Multi-Objective Optimization of Uranium Target Assembly–3: A Comparison of Genetic and Traditional Methods

Noel B. Nelson,[*,a] Andrew J. Conant,[a] Jorge Navarro,[a] Robert N. Wahlen,[b] Chad P. Denbrock,[b] and Terry L. Grimm[b]

[a] *Oak Ridge National Laboratory, Nuclear Energy and Fuel Cycle Division PO Box 2008, MS6003 Oak Ridge, TN 37831-6003*

[b] *Niowave, Inc. 1012 N Walnut Street Lansing, MI 48906*

[*]Email: nelsonnb@ornl.gov

Number of pages: 45
Number of tables: 5
Number of figures: 18

Abstract

Commonly produced as a byproduct of uranium fission, ^{99}Mo is a key medical isotope that is in high demand in the United States. An international goal is to switch from medical isotope production technologies that require highly enriched uranium to medical isotope production technologies that require only low-enriched uranium. Niowave Inc. is contributing to this goal by developing an accelerator-driven subcritical assembly called the Uranium Target Assembly (UTA). This work compares the performance of Dakota's Multi-Objective Genetic Algorithm (MOGA) against traditional sensitivity analysis in the neutronic optimization of the UTA-3 system. The design objectives are k-eigenvalue (k_{eff}) and natural uranium fission power, which are directly correlated with the amount of ^{99}Mo produced. Dakota:MOGA did not perform as well as human engineering ingenuity in optimization studies with high numbers of input parameters, such as fuel rod type selection and fuel rod placement. However, Dakota:MOGA did outperform traditional sensitivity analysis in optimization studies with fewer than 20 parameters and revealed the degree to which each parameter influences the optimal design space for k_{eff} and natural uranium fission power (to a lesser extent). As the design model became more complex in the final stage of design, the computational resources required to calculate the design objective values in the Monte Carlo N-Particle transport code from selected input parameter combinations limited Dakota:MOGA's performance, and, unfortunately, human intervention was required to discern the optimal design space. Future work will attempt to reduce computational resource constraints by incorporating a reduced-order neutronics model into the optimization cycle.

Keywords — Accelerator-Driven Systems, Subcritical Assembly, Photonuclear, Genetic Algorithm, Optimization

I. INTRODUCTION

In 2012, Congress passed the American Medical Isotopes Production Act, which "support[s] the development of non-HEU [highly enriched uranium] ^{99}Mo production technologies" and the establishment of a domestic supply of ^{99}Mo in the United States [1]. Old ^{99}Mo production technologies relied on the irradiation and fission of HEU targets in reactors, but most international suppliers now rely on low-enriched uranium (LEU) technology. Molybdenum-99 decays into ^{99m}Tc, which is a key medical isotope used every year in over "40,000 medical procedures in the United States" [2].

One of the most common ways to produce ^{99}Mo is through the fission of uranium. A by-product, ^{99}Mo is produced at a yield rate of 6% per fission. To produce ^{99}Mo, Niowave Inc. began research and development of their Uranium Target Assembly (UTA) following a favorable parametric study by Bencomo that suggested that a light-water accelerator-driven subcritical assembly (ADSA) could feasibly meet 12.25% of domestic demand [3]. An ADSA functions quite similarly to a traditional nuclear reactor. However, instead of a self-sustaining fission chain reaction being maintained, a subcritical fission chain reaction is maintained by an accelerator-driven neutron source. Several such systems have been designed for various applications, including the destruction of nuclear waste, power production, and medical isotope production ([4]; [5]; [6]; [7]).

Niowave's UTA ADSA has undergone three phases of development. The accelerator and assembly power and isotope production amounts (or targets for UTA-2 and UTA-3) of each design is shown in Table I.

UTA-1 is a licensed, deeply subcritical, small-scale demonstration facility, which can produce up to 1 mCi batches of ^{99}Mo [8]. It largely served to validate the complex lead bismuth eutectic (LBE) photoneutron source at the center of the assembly and the 2.3 W fission power.

TABLE I
Key parameters from each UTA design phase.

Parameter	UTA-1	UTA-2	UTA-3
LINAC Energy, Power	10 MeV, 10 kW	20 MeV, 10.4 kW	40 MeV, 200 kW
Thermal Power	2.3 W	230 W	330 kW
k_{eff}	0.43	0.63	0.95
Neutron Source (n/s)	7.2×10^{11}	7.0×10^{12}	2.0×10^{15}
Mass of LEU (kgU)	1.6	4.5	36
LEU enrichment (%)	6.4	9.75	9.75
Mass of NU (kgU)	4.6	13.5	60.0
^{99}Mo Production (Ci/week)	1 mCi/batch	10	2,000

UTA-2 was designed to be a larger-scale pilot facility loaded with uranium for more extensive testing, but it has not yet been constructed [9]. Several important lessons were learned from the UTA-2 core optimization and radiation shielding design. First, optimal fuel enrichment (9.75%), LBE target thickness (3 cm), and location (near the center), and assembly configuration trends were determined via a traditional parameter search. Second, substantial radiation shielding would be required for UTA-2 (especially behind the assembly, directly opposite of the accelerator beamline). Third and finally, the current state of the art stochastic radiation transport models (SCALE and MCNP) were ill-suited to handle deep penetration radiation shielding calculations involving photonuclear physics. MCNP does not have a good adjoint-driven importance sampling variance reduction method, and SCALE/MAVRIC does not have a photonuclear physics model. An initial reduced order model of the LBE photoneutron source was developed [9] and has been greatly improved upon since by Wilkerson [10] during the design of UTA-3.

The present work focuses on the core design of UTA-3, which is intended to be a full-scale pilot facility. The general concept of how UTA-3 functions is shown in Fig. 1.

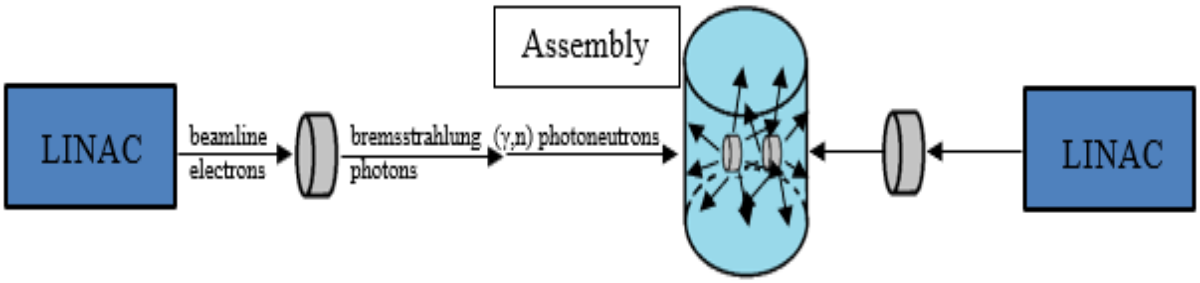

Fig. 1. Diagram of Niowave's photonuclear-driven subcritical assembly concept.

As shown in Fig. 1, two 250 kW continuous-wave linacs emit steady beams of 40 MeV electrons at heavy-metal targets. Those heavy-metal targets scatter the electron beams from the accelerator, producing bremsstrahlung. The bremsstrahlung photons created are well above the photonuclear cross-section thresholds for the target (6–10 MeV for most materials) [11] and produce a small amount of neutrons (i.e., 3% of the original electron beam's intensity). These neutrons then decelerate and are eventually absorbed within the assembly or the environment surrounding the assembly (i.e., leakage), as in a traditional light-water reactor (LWR). Some of these neutrons are absorbed within the uranium fuel pins, causing fission and thereby multiplying the neutron population and further amplifying fission rates within the core. Unlike a typical LWR, the target design maintains a k-eigenvalue (k_{eff}) of 0.97 or less instead of a k_{eff} close to 1. This constraint is a soft target for the majority of the study, but becomes a hard constraint in the final phase of design (Stage 4).

Prior studies have shown that a few Dakota optimization methods could assist with the development of the assembly design [12] and offer further insight into traditional methods [13]. Expanding upon those studies, a direct comparison between Monte Carlo N-Particle (MCNP) parameter studies and Dakota-assisted optimization was performed using the Multi-Objective Genetic Algorithm (MOGA). The primary goals of this work were to develop an assembly design for UTA-3 that maximizes k_{eff} (i.e., neutron multiplication) up to 0.97 (ideally) and fission power in natural uranium (NU) fuel rods (a key difference from previous design iterations) and to compare the performance of Dakota's MOGA [14] with that of a traditional design optimization based on sensitivity studies.

In theory, k_{eff} is positively correlated with fission power, and so, early phases of the design only optimizedk_{eff}. However, this assumption has not been verified for a light-water moderated, mixed LEU-NU ADSA. Later phases of the design sought to test this assumption by adding NU fission fraction as an optimization metric, however this added an additional assumption, that the NU fission fraction would be proportional to the NU fission power (in order reduce the number of fixed source radiation transport calculations). However, NU fission fraction proved to be a poor choice of an optimization metric, and those results were only left in this work to further comparisons between the MCNP parameter search and Dakota:MOGA. The final phase of the study confirmed that there is no substitute for the NU fission power and that it is indeed positively correlated with

k_{eff}. For further details, please see Section III.

Prior UTA studies have also focused on optimizing ADSA designs that use a pure LEU core. Because of stricter licensing regulations for facilities that manufacture and/or reprocess LEU fuel, Niowave has decided to focus solely on producing and extracting ^{99}Mo from NU fuel rods instead of LEU fuel rods. The current UTA-3 core design still contains both types of rods (NU and LEU) so that a near-critical state can be reached without the need for heavy water, but only the NU rods will be removed for isotopic extraction.

II. METHODOLOGY

The design of UTA-3 was optimized using two separate methods that are compared throughout this report. The first method was a traditional sensitivity study of various combinations of the core parameters. In this method, MCNP transport code version 6.2 [15] was used to calculate the quantities of interest (k_{eff} and the total fission rate within the NU rods) for each combination of parameters. The second method also used MCNP to calculate the objective quantities, but Dakota:MOGA chose and optimized the parameters automatically. As a reminder, the objective of both methods was to maximize k_{eff} up to 0.97 and to maximize the NU fission power, P^{NU}, meaning that the objective function was effectively the Boltzmann radiation transport equations.

