evaluating all the training data and applying a stochastic gradient descent algorithm to the weights used at each unit within the ANN. The ANN trained for this effort attempts to minimize the absolute value of the error between the predicted k_{eff} and the true canister k_{eff} . The Tensorflow 2 machine learning platform was used for this work. Tensorflow 2 allows for rapid implementation of ANNs by providing the streamlined Keras user interface. All calculations are performed using Python scripts.

For this year's work additional high and low reactivity canister loading patterns are considered. The same assemblies as were used with the randomized loading are reoriented in the canisters into the configurations shown in Figure 7 for the high reactivity configuration and Figure 8 for the low reactivity configuration. These data are then added to the training data in the same way that the randomized loadings are.




Figure 5. Diagram of the neural network used with the input, output and hidden layers labeled.




Figure 6. Burnup and enrichment probability distribution of the 17×17 assemblies in the UDB.




Figure 7. Loading scheme used to generate the “maximum” case canister k_{eff} calculations.

		5	4	4	5	
5	3	2	2	3	5	
4	2	1	1	2	4	
4	2	1	1	2	4	
5	3	2	2	3	5	
		5	4	4	5	

Figure 8. Loading scheme used to generate the “minimum” case canister k_{eff} calculations.

2.5 Storage and Transportation Fuel Loading Optimization in UNF-ST&DARDS

The assemblies and associated dry storage systems must be comprehensively characterized to support dry storage loading, extended storage beyond the initial 20 years of the storage license, shipping and handling after extended storage, and eventual disposal of the loaded canisters. Typically, dry casks have been primarily loaded according to dry storage requirements outlined in a dry cask Certificate of Compliance (CoC). This may complicate transportation in the future because transportation casks have a different set of loading requirements. For example, a CoC for a transport package may apply more restrictive criticality, shielding and thermal requirements. Criticality requirements for a transport package often include criteria expressed in terms of maximum enrichment vs. minimum burnup (or burnup curve) that may be more restrictive than the requirements for dry cask storage. To meet shielding requirements, such as surface dose rates or dose rates at a distance from a transport package, loadings for transport may require a different placement of assemblies and inserts within the fuel basket when compared with placements limitations in dry storage. Similarly, limits on basket cell decay heat and total canister decay heat may vary in dry storage and in a transport cask. Further, the loading may become even more restrictive when long-term storage or disposal of the fuel is considered. Therefore, with a rapidly growing UNF inventory, a primary area of need in dry storage planning is to apply methodologies to optimize dry cask loading with respect to long-term storage and transportation needs.

Integration of the dry storage optimized loading optimization tool called GAMMA-PC into UNF-ST&DARDS facilitates the formulation of cost-effective solutions, with improved worker safety and efficient transportation timelines. The tool generates optimal loading patterns to balance the competing goals of higher loading capacities, earlier fill and transport times, and overall lower dose rates.

2.5.1 Typical Acceptance Criteria for Loading into Dry Storage and Transport

Spent nuclear fuel management in a nuclear power plant includes planning to load SNF assemblies into dry storage over the lifetime of the plant including subsequent SNF transportation off-site. The requirements for loading into dry storage are defined in an associated dry storage cask CoC. Below is a list of typical acceptance criteria to be satisfied during dry cask storage loading:

- *Intended Loading Date* – Loading schedule is determined based on spent fuel pool capacities, availability of dry cask storage, requirements to unload wet storage due to plant shutdown, and other plant needs to move fuel from wet storage into dry storage.
- *Minimum Assembly Cooling Time* - For each assembly in the canister, the cooling time is calculated as the time between fuel discharge and the cask loading date. The lowest acceptable cooling time is typically provided in the CoC from the cask vendor. The minimum cooling time requirement is typically driven by limitation on dose rates from a cask and/or a dry storage installation.
- *Maximum Planar Average Initial Enrichment (MPAIE)* – as a part of criticality control, CoC may specify MPAIE based on assembly type being loaded into fuel basket.
- *Assembly Types* – To maintain subcriticality, the CoC may restrict only certain fuel assembly types into specific baskets or restrict mixing of certain fuel types in the same basket.
- *Damaged Assembly Location* – Assemblies that are not intact and contain certain damage may still be loaded into a dry cask canister but may be limited to certain locations in the fuel basket.
- *Maximum Burnup* – The maximum assembly burnup affects assembly residual decay heat and radiation source terms, therefore restrictions on assembly burnup are placed in the CoC.
- *Decay Heat* – To maintain peak fuel cladding temperature, the CoC provides limits on assembly decay heat per basket cell.
- *Operational Considerations:*
 - Availability of fuel for loading – Some fuel assemblies in a spent fuel pool may require use of special retrieval tools. These assemblies may not be readily available for loading until later in the plant life which may affect intended cask loading date
 - Supplemental cooling – During onsite fuel transfer supplemental cooling maybe required if cask decay heat exceeds decay heat limits specified in the CoC.
 - Supplemental shielding – Supplemental shielding maybe required to protect workers during cask loading. Due to an insulative effect of the supplemental shielding blankets or bricks these may further restrict decay heat limits on loaded assemblies.

