Progress in Nuclear Energy 108 (2018) 11–25

Contents lists available at ScienceDirect

Progress in Nuclear Energy

journal homepage: www.elsevier.com/locate/pnucene

Optimization of dry cask loadings for used nuclear fuel management strategies

Kristina Yancey Spencer[a,*], Pavel V. Tsvetkov[a], Joshua J. Jarrell[b,c]

[a] Texas A&M University, 3133 TAMU, College Station, TX 77843-3133, USA
[b] Oak Ridge National Laboratory, Oak Ridge, TN, USA
[c] Idaho National Laboratory, Idaho Falls, ID, USA

ARTICLE INFO

Keywords:
Nuclear power
Used nuclear fuel
Dry storage
Waste management
Multiobjective optimization
Memetic algorithm

ABSTRACT

To address the evolving needs of dry storage, this research developed a new methodology to identify loading configurations that minimize the number of casks, their heat load, and the time at which they meet transportation requirements. The dry cask loading problem was formulated as an adaptable problem, accommodating site and cask limits in broadly-defined constraints. A new method was developed to address its complexities, embedding greedy randomized adaptive search procedures (GRASP) in a multiobjective evolutionary algorithm. It was named the GRASP-enabled adaptive multiobjective memetic algorithm with partial clustering (GAMMA-PC). GAMMA-PC was demonstrated through integration with the unified database from the Used Fuel Systems group at Oak Ridge National Laboratory to optimize simulated loading configurations for Vermont Yankee. Its performance was evaluated through comparisons to test solutions representing the oldest and coldest loading strategy. GAMMA-PC produced diverse solutions that dominated the testing set. The improvement was concentrated in the average heat load, and the third objective function was shown to be sensitive to individual assembly characteristics. This research contributes one of the first in-depth studies of the dry cask loading problem. It expands the current treatment of assembly selection over longer timeframes and meets user-defined requirements. Future work should focus on refining the objective functions. With its adaptable structure, GAMMA-PC is a promising new method for this task.

1. Introduction

Nuclear reactors produce large amounts of clean, reliable power and leave behind relatively little waste, which remains highly radioactive and thermally hot for many years. When the majority of the current reactor fleet was built, it was envisioned that either reprocessing or a long-term repository would be operational before spent fuel pools reached capacity, so pools were not designed to hold used fuel generated over the entire reactor lifetime (BRC, 2012). However, used fuel has remained on-site and in increasingly crowded pools, making additional storage capacity necessary. Consequently, the majority of nuclear power plants in the United States (U.S.) have established an Independent Spent Fuel Storage Installation (ISFSI) (NRC, 2015). To plan for both the possibility of long-term storage as well as transportation, caretakers of used fuel would benefit from considering the competing and cooperating aspects of these pathways.

In 2006, the IAEA recommended that nuclear utilities move away from a coldest and oldest-first dry cask loading strategy and toward a more diverse selection of used fuel assemblies (IAEA, 2006). Since then relatively little research has been devoted to the optimization of the loading problem. With the development of the Used Nuclear Fuel-Storage, Transportation & Disposal Analysis Resource and Data System (UNF-ST&DARDS) (Banerjee et al., 2014), which contains the most accurate information about U.S. used fuel available, a more nuanced approach is possible. Recent work has been done to develop an optimization tool for UNF-ST&DARDS that explores this recommendation and finds optimal loading patterns for metrics relevant to worker safety and transportation, accounting for current and future inventory.

The literature on existing dry cask loading programs is limited. Professional software programs aid utilities in assembly selection, but papers published about Studsvik's MARLA (Knott and Oyarzun, 2011) and the Electric Power Research Institute's CASKLOADER (EPRI, 2014) do not disclose the method used to perform their optimization. Both of the programs also only select used fuel from the current pool inventory. On the other hand, academic studies do describe their method, but no program found in the literature uses a state-of-the-art optimization methodology that would enable efficient computation in such a large search space (Žerovnik et al., 2009; Cho and Jeong, 2014). Moreover,

* Corresponding author.
E-mail address: kristina.yancey@gmail.com (K.Y. Spencer).

https://doi.org/10.1016/j.pnucene.2018.04.029
Received 15 January 2018; Received in revised form 26 April 2018; Accepted 30 April 2018
Available online 15 May 2018

they do not include regulatory requirements as constraints. MARLA and CASKLOADER are the only choices that help utilities pick assemblies that meet the regulatory requirements, but they focus on short-term goals, and it is uncertain how flexible their algorithms are.

Traditionally, this type of combinatorial problem has been solved using approximate algorithms, such as the *First-Fit* and *Best-Fit* methods (Martello and Toth, 1990). Exact algorithms exist for very small problems but become too costly as the problem size increases from medium to large packing problems. Approximate algorithms, on the other hand, produce near-optimal solutions in $O(n \log n)$ time, significantly reducing the computational burden.

The loading study performed by Žerovnik et al. used this type of approach to pack spent fuel assemblies from the Krško nuclear power plant into canisters for deep geological disposition. The inventory data for the optimization was generated using the CORD-2 package for burnup, mass, and initial material composition and the ORIGEN 2.1 code for thermal power as a function of cooling time (Žerovnik et al., 2009). The assemblies were then packed into canisters using variations of a largest-fit heuristic constrained by a maximum thermal limit. The study compared the heuristic results to a theoretical minimum number of canisters and showed that the heuristic achieved solutions within 5% of the theoretical minimum. However, for multiobjective problems with more complex constraints, approximate algorithms work best when combined with a method suited for the multidimensional objective space.