II.A. OBJECTIVE FUNCTION: THE COUPLED BOLTZMANN RADIATION TRANSPORT EQUATIONS

As discussed in the previous section, the UTA-3 core operates on the principles of particle kinetics and fission nuclear reactions to produce ^{99}Mo in a manner similar to traditional nuclear reactors. Unlike traditional reactors that depend solely on neutron physics, UTA-3 depends on additional electron and photon physics from the linacs and photonuclear sources to induce fission in the subcritical assembly. The particle balance of the UTA-3 system can be adequately described by three coupled fixed-source Boltzmann particle transport equations. The UTA-3 system's electrons are represented in Eq. (1), its photons are represented in Eq. (2), and its neutrons are represented in Eq. (3):

$$\hat{\Omega}\cdot\nabla\psi_e(\vec{r},E,\hat{\Omega})+\sigma_t\psi_e(\vec{r},E,\hat{\Omega})=\int_0^\infty\int_0^{2\pi}\int_{-1}^{1}\sigma_s(E'\to E,\mu_0)\ \times\ \psi_e(\vec{r},\mu',\varphi',E')d\mu'd\varphi'dE'+$$
$$+\frac{\alpha}{2}\{\frac{\partial}{\partial\mu}[(1-\mu^2)\frac{\partial\psi_e(\vec{r},E,\hat{\Omega})}{\partial\mu}]+\frac{1}{1-\mu^2}\frac{\partial^2\psi_e(\vec{r},E,\hat{\Omega})}{\partial\varphi^2}\}+\frac{\partial}{\partial E}[S\psi_e(\vec{r},e,\hat{\Omega})]+Q_e(\vec{r},E,\hat{\Omega}),\ \vec{r}\in V, \tag{1}$$

$$\hat{\Omega}\cdot\nabla\psi_p(\vec{r},E,\hat{\Omega})+\mu_t(\vec{r},E)\psi_p(\vec{r},E,\hat{\Omega})=\int_0^\infty\int_{4\pi}\mu_s(E'\to E,\hat{\Omega}'\to\hat{\Omega})\psi_p(\vec{r},E',\hat{\Omega}')d\hat{\Omega}'dE'+$$
$$+R_{pf}(\vec{r},E,\hat{\Omega})+Q_p(\vec{r},E,\hat{\Omega}),\vec{r}\in V, \tag{2}$$

$$\hat{\Omega}\cdot\nabla\psi_n(\vec{r},E,\hat{\Omega})+\Sigma_t(\vec{r},E)\psi_n(\vec{r},E,\hat{\Omega})=\int_0^\infty\int_{4\pi}\Sigma_s(E'\to E,\hat{\Omega}'\to\hat{\Omega})\psi_n(\vec{r},E',\hat{\Omega}')d\hat{\Omega}'dE'+$$
$$+\chi(E)\int_0^\infty\nu_n(E')\Sigma_f(\vec{r},E')\phi_n(\vec{r},E')dE'+Q_n(\vec{r},E,\hat{\Omega}),\vec{r}\in V, \tag{3}$$

where

$$\mu_0=\mu'\mu+[(1-\mu'^2)(1-\mu^2)]^{1/2}cos(\varphi'-\varphi) \tag{4}$$

with appropriate vacuum boundary conditions,

$$\psi_e(\vec{r},E,\hat{\Omega})=\psi_p(\vec{r},E,\hat{\Omega})=\psi_n(\vec{r},E,\hat{\Omega})=0,\ \hat{n}\cdot\hat{\Omega}<0,\ \vec{r}\in\Gamma, \tag{5}$$

and where

$\hat{\Omega}$ = 3D unit vector

ψ_e = electron angular flux (cm^{-2} s^{-1} sr^{-1} eV^{-1})

$\vec{r}$ = coordinate point in 3D space

E = energy variable (eV)

V = volume of the system (cm^3)

Γ = outer surface of the system (cm^2)

σ_t = smooth-component total electron cross section (cm^{-1})

σ_s = smooth-component scattering electron cross section (cm^{-1})

μ = cosine of the directional polar angle

φ = directional azimuthal angle (rad)

α = restricted momentum transfer

S = restricted stopping power

Q_e = electron source distribution (cm^{-3} s^{-1} sr^{-1} eV^{-1})

ψ_p = photon angular flux (cm^{-2} s^{-1} sr^{-1} eV^{-1})

μ_t = linear attenuation coefficient (cm^{-1})

μ_s = photon scattering coefficient (cm^{-1})

R_{pf} = photon gain rate term from photofission reactions (cm^{-3} s^{-1} sr^{-1} eV^{-1})

Q_p = photon source distribution (includes bremsstrahlung and neutron fission gammas) (cm^{-3} s^{-1} sr^{-1} eV^{-1})

ψ_n = neutron angular flux (cm^{-2} s^{-1} sr^{-1} eV^{-1})

Σ_t = total macroscopic neutron cross section (cm^{-1})

Σ_s = scattering macroscopic neutron cross section (cm^{-1})

χ = neutron fission energy spectrum

ν_n = average number of neutrons produced per fission

Σ_f = fission macroscopic neutron cross section (cm^{-1})

ϕ_n = neutron scalar flux (cm^{-2} s^{-1} eV^{-1})

Q_n = neutron source distribution (includes photoneutrons) (cm^{-3} s^{-1} sr^{-1} eV^{-1}) ([16]; [17]; [18]).

Additionally, the constant electron source strength, $Q_e(\vec{r}, E, \hat{\Omega})$, is simply the accelerator power divided by the beam energy, or $\frac{P\,[kW]}{E[J]}$. The Monte Carlo method, specifically MCNP 6.2, can be used to solve for neutron scalar flux, ϕ_n, in Eqs. (1)–(3) [15]. Consequently, the first optimization

objective, NU fission power, can be determined as follows:

$$P^{\mathrm{NU}} = \omega_f \cdot \int_V d^3r \int_0^\infty \Sigma_f^{\mathrm{NU}}(E)\phi_n(r,E)dE, \tag{6}$$

where $\Sigma_f^{\mathrm{NU}}(E)$ is the NU macroscopic fission cross section, and ω_f is the fission energy released per fission (approximately 200 MeV).

The second UTA-3 optimization objective, k_{eff}, requires the separate solution of the eigenvalue, k, neutron transport equation [18]:

$$\hat{\Omega}\cdot\nabla\psi_n(r,E,\hat{\Omega}) + \Sigma_t(r,E)\psi_n(r,E,\hat{\Omega}) = \int_0^\infty \int_{4\pi} \Sigma_s(E'\to E, \hat{\Omega}'\to\hat{\Omega})\psi_n(r,E',\hat{\Omega}')d\hat{\Omega}'dE' + \\ + \frac{\chi(E)}{k}\int_0^\infty \nu_n(E')\Sigma_f(r,E')\phi_n(r,E')dE', \tag{7}$$

This equation can be solved via the power iteration method in MCNP's KCODE [19]. Early phases of this study (Stages 1–3) relied on the k_{eff} calculation alone and used the NU fission fraction (the NU fission rate divided by the sum of the NU fission rate and LEU fission rate) as the first objective function (instead of NU fission power). However, the NU fission fraction was later proven to be a counterproductive goal, as discussed in Section III.B.3. The per-source neutron fission rates were obtained from tallies of the same MCNP KCODE calculation that was used to determine the k_{eff}.

II.B. DAKOTA AUTOMATED OPTIMIZATION FRAMEWORK

The objective functions and their solution methods are described in Section II.A. This section focuses on the stochastic automated optimization method used in this study: Dakota's MOGA [14]. The overall optimization scheme is shown in Fig. 2 and is very similar to schemes from previous work [12], but a multi-objective algorithm is used instead of the single-objective genetic algorithm.

Design optimization searches carried out by Dakota:MOGA prior to Stage 4 did not include the fixed-source MCNP calculation of the NU fission power step (lower right-hand corner in Fig. 2). Previous UTA-3 designs were based on the optimization of the NU fission fraction (NU fission rate per source neutron divided by total fission rate per source neutron), which was determined from tallies taken during the k_{eff} calculation alone.

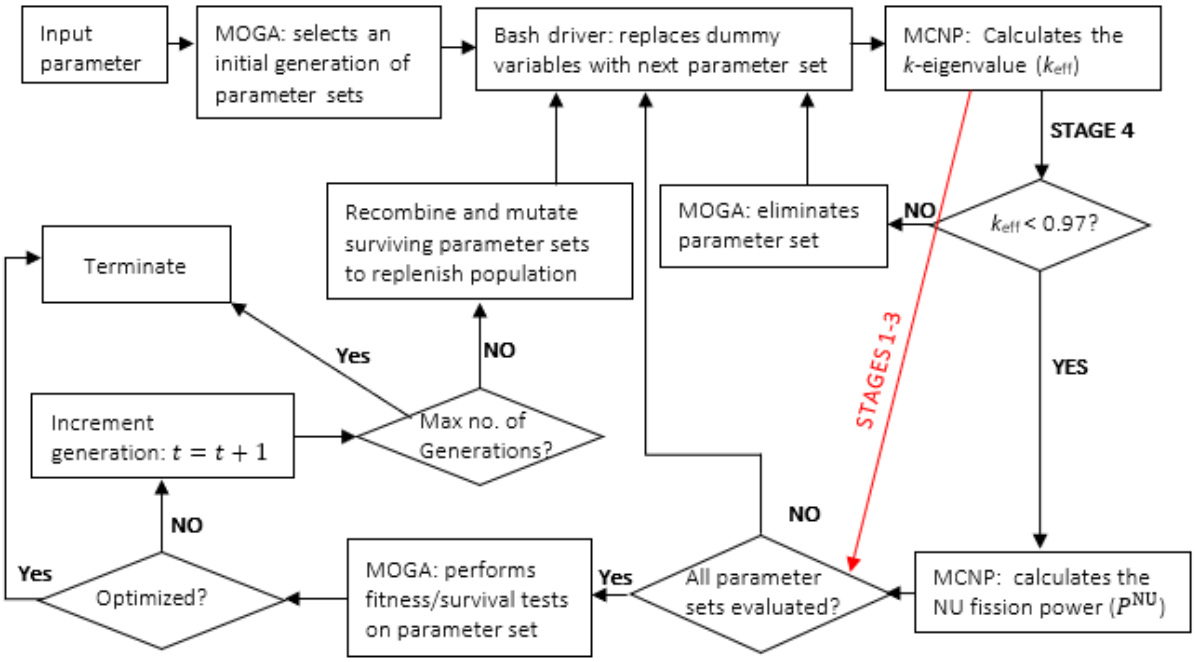

Fig. 2. Dakota optimization subroutine flowchart.