To transport SNF off-site, the DPC inventory must meet a transport cask CoC acceptance criteria. These requirements are often similar to those mentioned for loading into a dry storage system above but may be more restrictive. Typical examples or transport loading criteria include:

- *Intended Transport Date* – May be affected by limits on cooling time for the purpose of meeting dose rate limits and by meeting decay heat limits that may be more restrictive in transport casks. Therefore, there is typically a delay between the time is cask placed into dry storage and time when the cask is available for transport.
- *Minimum Burnup* – Because of more restrictive requirements in transport assembly minimum burnup limit is often required to limit reactivity of assemblies in the canister. These requirements are often provided in the CoC in terms of a burnup curve i.e., minimum burnup requirement as a function of the initial enrichment.

- *Acceptable Loading Configurations for Fuel Assemblies* – transport casks may have different or more restrictive requirements on loading of damaged fuel assemblies or fuel debris than storage casks
- *Acceptable Loading Configurations for Inserts* – due to different dose limit requirements in transport, loading of activated inserts may be more restrictive in transport casks than in dry storage.

2.5.2 Fuel Loading Optimization Methodology in UNF-ST&DARDS

As outlined above, there are several inputs and requirements to the dry cask loading problem resulting in a multivariable problem. The solution should not only be able to handle loading of thousands of assemblies into dry casks, but it also should have considerations for all the licensing and operational constraints related to heat loads, dose rates, and cask loading timelines. Furthermore, the utilities plan for transport of dual-purpose canisters off-site within a certain time in dry storage. Therefore, the dry cask loading problem has three competing goals: (1) minimizing the number of casks needed, (2) reducing the initial heat loads in the cask, and (3) minimizing the storage time required before the casks can be ready for transport. This essentially makes it a minimization problem.

$$\text{Solution} = \text{minimize } [(\# \text{ of casks}), (\text{initial heat load}), (\text{transport time})] \quad (1)$$

Since multiobjective problems have competing goals, a single solution cannot fulfill all the objective functions. The cask loading problem is categorized as a bin-packing problem in which there are N number of items to be packed into as few bins as possible within the problem constraints applied. Traditional bin packing problems often used first fit or best fit methods to generate a single solution. Such methods are useful for smaller problems, but they are computationally expensive and ineffective for larger problems. Approximate dynamic algorithms like the one used for GAMMA-PC produce a set of near optimal solutions within the defined constraints while significantly reducing problem time. Using the premise of the bin packing problem, GAMMA-PC uses certain decision variables to formulate an objective solution space. The problem in Eq. (1) can now be represented as follows: $X = [x, y, t_{fill}]$, where x is a matrix of assemblies to be loaded in the cask, y is an array of casks available to load the assemblies, and t_{fill} is a continuous variable indicating the time each cask is loaded with fuel assemblies. Other loading constraints (burnup limits, decay heat limits, required loading dates etc.) that were discussed above are added to further define the problem set. With the constraints defined and the solution set outlined, GAMMA-PC runs a population-based metaheuristic algorithm in which a greedy randomized adaptive search procedure (GRASP) governs the packing and local search, and various packing methods and other heuristics are applied. At the end of the computation, a set of solutions that meets all the specified loading constraints and fulfills optimization objectives is generated. To compare the solution set from GAMMA-PC against the traditional first fit methods, a set of solutions based on the oldest, coldest method are also developed. The oldest, coldest strategy is a simple first fit approach based on loading of the oldest and coldest fuel assemblies that were first placed in the cask.

2.5.3 Data Input from UNF-ST&DARDS to Fuel Loading Optimization

UNF-ST&DARDS contains a UNF database (UDB) and uses data from this database for various analyses. The UDB stores data required to execute the dry storage loading tool, including:

- Site-specific discharge inventory (e.g., assembly ID, discharge burnup, initial enrichment, discharge date)
- Information on assemblies that are already in dry storage systems to exclude them from the loading optimization,
- Assembly-specific decay heat of discharge inventory as a function of time
- Assembly-specific dose rates as a function of time to evaluate overall cask dose limits
- Projected discharge inventory for various projection scenarios

- Decay heat of projected discharge inventory.