This research makes two main contributions. First, it provides a formal optimization framework for the dry cask loading problem, covering longer timeframes than previous studies, and including regulatory limits within the problem architecture. The second is a new metaheuristic algorithm developed to handle the complexities of the dry cask loading problem. It incorporates greedy randomized adaptive search procedure (GRASP) (Feo and Resende, 1995) heuristics into an evolutionary framework based on the nondominated sorting genetic algorithm II (NSGA-II) (Deb et al., 2002) and features partial decomposition of the objective space to group similar solutions together during crossover operations. The algorithm was named the GRASP-enabled adaptive multiobjective memetic algorithm with partial clustering, or GAMMA-PC.

This paper formulates the dry cask loading problem in Sec. 2 and describes GAMMA-PC in Sec. 3. Sec. 4 compares the performance of GAMMA-PC to solutions produced by an approximate algorithm alone. Sec. 5 discusses future work needed in this area.

2. The dry cask loading problem

The dry cask loading problem fits within the mathematical class of bin packing problems, a type of NP-hard combinatorial problem (Anily et al., 1994). The general goal of bin packing problems is to pack a number of items N into as few bins as possible within constraints such as weight or size limits. Applied to the dry cask loading problem, this goal minimizes the number of casks needed to store a set of used fuel assemblies.

In taking a long-term view of used fuel storage systems, this bin packing goal was combined with competing objectives to improve worker safety and the future transportability of dry casks. In total, the dry cask loading problem seeks to minimize three objectives:

1. The number of casks needed to store the fuel,
2. The average initial heat load in each cask, and
3. The length of time for the casks to meet transportation dose limits.

While the main goal of the second objective is to reduce the average heat load and thereby improve thermal performance of the casks, it also serves as a convenient proxy for the dose that workers receive during transfer procedures. Both the decay heat and the radioactivity of the fuel are functions of nuclide decay, and their magnitudes tend to decrease together (Yancey and Tsvetkov, 2014). While it is not a perfect proxy, the improvement in computational efficiency for such a large multiobjective optimization problem makes decay heat a practical subsitute for dose. By including the third objective in the initial loading of the casks, it should aid in the future management of the fuel during decommissioning, which positively correlates with personnel dose, public exposure, and accident risks (Parks, 2005).

These three objectives reflect economic, safety, and social concerns. However, they may not be applicable to all stakeholders' goals. The first contribution of this research is the framework to optimize the loading patterns. Similar to the modular structure of the methodology developed in Sec. 3, this problem can accommodate new objective functions or constraints. New objectives should be developed to align with individual user priorities and to reflect the multifaceted nature of dry storage.

The mathematical paradigm here uses the standard multiobjective problem format, given in Eq. (1).

$$\begin{aligned} &\textit{minimize} \quad \mathbf{F}(\mathbf{x}) = (f_1(\mathbf{x}), f_2(\mathbf{x}), f_3(\mathbf{x})), \\ &s.\ t. \qquad\qquad \mathbf{x} \in \Omega. \end{aligned} \tag{1}$$

Here, the loading configuration is expressed mathematically through the decision vector $\mathbf{x}$, which belongs to the feasible region Ω. The objective vector $\mathbf{F}$ translates $\mathbf{x}$ into the objective space through the three objective functions. The best solutions in Ω involve trade-offs among the competing objectives. Comparisons between solutions are made on the basis of domination, where a solution with objective vector $u = (u_1, \ldots, u_m)^T$ **dominates** another with vector $v = (v_1, \ldots, v_m)^T$ iff $\forall\ \theta \in \{1, \ldots, m\}$, $u_\theta \le v_\theta$ and $u \ne v$ (Zhou et al., 2011). A feasible solution is optimal if it is not dominated by any other solution.

2.1. The decision vector

To optimize the selection of used fuel assemblies, it is necessary to consider how many casks to use, which assemblies to place into each cask, and when to perform transfer procedures. The formal dry cask loading problem includes both the current used fuel assemblies in the spent fuel pool and those projected over the reactor lifetime. Fig. 1 illustrates this setup with a small example problem. A pool with capacity C_{pool} stores used fuel assemblies, which are moved into dry storage from $t_{fill,1}$ to $t_{fill,4}$. Each of the casks can hold C_{cask} assemblies, and casks may have empty spaces if the total number of assemblies is not divisible by C_{cask}. Casks that have already been loaded at a site are a sunk cost and are not considered during the optimization.

Equation (2) combines these variables into the general form of the decision vector $\mathbf{x}$.

$$\mathbf{x} = [x \quad y \quad t_{fill}] \tag{2}$$

This combination expands on the standard bin packing problem decision vector format with the inclusion of the continuous time variable t_{fill}, an array denoting the time each cask is loaded with used fuel assemblies. The decision variables x and y are binary matrices and follow the standard format of bin packing problems, as shown in Eqs. (3) and (4).

$$x_{ij} \in \{0,1\}, \quad \forall\ i = \{1, \ldots, \overline{M}\}, j = \{1, \ldots, N\} \tag{3}$$

$$y_i \in \{0,1\}, \ \forall\ i \in \{1, \ldots, \overline{M}\} \tag{4}$$

Here, if item j (i.e. assembly j) is assigned to bin i (i.e. cask i), then $x_{ij} = 1$. Otherwise, the variable is set to 0. For the dry cask loading problem, the column indices within x match the keys of a dictionary containing individual assembly data, which is used to ensure that the selection process does not violate cask constraints. The array y indicates which bins are in use ($y_i = 1$) and which are closed ($y_i = 0$). The constant $\overline{M}$ represents the theoretical maximum number of bins.

While many bin packing problems set the theoretical maximum at N, this formulation uses a smaller value. Loading used nuclear fuel into