II.B.1. Genetic Algorithms

Genetic algorithms are a subclass of evolutionary algorithms and can be generally described by the following definition from Back [20] (Chapter 2, Algorithm 8):

Algorithm 1 (Outline of a Genetic Algorithm)

t:=0;

initialize $P(0) := \{\vec{a}_1(0), \ldots, \vec{a}_\mu(0)\} \in I^\mu$

where $I = \{0,1\}^l$;

evaluate $P(0) : \{\Phi(\vec{a}_1(0)), \ldots, \vec{a}_\mu(0)\}$

where $\Phi(\vec{a}_1(0)) = \delta(f(\Gamma(\vec{a}_k(0))), P(0))$;

while $(\iota(P(t)) \neq$ **true**$)$ **do**

recombine: $\vec{a}'_k(i) := r'_{\{p_c,z\}}(P(t)) \,\forall\, k \in \{1, \ldots, \mu\}$;

mutate: $\vec{a}''_k(t) := m'_{\{p_m\}}(\vec{a}'_k(t)) \,\forall\, k \in \{1, \ldots, \mu\}$;

evaluate: $P''(t) := \{\vec{a}''_1(t), \ldots, \vec{a}''_\mu(t)\}$:

$\{\Phi(\vec{a}''_1(t)), \ldots, \Phi(\vec{a}''_\mu(t)\}$ **where**

$\Phi(\vec{a}''_k(t)) = \delta(f(\Gamma(\vec{a}''_k(t))), P(t - \omega))$;

select: $P(t+1) := s(P''(t))$

where $p_s(\vec{a}_k''(t)) = \Phi(\vec{a}_k''(t))/\sum_{j=1}^{\mu}\Phi(\vec{a}_j''(t))$;

$t := t+1$;

od

where $P(t)$ is the population at generation t; $\vec{a}_\mu(t)$ is the bit-string encoded parameter vector of length l at generation t and of population size μ; γ is the function that decodes the binary bit string to base-10 parameter numbers; and δ is the scaling function, which can vary. A few examples of scaling functions include linear, logarithmic, exponential, and sigma truncation (which is weighted by the standard deviation for the set of objective function values of the population sample). Furthermore, $r'_{p_c,z}$ and m'_{p_m} (Eq. [2.82] in Back [20]) are the recombination and mutation operators, respectively. The recombination operator chosen for this work was shuffle crossover [20]. Lastly, $_l(P(t)) \neq true$ is the convergence criterion, which is typically either a minimum population diversity threshold (as shown in Eq. [2.90] and [2.91] in Back [20]) or a maximum number of generations.

II.B.2. Dakota:MOGA Tuning and Settings

MOGAs operate similarly to classical genetic algorithms but require specialized fitness functions and further adjustments to promote solution diversity [21]. For this study, the default mutation and crossover rates of 0.08 and 0.8, respectively, were selected for MOGA. Solution chromosomes are real-value encoded, and the crossover type is random shuffle. The maximum number of generations was typically 2,000, but this number was increased to 10,000 in one case (the full rod pattern search in Section III.A.2) and was decreased to 1,000 in Stage 4 (see Section III.C.1) because of computational resource and time constraints.

The stopping criterion for Dakota:MOGA was one of three conditions (whichever occurred first): (1) the average fitness falling below 10% after 10 generations, (2) reaching a convergence tolerance of 1×10^{-4} (default) in Stage 3 or of 0.06 in Stage 4 of the design stage (the tolerance is the relative change of the objective function per iteration), or (3) reaching the maximum number of generations or function evaluations (in the present case) [22]. Throughout the study, stopping criteria (3) was applied to avoid running hundreds of thousands to millions of function evaluations (each of which takes on the order of minutes), which was infeasible due to limited time and

computational resources.

II.C. UTA-3 EARLY DESIGN MODELS AND PROGRESSION

The design of UTA-3 has evolved across four phases, each with its own sets of target parameters to be optimized and studied. Fig. 3 provides an overview of the assembly models, goals, and parameters studied in each stage of development. Again, prior studies focused on the comparison of the performance of several DAKOTA algorithms on the optimization of the UTA-3 Stage 1 design [12] and the optimization of the Stage 2 and 3 design with MCNP [13]. This work largely focuses on the comparison of DAKOTA:MOGA to the prior results of the traditional parameter search [13] and the optimization of the UTA-3 Stage 3 and 4 design.

For each phase, the conditions are design constraints that are implemented and the variables are the parameters that are optimized. The optimization metrics are the objective functions to be maximized, and the outcomes are the overarching goals of the study.

The purpose of the first stage of development was largely to scale up the old UTA-2 test-scale model ([9], [12]) to a commercial production model capable of meeting Niowave's projected ^{99}Mo demand. Because of stricter regulations on processing LEU, Niowave also chose to target production of ^{99}Mo purely in NU. To counter the expected loss of reactivity by adding NU rods to a thermal light-water system, pitch and reflector materials and thicknesses were investigated through parameter studies by Niowave and Oak Ridge National Laboratory. See a conference paper that details part of these studies [13].

During Stage 2 of the UTA-3 optimization studies, the model was refined to include more realistic reflector blocks and the optimum reflector material, BeO. The study largely focused on determining the optimal rod configuration, which turned out to be a cylindrical bundle of LEU rods with an outer annulus of NU rods. Unfortunately, the large number of parameters proved too much for Dakota:MOGA to handle without a much larger number of MCNP evaluations. Further details can be found in Section III.A.

The UTA-3 reflector component was removed between Stages 2 and 3 because the marginal gains in criticality were not substantial enough for Niowave to justify the higher costs associated with constructing and maintaining this component. In Stages 1 and 3, the complex LBE target and horn structures were not included in order to simply minimize the computational costs of an

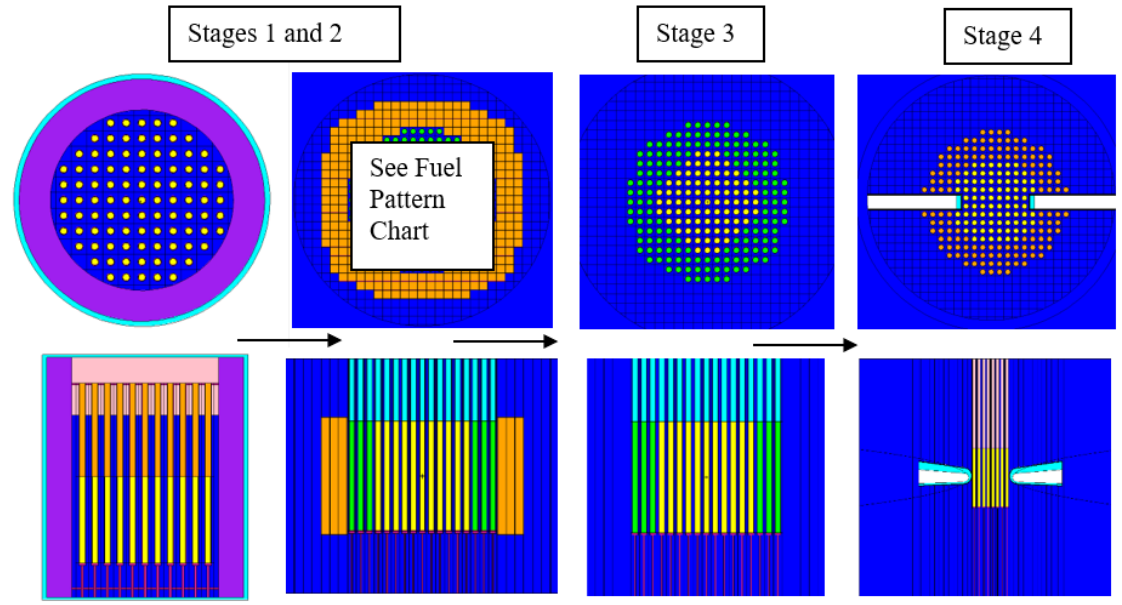

Stage 1: LEU
Goal: Gain understanding about LEU core

Conditions:
- Constant LEU mass
- No NU
- No lead-bismuth eutectic (LBE) electron converter
- Full height = core diameter

Variables:
- Reflector (Be metal, BeO, graphite, H_2O)
- Reflector thickness
- LEU pitch

Optimization Metric:
- k_{eff}

Outcomes:
- Optimal reflector material and thickness
- Optimal LEU pitch

Stage 2: LEU + NU Core
Goal: Gain understanding about core when NU is added

Conditions:
- Constant LEU and NU mass
- No electron converter
- Reflector material and thickness from Stage 1
- NU pitch = LEU pitch
- NU height = LEU height

Variables:
- Fuel rod location(s)

Optimization Metrics:
- k_{eff}
- Fission fraction in NU
- Fission in NU per LBE neutron

Outcomes:
- Optimal core pitch
- Optimal core H/D

Stage 3: LEU + NU Core (reconfigured)
Goal: Optimize the core and determine whether a unique NU pitch is beneficial

Conditions:
- LEU and NU maximum mass
- No electron converter
- Thick water reflector
- LEU pitch from Stage 2
- LEU height from Stage 2

Variables:
- LEU and NU masses
- LEU and NU pitches
- Fuel density, height, and radius

Optimization Metrics:
- k_{eff}
- Fission fraction in NU

Outcomes:
- Optimal NU configuration

Stage 4: Core + Horns
Goal: Optimize the full core with LBE horns to determine whether more LEU is needed

Conditions:
- Same conditions from Stage 3
- Optimal NU configuration (Stage 3)

Variables:
- LEU and NU masses
- LEU and NU pitches
- Fuel density, height, and radius
- Horn location (run electron transport problem)

Optimization Metrics:
- $k_{eff} \approx 0.95$
- Fission in NU per LBE neutron
- Fission power in NU

Outcomes:
- Understand whether more LEU mass is needed

Fig. 3. Four stages of the UTA-3 development cycle with example model geometry schematics.