To generate optimized loading maps, the fuel loading optimization requires detailed information about the current and projected assemblies, the casks into which these assemblies will be loaded, the site and spent fuel pool information, and other relevant loading constraints. The UDB contains all the necessary information needed as shown in an infographic layout of the process in Figure 9.




Figure 9. UNF-ST&DARDS data input to Loading Optimization (GAMMA-PC)

3. RESULTS AND DISCUSSION

3.1 Single Site Reloading Results

To determine the potential reactivity reduction associated reorganizing the contents of each canister, the assemblies in the Zion casks were reloaded with the algorithm described in Section 2.3. The reloading algorithm was run using loading schemes that divided the population of available of intact fuel assemblies into one, two, three, and four bins, to determine the impact of distributing the highly reactive fuel across a range of DPCs. In the one-bin case, assemblies were sequentially drawn from the top and bottom of the distribution so that half of the first canister was filled with the most reactive fuel, and the other half was filled with the least reactive fuel. The increasing number of bins resulted in a more uniform distribution of high reactivity, low reactivity, and intermediate reactivity assemblies across all of the canisters. The results of the calculations are shown in the box plots in Figure 10, and the summary statistics are provided in Table 1. The as-loaded configuration is presented as the first entry in both Figure 10 and Table 1. The following entries correspond to loading schemes using increasing numbers of bins in the assembly selection algorithm. The results show that even a single bin provides significant reduction compared to the as-loaded condition. However, the most reactive canister still has a k_{eff} of ~ 1.03 . A two-bin assembly selection results in a reduction in both the average canister reactivity and the peak canister reactivity to a k_{eff} of approximately 0.99. Therefore, a two-bin selection approach would make all DPCs loaded at Zion site disposable without a criticality concern. Increasing the number of bins used beyond two did not result in substantial canister reactivity reduction.

Upon examining the highly reactive canisters, it was found that they all appear to be loaded with damaged fuel. Because the purpose of this analysis is to show that repositioning the fuel can result in substantial reactivity reduction, further investigation is not warranted.




Figure 10. Box plots for the single site reloading analysis.

Table 1 – Summary statistics for the single site reloading results.

	As-Loaded	1 Bin	2 Bins	3 Bins	4 Bins
Mean	1.0009	0.9736	0.9666	0.9653	0.9645
Standard Deviation	0.0397	0.0219	0.0116	0.0107	0.0110
Maximum	1.0862	1.0311	0.9896	0.9913	0.9906
Minimum	0.9337	0.9485	0.9487	0.9510	0.9506

3.2 Development of ANN for canister reactivity prediction

As discussed in Section 2.4, the first step in developing an ANN that can predict canister k_{eff} is to develop the training data. The training data used here are from 1,000 canisters with known contents. The data were developed by retrieving all the PWR fuel burnups and enrichments from the UDB and then using that data to develop a KDE. From the KDE, 32,000 assembly burnups and enrichments were randomly sampled and assigned to canisters. Assembly reactivity calculations were then run for each of the assemblies, and canister criticality calculations were run for each canister. The generated assembly reactivity values, shown in Figure 11, have a mean and standard deviation of 1.1129 ± 0.0407 and range from 0.7817 to

1.4171. Qualitatively, the assembly reactivity values in Figure 11 look more center peaked than expected from a normal distribution with long skinny tails. Four sets of canister k_{eff} results are used to train ANNs and are distinct from one another based on the canister model used and the loading patterns incorporated. The canister models used are the MPC-32 and the NUHOMS DPCs and the loading used include only those with random fuel placement and those with random fuel placement and high and low reactivity fuel placement corresponding to the configurations specified in Figure 7 and Figure 8. The data for the MPC-32 calculations are presented in Figure 12 and the data for the NUHOMS canisters are shown in Figure 13. In both Figure 12 and Figure 13 the data set with only randomized loadings is presented in the upper portion of the figures and the data sets which include the randomized, maximum, and minimum reactivity configurations are shown in the lower portion of the figures.




Figure 11. Histogram of assembly reactivity values used to train the ANN.




Figure 12. Histogram of MPC-32 canister k_{eff} values used to train the ANN.




Figure 13. Histogram of the NUHOMS canister k_{eff} values used to train the ANN.

The training data were then used to develop an ANN with the assembly reactivity values as the inputs and the canister k_{eff} as an output. The ANN model that provided the best results was a three-hidden-layer network with 64 units per layer. The input layer was 32 units corresponding to the assembly reactivity of each storage location in the canister, and the output layer was a single unit corresponding to the canister k_{eff} . The data were divided into the training data with 700 casks, the validation data with 100 casks, and the testing data with 200 casks. The training data were fed to the neural network for a series of 4,000 epochs, using the scaled exponential linear unit as the activation function for all units. Figure 14 shows the training loss (blue) and the validation loss (orange). The validation error served as a metric to determine when enough epochs had been used to maximize the accuracy of the model. The data in Figure

14 indicate that the model appears to be fully converged at about 2,500 epochs. The average training error over the last 500 generations was 0.0014, and the average validation error over the last 500 epochs was 0.0030.




Figure 14. Training and validation errors for the neural network.

Once the training process was complete, the testing data set was used to quantify the accuracy of the model. This was done by making predictions of canister k_{eff} based on the assembly reactivity values of the contents of the testing data set with the finalized model and then comparing the predictions to the KENO-calculated values. The predicted k_{eff} as a function of the KENO-calculated k_{eff} is plotted in the left half of Figure 15. In general, the prediction results trend with the true k_{eff} results and are reasonably tightly scattered around the linear fit to the data. It is notable that the ANN predictions depart most from the true k_{eff} values in the higher reactivity canisters. This is both correctable and less important than having good predictions of the lower reactivity canisters. It is likely correctable by including canisters with high reactivity fuel in the center of the canister. It is also less important than having good predictions for lower reactivity canisters since they are still generally identified as being high reactivity and would need some amount of adjustment to satisfy disposal criticality concerns. The right portion of Figure 15 shows a histogram of the differences between the predicted and true k_{eff} values. These values look to be reasonably normally distributed, with a root mean square error of 0.0038. Of the 200 test canisters considered, 175 had errors less than 0.0050, and 25 had errors greater than 0.0050.




Figure 15. Predicted k_{eff} as a function of testing k_{eff} (left) and frequency of error (right) for the testing data set.

As it was covered earlier in this section, an activation function is applied to each layer. Depending on the activation functions, weights, and biases in each layer, the values for subsequent activation layers are calculated. There are multiple activation functions available in the Keras framework and the function we chose has an impact on the results. In this work two activation functions were explored; scaled exponential linear unit (SELU) and Gaussian exponential linear unit (GELU). A data set of 3,000 MPC-32 canister loading options were used representing the most reactive, least reactive, and random loading pattern. Three layers were used in the model and the density of the network was varied from 2 to 512, doubling in each iteration. The aim of the exercise was to identify the activation layer and network density that resulted in minimizing the error. Error was calculated by summing the deviation of the predicted k_{eff} from the true k_{eff} for all canisters and dividing by the number of canisters. The k_{eff} error results using the two activation functions are shown in Table 2.

It can be observed from the results in Table 2, the error while using the SELU activation model are fairly similar across a wide range of layer densities however the same cannot be said for the GELU model where the error increases by almost an order of magnitude when the layer density crosses 256. From this exercise it was realized that the GELU model outperforms the SELU model for low layer densities while the minimum error value is obtained while using the SELU model at a layer density of 16. Although this optimization work was performed, the remainder of the report uses a SELU model employing a layer density of 64. This was done to ensure that the same FY21 model parameters are used and to provide the team an opportunity to understand the holistic implications of any model changes implemented in this work.

It is also worth mentioning that since these are models and not algorithms, each time a model is run the user is bound to obtain a slightly different result. The same has been observed in our work as well where a change in the fourth decimal of the k_{eff} was observed.

Table 2 – k_{eff} error using the SELU and GELU activation models for various layer densities

Density	SELU Error	GELU Error
2	0.00278	0.00278
4	0.00256	0.00224
8	0.00231	0.00227
16	0.00212	0.00279
32	0.00252	0.00267
64	0.00220	0.00327
128	0.00247	0.00662
256	0.00214	0.02383
512	0.00251	0.08806

3.3 Impacts of training data on k_{eff} predictions

For ANN models, the predictions made are heavily dependent on the training data fed into the model. Therefore, it may be important to feed a wider range of data to the model for accurate k_{eff} predictions. To investigate this, it was decided to use two types of data sets, one with random canister loading and the other data set that includes a combination of optimal, worse, and random assembly loadings. To study the impact of this, both MPC-32 and NUHOMS-32 canisters were used. Figure 16 (a) shows the testing prediction error for MPC-32 canisters when only random assembly loading data was used during training. Figure 16 (b) shows the prediction error when the training data contained worse case, optimal, and random training data. As it can be observed, having a wide range of data training data helps in arriving at better predictions because the ANN model is better equipped to accurately predict for a wide range of canister loading scenarios. Figure 16 also suggests that it may be beneficial to have a wide range of input scenarios (in this case, assembly loading options) in the training data if we wish to minimize prediction errors. This demonstrates the general concept using both MPC-32 and NUHOMS-32 canisters it is beneficial to use a wide range of input data sets to minimize the prediction k_{eff} error. In order to investigate this effect for as-loaded inventories a validation study is presented in Section 3.4.