MCNP evaluation and focus on the optimization of the assembly only. During Stage 2 rod pattern optimization, a simplified UTA-2 LBE target and beamline were placed in the initial Stage 2 MCNP model to account for neutron leakage into the beamline, which could affect the shape of the optimal rod map. The beamline and target were later removed from the Dakota optimization model merely to reduce the MCNP calculation time because a proportionally large number of MCNP evaluations were expected to be needed to balance the large number of rod location parameters.

Mainly, MCNP parameter studies were utilized in Stage 2. Fig. 4 shows a representative set of fuel patterns investigated.

Each fuel pattern was restricted to the same number of LEU and NU rods, constant NU and LEU fuel rod dimensions, and constant pitches and optimized in terms of k_{eff}. Fuel rod patterns were also briefly investigated with Dakota's single objective genetic algorithm (SOGA) under the same convergence parameter settings as MOGA, but this method was found to be impractical because it required too many generations to converge properly, even with core symmetry assumptions (as discussed in Section III.A.2). However, including the automatic rod placement algorithm described in Section II.D sufficiently reduced the dimensionality of the problem to allow reasonable optimization search convergence and execution times by Dakota:MOGA.

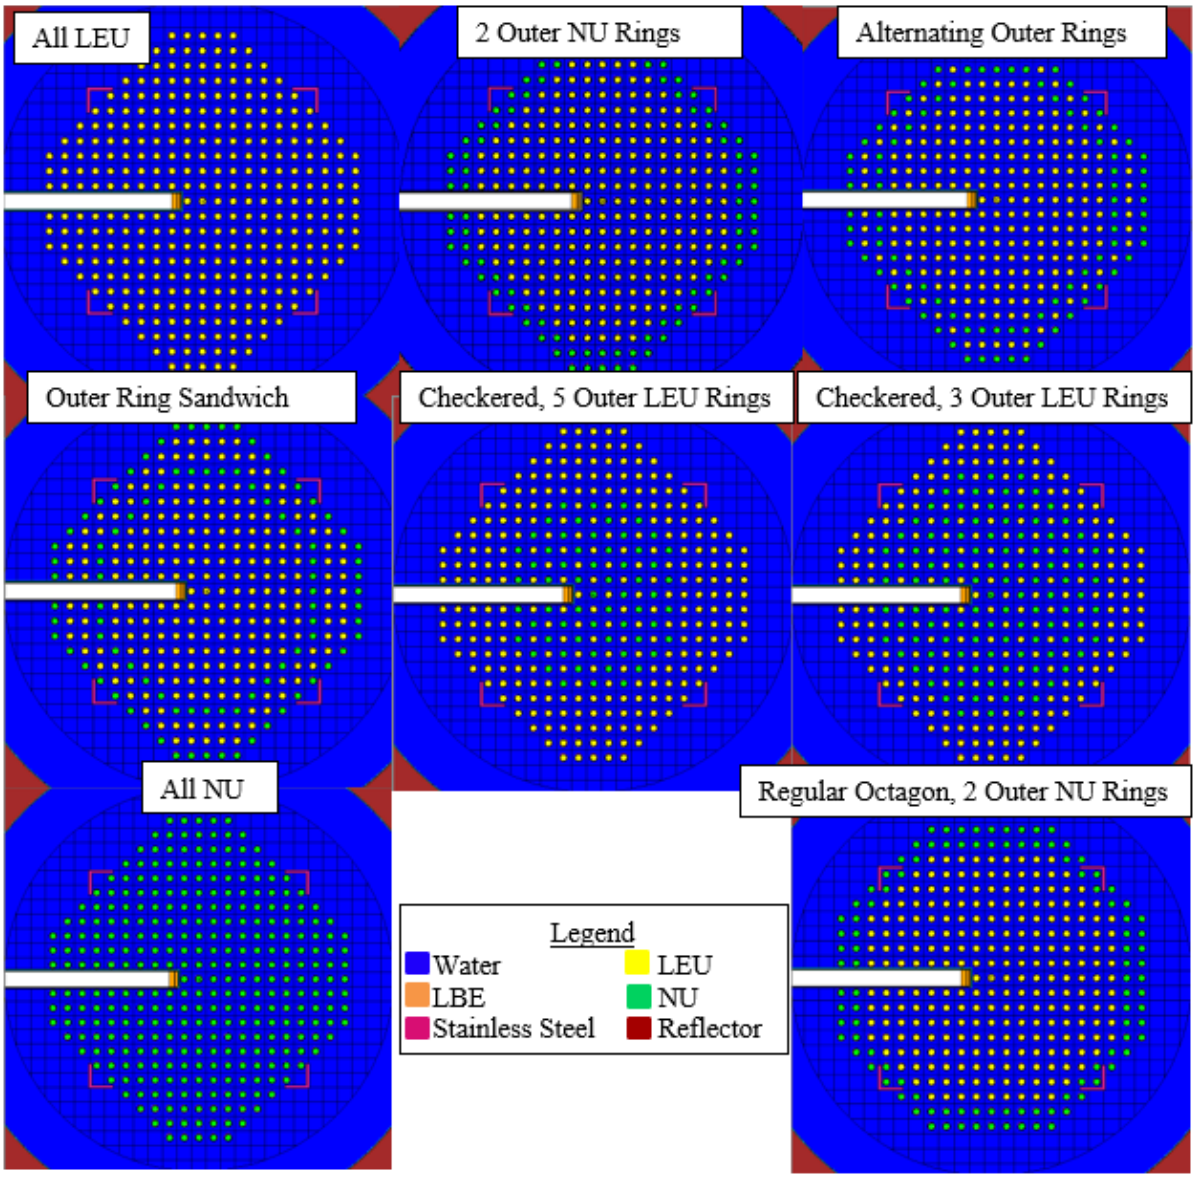

Fig. 4. UTA-3 Stage 2 fuel rod pattern chart.

To quantify the randomness of the patterns chosen by Dakota:MOGA, Moran's test was used [23]:

$$I = \frac{n \sum_{i=1}^{n} \sum_{j=1}^{n} w_{ij} z_i z_j}{S_0 \sum_{i=1}^{n} z_i^2}, \tag{8}$$

where x is the variable of interest

$$z_i = x_i - \bar{x}, \tag{9}$$

$$z_j = x_j - \bar{x}, \tag{10}$$

$$\bar{x} = \sum_{i=1}^{n} \frac{x_i}{n}, \tag{11}$$

$$S_0 = \sum_{i=1}^{n} \sum_{j=1}^{n} w_{ij}, \tag{12}$$

and the nearest-neighbor assumption is applied to the weight matrix, w_{ij}.

$$w_{ij} = \begin{cases} 1 & \text{if } x_j = x_{k\pm1,\ell} \text{ or } x_{k,\ell\pm1}, i \in 1, ...n \\ 0 & \text{otherwise,} \end{cases} \tag{13}$$

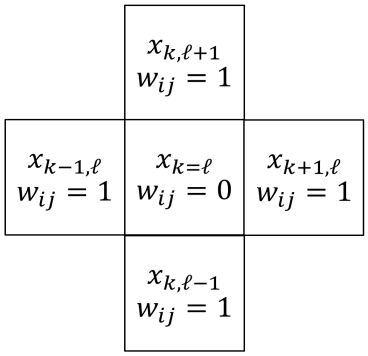

Additionally, the Dakota:SOGA rod pattern optimization search was constrained to the same total fuel mass as the human-determined patterns by adjusting density accordingly. The $n \times n$ parameter matrix shown in Fig. 5 is a map of all possible rod positions, $(R_{i,j})$ and each individual rod may contain one of four possible components.

$$\begin{array}{ccccccc}
R_{1,1}, & R_{1,2}, & R_{1,3}, & \ldots, & R_{1,n-2}, & R_{1,n-1}, & R_{1,n} \\
R_{2,1}, & \ddots & & & & & R_{2,n} \\
R_{3,1}, & & & & & & R_{3,n} \\
\vdots & & & \ddots & & & \vdots \\
R_{n-2,1}, & & & & & & R_{n-2,n} \\
R_{n-1,1}, & & & & & \ddots & R_{n-1,n} \\
R_{n,1}, & R_{n,2}, & R_{n,3}, & \ldots, & R_{n,n-2}, & R_{n,n-1}, & R_{n,n}
\end{array}$$

Fig. 5. Rod position matrix in which each rod position, $R_{i,j}$, may contain a water channel, reflector block, NU fuel rod, or LEU fuel rod.

Stage 3 focuses on optimizing the full assembly (excluding the LBE horns) under the optimal

rod configuration from Stage 2 and determining whether there is an advantage to having an NU fuel rod pitch that is larger or smaller than the LEU pitch. During this phase of the study, both MCNP parameter studies and Dakota:MOGA were fully utilized. These approaches are compared in Sections III.B.

Table I compares the target parameter ranges used in the Stage 3 MCNP grid search to those used in the Stage 3 and 4 Dakota:MOGA optimization. Other important design constants are included in Table II for reference.

TABLE I
Optimization parameter ranges listed by design stage and optimization method

Optimization Method	Stage 3		Stage 4	
	Parameter	Range	Parameter	Range
MCNP Grid Search	NU pitch (cm)	1.7 - 3.7	None	
	LEU pitch (cm)	2.5 - 3.9		
	Fuel density (g/cm^3)	2.5 - 5.5		
	Fuel height (cm)	30 - 60		
	Fuel radius (cm)	0.6 - 1.2		
	Total NU mass (kg)	60 - 61		
	Total LEU mass (kg)	36-37		
Dakota: MOGA			LBE horn x-position (column #)	-10 - 10
	NU=LEU pitch (cm)	2.5 - 5.0	Same as Stage 3	
	Fuel density (g/cm^3)	2.4 - 8.3	fuel density (g/cm^3)	2.49
	Fuel height (cm)	30 - 99	Same as Stage 3	
	Fuel radius (cm)	0.81 - 1.25	Same as Stage 3	
	Total NU mass (kg)	4 - 60	Same as Stage 3	
	Total LEU mass (kg)	4-36	Same as Stage 3	

TABLE II
Key design parameter constants.

Parameter	Value
Fuel form	U_3O_8
Fuel rod gap thickness	0.1 μm
Cladding material	Al
Cladding thickness	0.124 cm
LBE density	10.3 g/cm^3
LBE thickness*	2-3 cm
LBE height*	7.2 cm
LBE width*	7.2 cm

*The rough dimensions of the active volume of the target, whereas the actual shape of the Stage 4 LBE target is a parabaloid waterfall (see Fig. 3).

Only a Dakota:MOGA optimization study was carried out in Stage 4 in order to attempt to evaluate a larger set of parameters with the fewest MCNP evaluations (see Section III.C). The purpose of the Stage 4 study was to fully optimize the UTA-3 system, including the LBE horns, observe how the optimal parameter trends were affected from Stage 3 (if at all), and determine whether they needed to add LEU mass to UTA-3 to counter the loss in reactivity from the addition of the LBE horns.

II.D. UTA-3 AUTOMATIC FUEL ROD PLACEMENT ALGORITHM

An automated script to generate fuel placement within the square lattice was created under the primary constraint of a fixed mass of fuel for both LEU and NU. Two main physics-informed assumptions were used. The first assumption was that circular loading is the best method for preserving neutron flux economy. This assumption is believed to be accurate because a circle is the ideal shape for maximizing loading within a given area; modern commercial nuclear reactor core designs, which tend to be configured into a circular shape from a top-down view, reinforce this assumption. The second assumption was that the LEU should be placed radially inward with respect to the NU because LEU is the main driver of neutron production. The NU predominantly acts as an absorber in a light-water environment.

The fuel-loading pattern, based on optimization of this circular loading, is constrained by the total fuel mass. The number of rods, N_{rods}, to load into the core can be defined as follows:

$$N_{\text{rods}} = \frac{M_{\text{U,total}}}{w_f \rho \pi R^2 H}, \tag{14}$$

where $M_{\text{U,total}}$ is the total mass of either the LEU or NU, w_f is the weight fraction of uranium in the fuel form (e.g., U_3O_8), ρ is the fuel density, R is the fuel radius, and H is the fuel height. For this study, the fuel height and radius for the LEU and NU were the same, but future studies can examine different values.

The number of rods can then be loaded into the core. The first rod loaded is at the center—that is, a position of $i = 0$ and $j = 0$, or (0,0). To maintain an approximately circular shape, the distance to the center pin is minimized via the Pythagorean theorem:

$$\min_{\forall i,j \in \mathbb{W}} D_{i,j} = P\sqrt{i^2 + j^2}, \tag{15}$$

where D is the distance from coordinate (i, j) in the lattice and P is the pitch (assuming uniformity between LEU and NU). Fig. 6 shows this optimization. The (i, j) coordinate can also be changed to be consistent with the outer boundary of each lattice cell, although the center-to-center distance was used in this study.

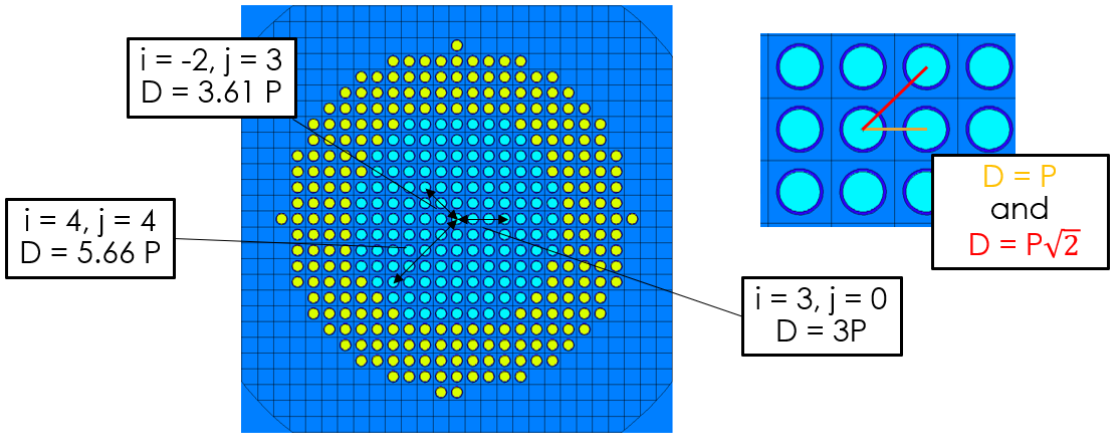

Fig. 6. Rod loading methodology of distance D from the central rod with three examples calculated from the Pythagorean theorem. In this case, LEU and NU fuel rods are cyan and yellow, respectively, and the distance, D, represents the number of pitches away based on a center-to-center distance.

III. RESULTS

III.A. STAGE 2 OPTIMIZATION STUDY

To scale up the UTA-2 model to a production scale, nearly 4× the amount of fuel was required. In Stage 1, a baseline model with pure LEU rods was studied. In Stage 2, NU rods were added purely as irradiation targets to produce ^{99}Mo. Because UTA-3 is a light-water subcritical assembly, these rods act more as neutron absorbers than as fission sources. Therefore, rod placement in the assembly requires careful planning so that fission within the rods is maximized (thereby optimizing ^{99}Mo production) and losses in neutron population and multiplication (k_{eff}) are minimized.

III.A.1. Human-Determined Fuel Rod Patterns

The authors selected a variety of fuel patterns that contained both rod types, including checkerboard and alternating rings of NU and LEU (Fig. 4 provides a visual reference). The pure NU and LEU patterns were modeled solely to determine the baseline bounding k_{eff}, and all criticality calculations were performed using MCNP [15]. One hundred fifty KCODE cycles were performed (with 20 skipped cycles), and 10,000 neutrons were sampled per cycle. ENDF/B VII.1 cross-section libraries were employed for all calculations. Table III shows the results from each pattern.

TABLE III
Calculated k_{eff} and absolute Monte Carlo uncertainties of the UTA-3 Stage 2 fuel patterns

Pattern	k_{eff}	Standard deviation (σ)
All LEU rods	0.99144	0.00062
Two outer NU rings	0.95181	0.00065
Alternating outer rings	0.94867	0.00069
Outer ring sandwich	0.93944	0.00052
Checkered, five outer LEU rings	0.87208	0.00066
Checkered, three outer LEU rings	0.81820	0.00062
All NU rods	0.29937	0.00023
Regular octagon, two outer NU rings	0.95792	0.00066

As expected, the high concentration of ^{238}U in the NU behaved more as a neutron poison or absorber than as a fission source. Therefore, concentrating any significant quantities of NU rods near the center of the core was highly detrimental to the neutron multiplicity and k_{eff}, as shown by the k_{eff} results of the checkered rod patterns. Aside from the pure NU rod baseline case (i.e., all NU

rods), the checkered rod patterns had the lowest k_{eff}. The outermost rings of the assembly were the most favorable location for the NU rods: reactivity was increased by 0.0025 over the next-highest pattern. Additionally, reshaping the core to better approximate a cylinder—as opposed to the original, more diamond-shaped pattern that was otherwise the same—created the largest increase in reactivity (nearly 0.06). The regular octagon, two outer NU rings pattern was the governing rod pattern used for all subsequent UTA-3 studies, in which only small deviations were made during the addition or subtraction of NU and LEU rods.

III.A.2. Machine-Determined Fuel Rod Patterns

As discussed in Section II.C, Dakota:SOGA was briefly employed in this phase of the design to optimize k_{eff} based on fuel rod location and make comparisons to the optimal human-informed rod pattern: regular octagonal, two outer NU rings. Again, the beamline and target were removed from the MCNP model to reduce calculation times in anticipation of the need for a large number rod combination evaluations by the genetic algorithm, as well as the observation in the human-informed pattern search that the centralization of U-235 outweighed the beamline leakage effects. Further, the authors hypothesized that the sheer number of possible rod combinations in a square arrangement would likely be too much for Dakota:SOGA to handle without a more than equivalent number of objective function evaluations to converge. For example, the 41×41 rod position map (shown in Fig. 5) would contain 1,681 positions and over 330 billion potential combinations of rods. Analyzing more than 10,000 MCNP objective function evaluations in a single Dakota:SOGA search was beyond the scope of this study.

Therefore, parameter reduction was attempted by imposing half- and quarter-core symmetry. However, even under quarter-core symmetry, the rod map was reduced to only 420 positions and approximately 1.3 billion combinations. Although this result was not encouraging, the authors still hoped that Dakota:SOGA would yield a few unique patterns that might be close to an optimal pattern. Unfortunately, it did not, as illustrated by Fig. 7.