Figure 16. (a) Prediction error for MPC-32 canisters while using random training data; (b) error in the k_{eff} prediction while using worse, optimal, and random assembly loading scenarios for MPC-32 canisters




Figure 17. a) Prediction error for NUHOMS-32 canisters while using random training data; b) error in the k_{eff} prediction while using worse, optimal, and random assembly loading scenarios for NUHOMS-32 canisters.

3.4 Validation of the ANN against as-loaded casks

To evaluate the performance of the reactivity predicting ANN for real cask inventories the canisters available in the UNF-ST&DARDS database from DPC reactivity assessment [4] were used. To perform these calculations the canister k_{eff} values for July 1, 22000 and the assembly reactivity values for 9500 years of cooling time were extracted from the UNF-ST&DARDS database for the 32-assembly casks available. The assembly reactivities, decay times and canister cooling times were selected to be consistent with the training data for the ANN and because peak disposal reactivity date is typically in the year 22000. The 32-assembly DPCs in the database are of the MPC-32 and NUHOMS 32PT and NUHOMS 32PTH designs. For disposal reactivity calculations the NUHOMS 32PT and 32PTH DPCs are equivalent and are collectively referred to as NUHOMS-32 DPCs for the remainder of the report.

The first attempt to validate the ANN against the as-loaded results was performed during FY2021. The FY 2021 work had 227 32-assembly as-loaded criticality results available for comparison. The results of the ANN prediction (“Predicted k_{eff} ”) are plotted against the as-loaded UNF-ST&DARDS criticality calculations (“True k_{eff} ”) in Figure 18. The ANN developed during FY 2021 was trained using only randomized Westinghouse 17×17 Standard fuel assemblies, however, the as-loaded data contains a variety of fuel designs and NUHOMS-32 and MPC-32 DPCs. The results in Figure 18 show that there is a clear bias for the NUHOMS-32 DPCs and that the results for the high and low reactivity canisters may be mispredicted, though there is no apparent misprediction associated with the fuel type.

This report further investigates the performance of the ANN in comparison to as-loaded criticality calculations by expanding the number of DPCs considered and examining the sensitivity of the mispredictions to variations in the ANN training data. The number of canisters in the validation set was increased to 432 because additional 32-assembly canisters were analyzed in the FY22 DPC reactivity assessment report [6]. In addition to expanding the total number of canisters in the validation set an increased diversity of fuel types are included. Three sets of calculations were performed with the expanded validation suite: the first using the MPC-32 ANN with randomized loadings, the second used ANNs for the MPC-32 and NUHOMS-32 DPCs with randomized loadings, and the third used MPC-32 and NUHOMS-32 ANNs with randomized loadings and the high and low reactivity configurations discussed in Section 3.2. The results of the comparisons are presented in Figure 19, Figure 20, and Figure 21 for the three ANNs considered. Additionally, the bias and standard deviation of the ANNs compared to the criticality calculations are shown in Table 3 for each DPC and fuel type combination.




Figure 18. The true and predicted k_{eff} values of a wide variety of fuel assemblies and canister designs using the MPC-32 ANN model

Examining the results in the first and second column of Table 3 and Figure 19 and Figure 20 it is apparent that a dedicated model for the NUHOMS-32 design shows a visible improvement in the k_{eff} predictions. The ANN predictions for the W 17×17 , the W 15×15 and the B&W 15×15 fuels in the NUHOMS-32 DPCs have a notable decrease in the overall bias (~ 0.003 vs ~ 0.05). However, the W 14×14 , CE 14×14 , CE 16×16 and Palisades fuel types still show a significant misprediction regardless of DPC. The W 14×14 , CE 14×14 , and Palisades fuel are under predicted and the CE 16×16 fuel is over predicted. This is not surprising given the differences in the fuel geometry in comparison to the W 17×17 . The CE 14×14 and CE 16×16 fuel types have five large guide tubes compared to the Westinghouse fuel, which has many smaller guide tubes and the Palisades fuel has no guide tubes. Based on this observation, future work will incorporate fuel type in the ANN.