All fuel patterns were loaded to within 3% of the maximum LEU fuel mass limit. Overall, the majority of the Dakota:SOGA-selected patterns—that is, patterns (b)–(d) in Fig. 7—appeared fairly random, and this trend was confirmed by their respective Moran's test values (between 0.32 and 0.33). In Moran's test, patterns that demonstrate high degrees of order produce values

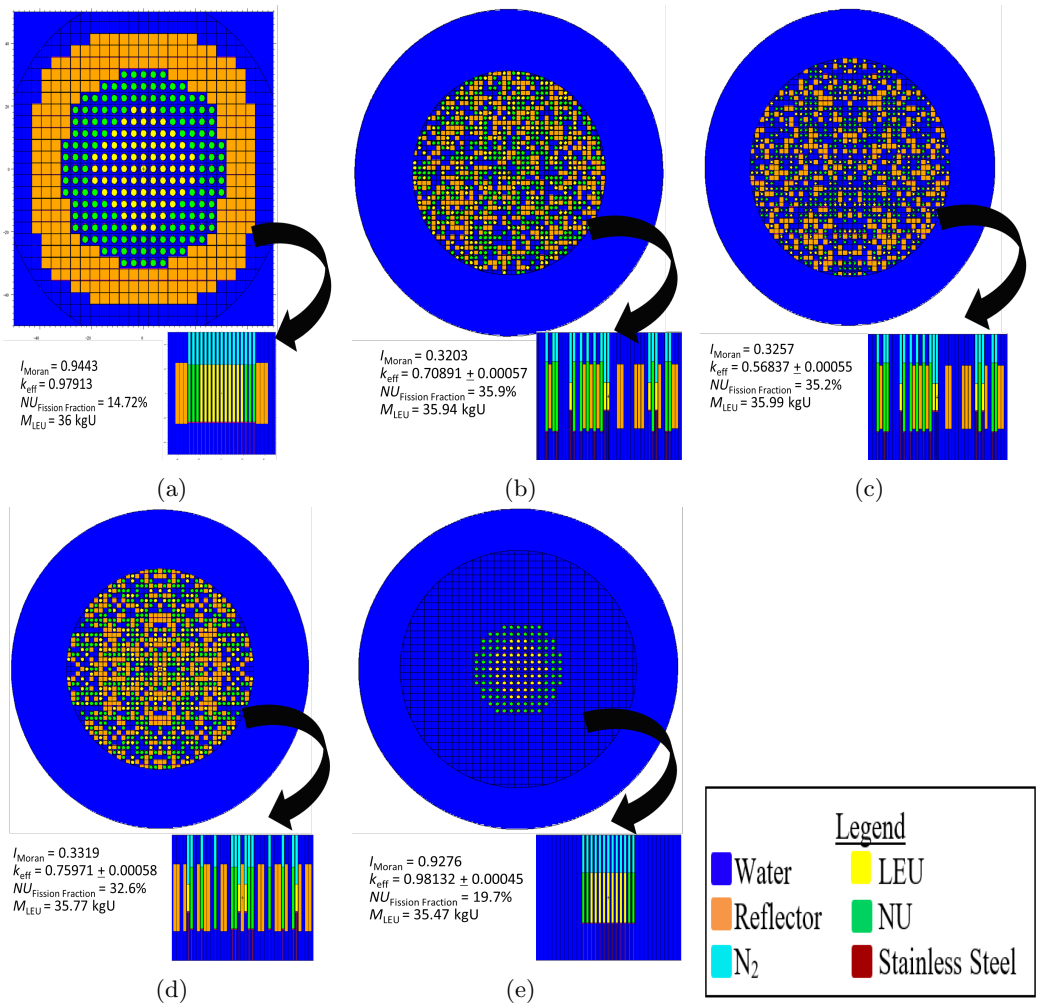

Fig. 7. UTA-3 Stage 3: (a) Niowave's baseline design pattern, (b) Dakota:SOGA full-core search, (c) Dakota:SOGA half-core symmetry, (d) Dakota:SOGA quarter-core symmetry, and (e) Dakota:SOGA with Conant's rod-loader algorithm.

close to 1, whereas patterns that demonstrate high degrees of randomness produce values close to 0. Patterns (b)–(d) were also far from optimal. The highest k_{eff} found in a search of 2,000 objective function evaluations was 0.75971, even under the quarter-core symmetry assumption. For these reasons, rod pattern searches with Dakota:SOGA were not explored further; these results demonstrated that many more function evaluations and significant computational resources would be necessary to produce viable, less random patterns.

Therefore, the optimal rod pattern (regular octagon, two outer NU rings) discussed in Section III.A.1 was implemented in all the later phases of the UTA-3 design. The rod position parameters in DAKOTA::SOGA were replaced by just two parameters: total mass of NU and LEU, and the UTA-3 automatic fuel rod placement algorithm was added to the Bash driver, which writes an MCNP input deck of the candidate UTA-3 core geometry. The rod placement algorithm calculates the total number of fuel rods available based upon mass, density, and rod volume, and then places them in a rough regular octagonal pattern starting with LEU in the center and ending with NU on the periphery. For further details on the algorithm, please see Section II.D. By switching to the octagonal pattern, Dakota:SOGA produced a much more promising result: pattern (e) in Fig. 7. Pattern (e) produced the highest k_{eff} and an NU fission fraction that was 5% higher than that of the previous best pattern, pattern (a), found by Niowave.

III.B. STAGE 3 OPTIMIZATION STUDY

The Stage 3 optimization focused on configuring the NU rods while maintaining fixed parameters for the LEU rods. Additionally, the mass of the NU was conserved.

III.B.1. MCNP P-STUDY NU Pitch Results

A selected value for a pitch can be optimized for the LEU rods, but an optimization is needed for a combination of LEU and NU rods. A primary constraint is that, other than enrichment, the fuel rod geometry and composition must be the same for the LEU and the NU. One design variable is the pitch of the NU rods. Because NU is not a multiplying medium by itself, further spacing between the NU rods decreases reactivity.

The NU pitch was optimized in MCNP using the P-STUDY capability. For this study, the authors considered the NU rods being configured solely with a rectangular pitch. However, two

unequal pitches for the LEU and NU created an incongruous geometry. The analysis was simplified by considering a square lattice of LEU rods and a surrounding annular square lattice of NU rods. Fig. 8 provides an example of two case studies. Because the lattices did not always overlap, a metric called *offset distance* was created. Offset distance is defined as the distance between the outer edge of the LEU lattice and the inner edge of the NU lattice. This value is always less than the NU pitch.

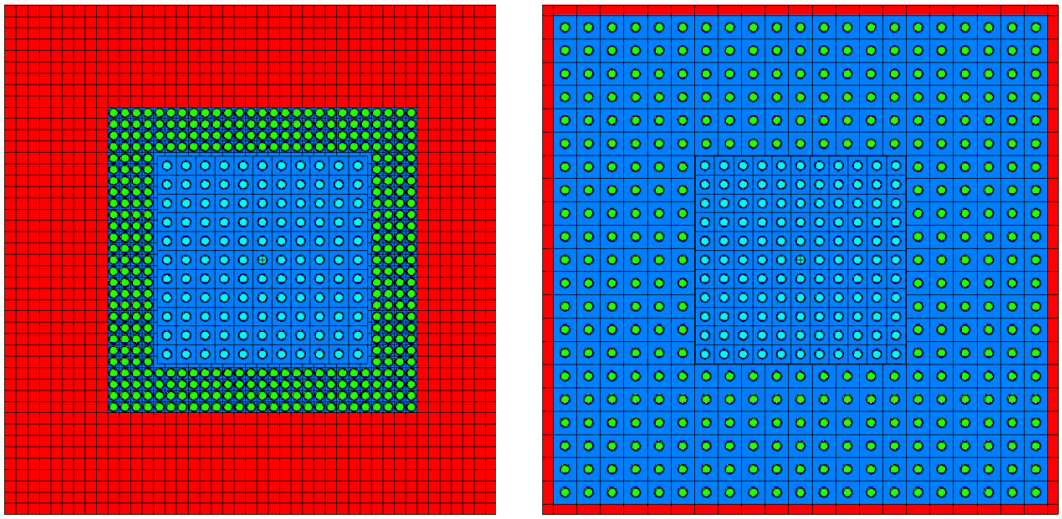

Fig. 8. Core designs in which the NU (green) and LEU (light blue) rods have separate pitches; one design has a lower NU pitch (left), and one design has a higher NU pitch (right).

Fig. 9 shows the results for the MCNP cases in which the NU pitch was varied. The k_{eff} values and NU fission fractions are shown. The k_{eff} and NU fission fraction tended to decrease as NU pitch increased. A larger offset distance resulted in fewer neutrons traversing to the NU rods from the LEU lattice and a lower k_{eff}.

III.B.2. MCNP P-STUDY Simultaneous Parameter Optimization

MCNP P-STUDY was used to generate a set of cases with varying fuel radii, heights, densities, and pitches (the same for LEU and NU in this case). The total number of cases was 2,655. Each value was bound by reasonably achievable values consistent with Niowave's fabrication as well as physical constraints. The number of rods was calculated using Eq. (14). The LEU and NU fuel

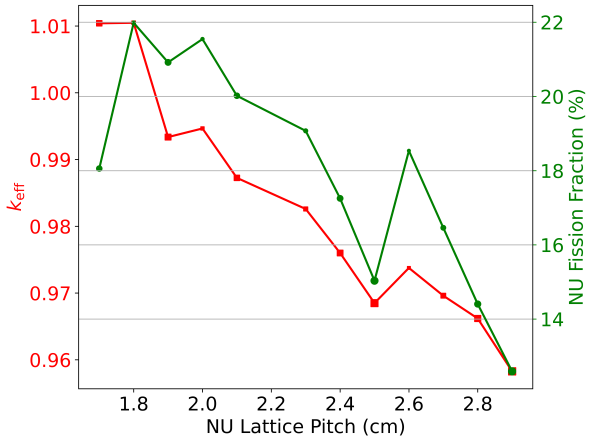

Fig. 9. NU fission fractions and k_{eff} values for various NU pitches with a fixed LEU pitch. The sizes of the markers are proportional to the offset distance. The LEU pitch was constant at 3.0 cm.

masses remained fixed for all cases, and the LEU and NU rods had the same parameters. The variations between the fuel types can be examined in a future study.

Fig. 10 shows the results of the 2,655 cases with varying geometries and material parameters. The NU fission fraction was inversely related to k_{eff} for significant NU fission production, with one exception: below 10% fission fraction, the relationship was the opposite. The inverse relationship can be explained by the increase in neutrons reaching the NU rods, which decreased the efficiency of multiplication. The exception can be explained by a subset of cases that all had low fuel densities and large pitches relative to the fuel radii. The lowest NU fission fraction achieved, for example, had the largest pitch and smallest radius. Any increase in radius or decrease in pitch increased both k_{eff} and the NU fission fraction.