Looking at the results in the second and third column of Table 3 and Figure 20 and Figure 21 there is no apparent improvement to the results associated with including the high and low reactivity configurations with models. These results will allow for reduced computational expenditure when generating additional models for future ANNs.

Table 3 – Bias and standard deviation of the comparison of the ANNs with the 432 DPC validation set.

Canister Type	Fuel Type	MPC-32 Random	MPC-32 and NUHOMS Random	MPC-32 and NUHOM Full
MPC-32	W 17 x 17	0.0003 ± 0.0063	0.0003 ± 0.0063	0.0018 ± 0.0064
MPC-32	CE 16 x 16	-0.0259 ± 0.0021	-0.0259 ± 0.0021	-0.0255 ± 0.0031
MPC-32	W 15 x 15	0.0048 ± 0.0130	0.0048 ± 0.0130	0.0064 ± 0.0143
NUHOMS-32	CE 14 x 14	-0.0191 ± 0.0164	0.0236 ± 0.0165	0.0201 ± 0.0144
NUHOMS-32	CE 16 x 16	-0.0841 ± 0.0055	-0.0385 ± 0.0060	-0.0378 ± 0.0045
NUHOMS-32	W 17 x 17	-0.0513 ± 0.0074	0.0032 ± 0.0087	0.0011 ± 0.0043
NUHOMS-32	W 15 x 15	-0.0478 ± 0.0022	0.0027 ± 0.0037	0.0031 ± 0.0038
NUHOMS-32	B&W 15 x	-0.04838 ± 0.0118	0.0034 ± 0.0046	0.0031 ± 0.0090
NUHOMS-32	W 14 x 14	-0.04219 ± 0.0025	0.0122 ± 0.0026	0.0114 ± 0.0031
NUHOMS-32	Palisades	-0.01186 ± 0.0045	0.0492 ± 0.0090	0.0421 ± 0.0076




Figure 19. Comparison of the UNF-ST&DARDS calculated and Randomized MPC-32 ANN predicted canister k_{eff} values for 432 DPC validation set.




Figure 20. Comparison of the UNF-ST&DARDS calculated and Randomized MPC-32 and NUHOMS-32 ANN predicted canister k_{eff} values for 432 DPC validation set.




Figure 21. Comparison of the UNF-ST&DARDS calculated and Expanded MPC-32 and NUHOMS-32 ANN predicted canister k_{eff} values for 432 DPC validation set.

3.5 Damaged Fuel Predictions

There is a recognized mismatch between the assembly reactivities calculated for damaged fuel and the modeling of these assemblies in the as-loaded criticality calculations. Currently the assembly reactivity calculations consider the assembly burnup and intact geometry of the assemblies regardless of damaged fuel designation, while the canister k_{eff} calculations retain fuel vendor assumptions of fresh fuel with widened pitch for these assemblies. For some canisters, it was observed that there is a relatively large difference between the true and predicted k_{eff} values. To quickly estimate the possible impact of this mismatch the ANN predictions were rerun for canisters containing damaged fuel with assembly reactivities of 1.4 for instead of the calculated values. Figure 22 shows the true and predicted k_{eff} values for six canisters along with the canister IDs. A suffix of ‘_o’ after the canister ID indicates the old k_{eff} predictions that were made with the actual k_{inf} of the damaged fuel assembly. While a suffix of ‘_n’ after the canister id indicates that a k_{inf} of 1.4 was used for the damaged fuel assemblies. As it can be observed, using a k_{inf} of 1.4 significantly improves the k_{eff} predictions for a few canisters (i.e., GC006, GC007, and GC0011). It is also worth pointing out that a decrease in error between real and predicted k_{eff} is noticed for all six canisters indicating that using an increased assembly reactivity for damaged fuel assemblies can help improve predictions.




Figure 22. True and predicted canister k_{eff} values for damaged assemblies while using predefined k_{inf} assembly reactivity and a k_{inf} of 1.4

3.6 Fuel Loading Optimization Case Study for a PWR Nuclear Plant

This section describes a case study for a selected PWR nuclear power plant (Vogtle) and loading of spent nuclear fuel with consideration for dry storage and transport. For this study the MPC-32 is selected. This canister has capacity of 32 SNF assemblies. The dry storage overpack is a vertical system HI-STORM 100 and the HI-STAR 100 transportation cask is used. The number of assemblies including fuel projection to a proposed reactor shut down date is shown in Table 4.

Table 4 – Summary of Input Data

Number of Assemblies (including projection)	6505
Projected Shutdown Date (Unit 1/Unit 2)	2047/2049
Pool Unload Date	5 years after shutdown
Cask Capacity	32
Minimum Number of Casks	204

The optimization algorithm minimizes the number of casks used over the lifetime of the plant, calculates total cask decay heat, and estimates year of transport. The optimization tool evaluated solutions using from 204 to 215 casks. Data in Table 5 show this variation in number of casks has only a small effect on decay heat and no effect on transportability of canisters. Therefore the solution set with 204 canisters is selected for further evaluation.

Table 5 – Solution Sets Generated

No. of Casks	Cask Heat Load (kW)	Latest Year of Transport Year
204	18.06	2077
205	18.00	2077
206	17.91	2077
207	17.83	2077
208	17.76	2077
209	17.69	2077
211	17.60	2077
212	17.53	2077
213	17.46	2077
214	17.38	2077
215	17.28	2077

Once the solution set is identified the detailed data are printed for each cask. A sample of data for the first ten casks are shown in Table 6.

Table 6 – Cask Information Output from the Optimization Run.

Cask ID	Fill Ready Date	Trans. Ready Date	Total Decay Heat	Loading Date
1	2/3/2015	2/11/2029	27.17	11/10/2042
2	4/7/2015	3/3/2026	23.50	11/10/2042
3	9/15/2015	3/8/2030	25.56	11/10/2042
4	10/22/2015	2/6/2030	25.14	11/10/2042
5	3/10/2017	9/17/2024	21.47	11/10/2042
6	3/30/2017	2/1/2031	23.90	11/10/2042
7	9/8/2017	5/7/2030	22.00	11/10/2042
8	11/4/2017	9/4/2030	20.62	11/10/2042
9	12/28/2017	9/17/2024	20.16	11/10/2042
10	8/21/2018	8/5/2030	24.56	11/10/2042

Subsequently, we generated detailed information about the assemblies contained within each cask, which is shown in Table 7. In addition to the tabular form, a visual representation of a cask is also printed for each cask as shown in Figure 23.

Table 7 – Assembly Information for the Individual Casks.

Fill Ready Date: 2015-02-03	Canister Model: MPC-32
Transport Available Date: 2029-02-11	Storage Model: HI-STORM-100
Total Decay Heat (kW): 27.172325	Transport Model: HI-STAR-100
Loading Date: 2042-11-10	Inventory: 32 assemblies
ISFSI Name: Vogtle ISFSI	Cask Number: Cask 1

Location	Assembly ID	Discharge Date	Burn Up (MWD/MTU)	Decay Heat (kW)	Enrichment
1	6P45	9/14/2008	48.63	0.67	4.4
2	6P58	9/14/2008	48.858	0.669	4.4
3	6N02	9/14/2008	46.76	0.639	3.99
4	6N79	9/14/2008	48.055	0.641	4.95
5	6P63	9/14/2008	48.917	0.67	4.41
6	6P08	9/14/2008	48.991	0.677	4.2
7	5W04	4/7/2001	50.87	0.611	4.19
8	5W02	4/7/2001	50.844	0.612	4.2
9	6P42	9/14/2008	48.418	0.665	4.4
10	6P83	9/14/2008	50.102	0.684	4.79
11	6P19	9/14/2008	48.603	0.669	4.2
12	5U76	4/7/2001	47.895	0.559	4.5
13	5U71	4/7/2001	47.669	0.556	4.5
14	5U70	4/7/2001	47.878	0.56	4.51
15	5U59	4/7/2001	47.769	0.557	4.51
16	6M56	9/14/2008	49.831	0.683	4.01
17	6N60	9/14/2008	46.078	0.609	4.95
18	5U51	4/7/2001	48.021	0.561	4.5
19	5U11	4/7/2001	47.363	0.555	4.2
20	5L16	9/17/2000	25.852	0.319	2.12
21	5L14	9/17/2000	25.473	0.314	2.11
22	6P20	9/14/2008	48.388	0.665	4.21
23	6P16	9/14/2008	47.185	0.644	4.2
24	6P79	9/14/2008	50.244	0.687	4.8
25	5K87	9/17/2000	51.561	0.613	4.21
26	5K83	9/17/2000	51.454	0.612	4.21
27	6P41	9/14/2008	47.69	0.651	4.21
28	6P24	9/14/2008	48.04	0.659	4.2
29	6P60	9/14/2008	48.98	0.672	4.4
30	6N55	9/14/2008	45.602	0.602	4.95
31	5W06	4/7/2001	51.233	0.616	4.19
32	5W05	4/7/2001	51.269	0.617	4.2




Figure 23. Loading map for MPC-32 Cask 1.