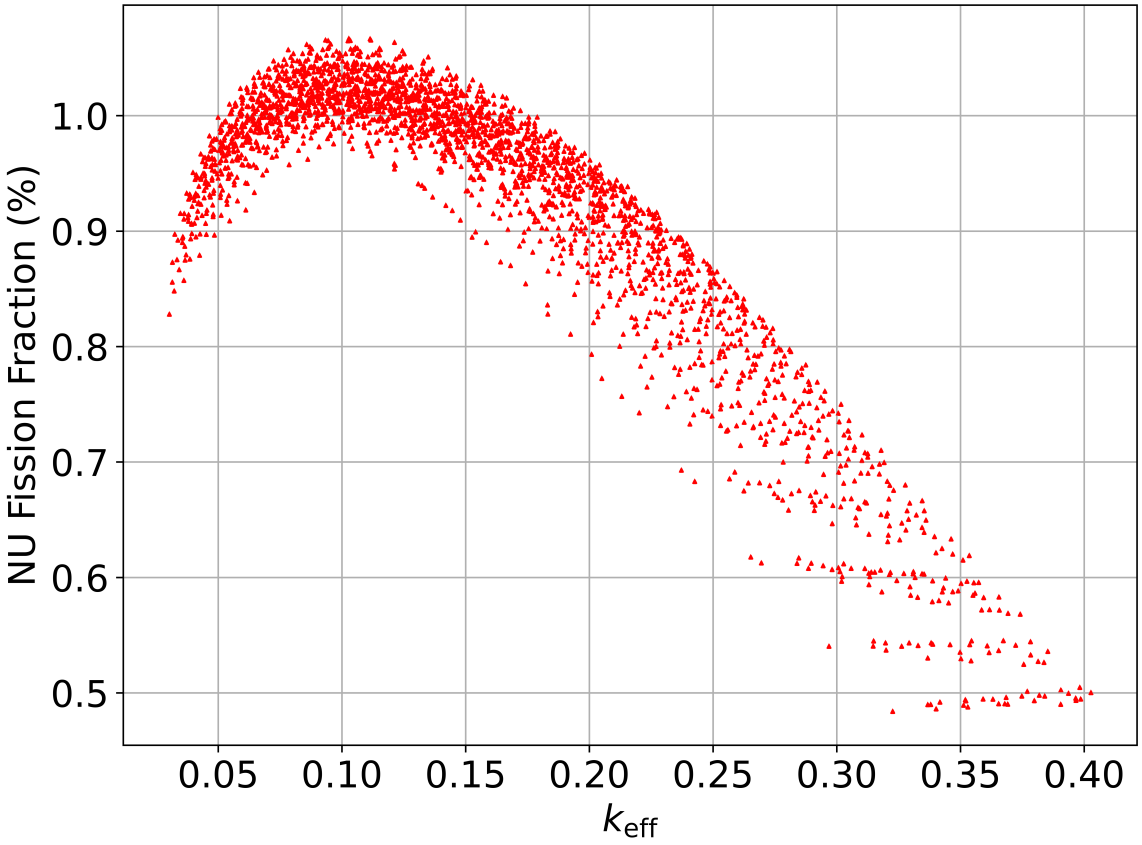

Fig. 10. Trade-off between k_{eff} and NU fission fraction for 2,655 MCNP P-STUDY cases of LEU and NU configurations with varying fuel radii, heights, densities, and pitches.

III.B.3. Machine-Determined Results

Dakota:MOGA did not reach its convergence tolerance because the default was set too low at 1×10^{-4}, whereas the objective function (k_{eff}) uncertainty was closer to 7×10^{-4}. Instead, the convergence criterion of 2,000 function (MCNP) evaluations was implemented. Dakota:MOGA performed well in mapping the optimization trends for both optimization objectives: k_{eff} and NU fission fraction. The effects of fuel rod height, pitch, and diameter on k_{eff} are shown in Fig. 11, and the effects of fuel density, pitch, and diameter on the NU fission fraction are shown in Fig. 12. Unlike in the study described in Section III.B.1, separate parameter values were not used for the NU rods.

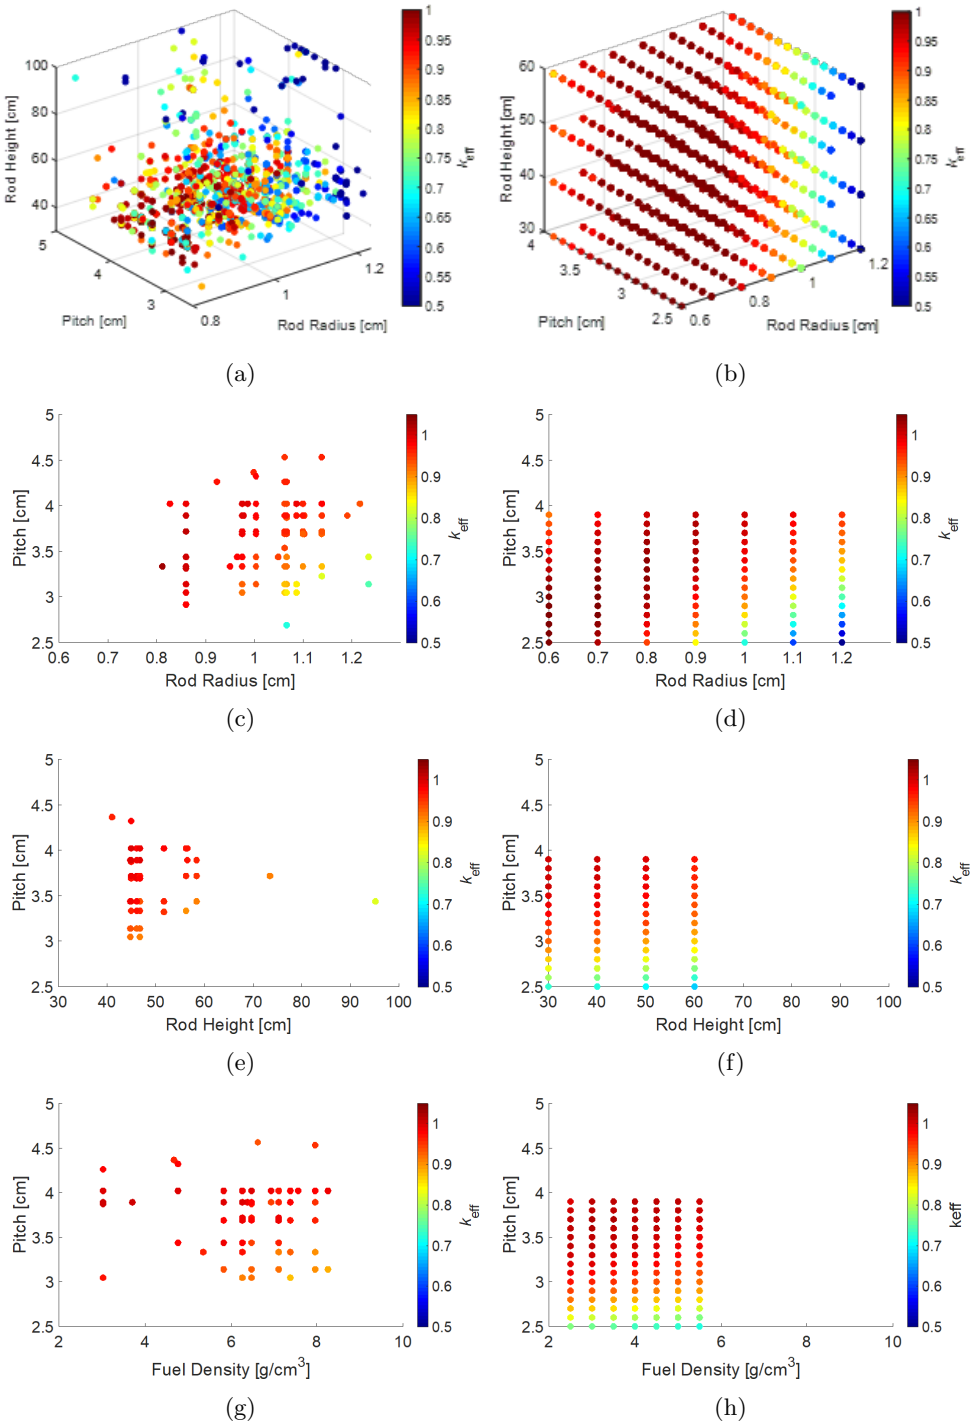

Fig. 11. Comparison of the calculated k_{eff} as a function of various parameters including fuel rod height, pitch, and radius between Dakota:MOGA (left) and the MCNP parameter study (right). For each 2D slice, the other design parameters were set to the following ranges, where applicable: $M_{LEU} = 29 - 36$ kg, $r_{\text{fuel}} = 0.95 - 1.05$ cm, $h_{\text{fuel}} = 32 - 48$ cm, and $\rho_{\text{fuel}} = 4.5 - 6.5$ g/cm^3.