As mentioned above, the GAMMA-PC optimization tool provides optimized solution and an oldest coldest packing methodology solution. Figure 24 shows that the distribution of heat loads for the optimized solution is skewed towards lower cask heat loads when compared with oldest-coldest packing methodology solution. In addition, the optimized solution has a lower number of casks in the highest heat load group (30-35 kW). Figure 25 shows a histogram of how many casks are ready for loading into the dry storage by a decade. It is shown that the optimized solution shifts cask loading to a later date thus enabling mixing older and colder fuel with more recently discharge, higher heat load assemblies, thus reducing average cask heat loads as shown in Figure 24.




Figure 24. Number of Casks in the Loading Plan by Decay Heat




Figure 25. Number of Casks Loaded per Decade.

4. CONCLUSIONS AND FUTURE WORK

The work described in this report continues the process of examining the potential for using future loading optimization as a disposal criticality suppression technique. Work conducted this year can be categorized into two areas. The first area is the refinement of the canister k_{eff} predicting ANN and the second phase investigated the storage and transportation-oriented loading algorithm.

The canister k_{eff} predicting ANN was compared to canisters considering their as-loaded configurations. A total of 432, 32-assembly canisters were considered based on the FY22 DPC reactivity assessment work. The original algorithm, developed last year, was shown to have significant biases when predicting the k_{eff} of the NUHOMS-32 DPCs. To rectify this problem an additional ANN was produced with NUHOMS-32 canister design and applied to those DPCs. The performance of the NUHOMS-32 DPC predicted k_{eff} values improved significantly with the incorporation of this model. Additionally, a wider range of DPC reactivities was considered by incorporating high and low reactivity canister loading maps in the training data. The incorporation of the additional loading maps did not significantly improve the predictive capability of the as-loaded canister configurations. Finally, adjustments to damaged fuel modeling within the ANN were investigated and showed improvement to the results for DPCs containing damaged fuel assemblies.

The second portion of work this year focused on the investigation of the GAMMA-PC loading algorithm for loading DPCs. The GAMMA-PC algorithm balances the initial loading decay heat, the number of empty spots in the canister, and the time between when the canister can be loaded and when it will satisfy the transportation requirements. A case study is provided for a PWR plant lifetime loading plan of SNF into dry storage with a subsequent transport. In this optimization study, the restrictions on loading included shielding and decay heat criteria. The results show that optimization algorithm results in more favorable distribution of fuel resulting in reduced total cask heat loads. In the future work, the results of this study could be amended by applying criticality criteria: a) in terms of a transport cask burnup credit curve or b) by incorporating the abovementioned ANN reactivity algorithm to predict reactivity canisters during cask loading optimization

Future work in this area will be focused on improving the canister k_{eff} predictive ANN algorithm, adapting the storage and transportation optimization tool for real world scenarios and coupling both algorithms such that post closure criticality can be considered in unison with other constraints. The canister k_{eff} ANN will be adapted to be applicable broader variety of fuel designs and will be broadened to work for more DPC designs. The storage and transportation algorithm should be able to directly consider personnel dose and potentially other factors.

5. REFERENCES

1. *Feasibility of Direct Disposal of Dual-Purpose Canisters: Options for Assuring Criticality Control*, EPRI Report 1016629, Electric Power Research Institute, Palo Alto, Calif., 2008.
2. *The Potential of Using Commercial Dual Purpose Canisters for Direct Disposal*, TDR-CRW-SE-000030 Rev 00, US Department of Energy, Office of Civilian Radioactive Waste Management, November 2003.
3. *Handbook of Neutron Absorber Materials for Spent Nuclear Fuel Transportation and Storage Applications*, Electric Power Research Institute, Palo Alto, Calif., 2009.
4. J. B. Clarity et al., *Dual Purpose Canister Reactivity and Groundwater Absorption Analyses*, Revision 6, ORNL/SPR-2020/1724, September 2020.
5. K.Y. Spencer, P.V. Tsvetkov, J.J. Jarrell, *Optimization of dry cask loadings for used nuclear fuel management strategies*, Prog. Nucl. Energy, 108 (2018), pp. 11-25
6. A. M. Shaw et al., *Dual Purpose Canister Reactivity and Groundwater Absorption Analyses*, Revision 8, ORNL/SPR-2022/2609, September 2022