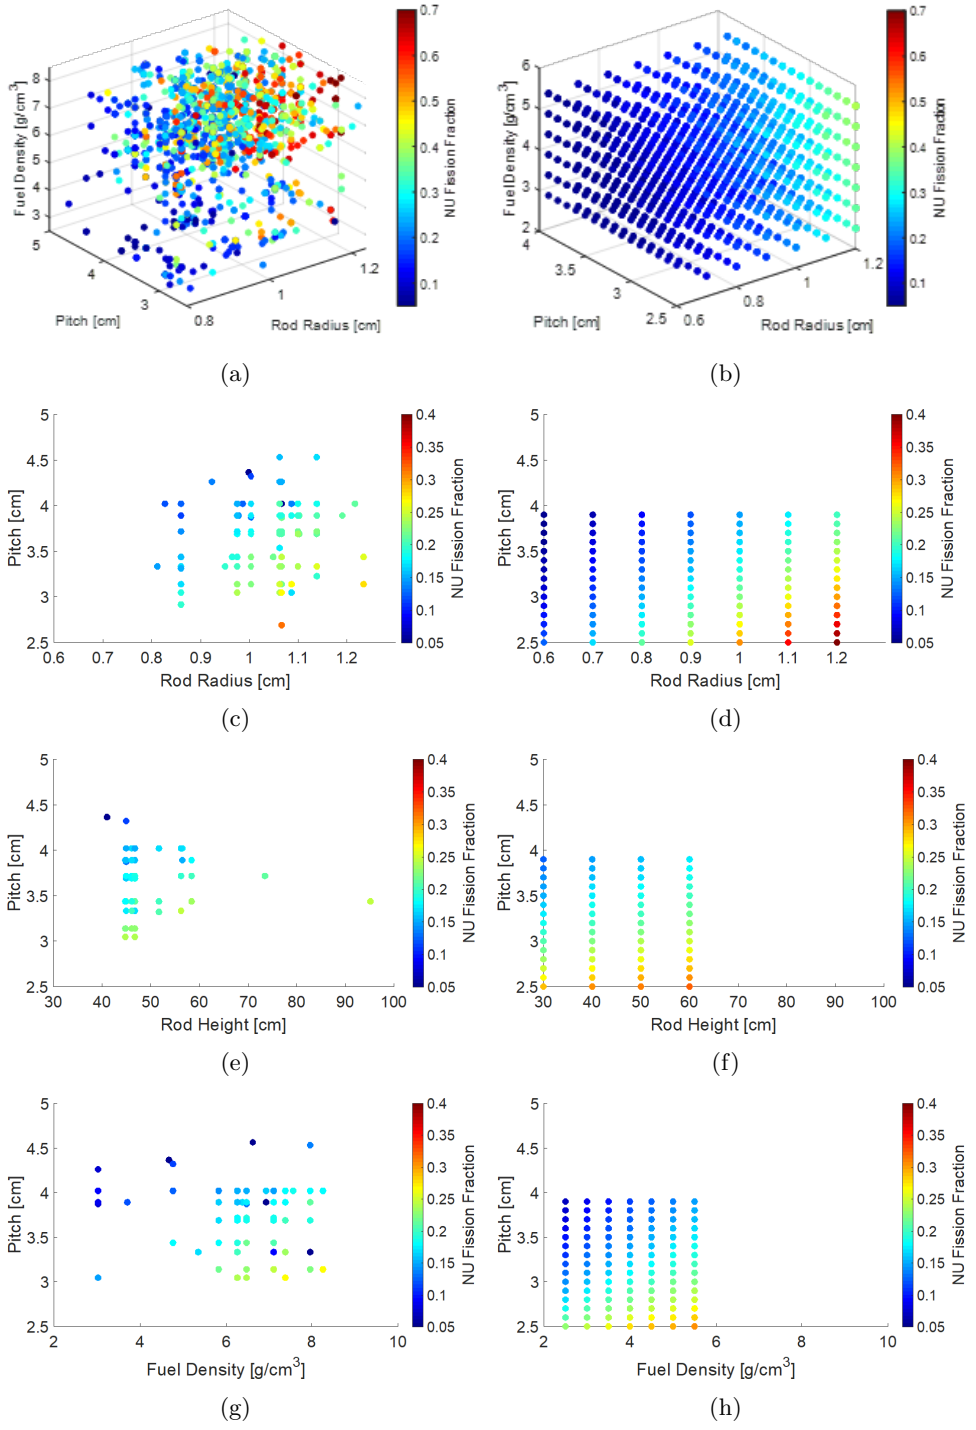

Fig. 12. Comparison of the calculated NU fission fraction as a function of various parameters including fuel rod height, pitch, and radius between Dakota:MOGA (left) and the MCNP parameter study (right). For each 2D slice, the other design parameters were set to the following ranges, where applicable: $M_{LEU} = 29 - 36$ kg, $r_{\text{fuel}} = 0.95 - 1.05$ cm, $h_{\text{fuel}} = 32 - 48$ cm, and $\rho_{\text{fuel}} = 4.5 - 6.5$ g/cm^3.

Aside from the differences of more randomized genetic sampling of the continuous parameters and a uniform discrete grid, the results from Dakota:MOGA and the parameter search were similar when the total LEU mass is held constant (see Figs. 11 (b)–(h) and 12 (b)–(h)). In Figs. 11(a) and 12(a), both the LEU and NU mass are allowed to vary, which does not affect the optimal k_{eff} region but broadens the optimal NU fission fraction region at higher densities. The broadening of the fission fraction is largely due to small LEU mass configurations that produce a high NU fission fraction and simultaneously a low k_{eff} and LEU fission rate. The k_{eff} was optimal between pitch-to-rod-diameter ratios of 1.5 and 2, and k_{eff} was relatively independent of the rod height.

When LEU and NU mass are constant (i.e., excluding sub-figure (a)), the optimal region of high NU fission fractions broadened as density increased and appeared to occur at a pitch-to-diameter ratio of less than 1:1. The same region produced a lower k_{eff} in Fig. 11. This relationship is easier to see through the Pareto front between the NU fission fraction and k_{eff} in Fig. 13.

Dakota:MOGA revealed slightly more optimal values than the MCNP parameter search above $k_{\text{eff}} = 0.9$ and demonstrated its power to explore the parameter space efficiently by revealing a significantly higher front below that threshold. Generally, the NU fission fraction decreased with increasing k_{eff}. This effect is logical because lowering the k_{eff} would lower the LEU fission rate faster than the NU fission rate because the NU sat on the periphery, whereas the LEU was the main source of neutron fission and multiplication for the assembly.

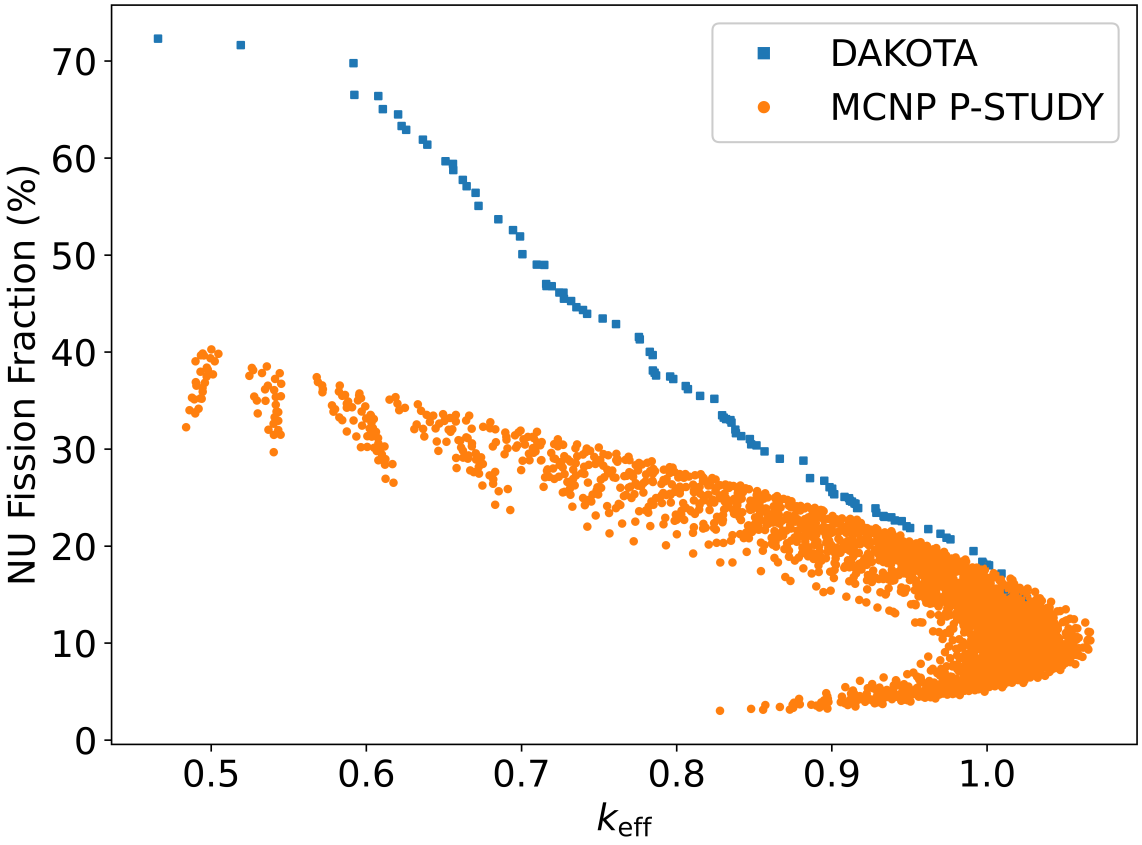

Fig. 13. Pareto front for NU fission fraction versus k_{eff}.

Clearly, having two competing objective functions was not a desirable outcome. The primary goal of the UTA-3 design is to maximize the absolute NU fission rate and thereby maximize the amount of ^{99}Mo produced within the NU rods. Maximizing the NU fission fraction is not guaranteed to maximize the overall NU fission rate. For example, the NU fission fraction can be easily increased by lowering the LEU fission rate (perhaps by lowering the LEU mass) and increasing the relative but not the absolute contribution of NU fission to the total. With this in mind, it quickly becomes apparent that NU fission fraction was a poor choice in optimization metric. As a result, the objective function was at first changed to the NU fission rate per source neutron in Stage 4 of the design because it would be independent from the LEU fission rate. Table IV shows the bounding results of the objective functions, including a point close to Niowave's UTA-3 Stage 3 target of $k_{\text{eff}} \leq 0.97$.

TABLE IV
Three bounding parameter sets from Dakota:MOGA's UTA-3 Stage 3 Pareto curve and their resulting k_{eff} and NU fission fraction values.

Fuel Density (g/cm^3)	Rod Height (cm)	Pitch (cm)	Rod Radius (cm)	Mass of NU (kg)	Mass of LEU (kg)	NU Fission[a] Fraction	k_{eff}[b]
8.26	51.7	2.95	1.234	58.8	6.44	0.7230	0.4658
7.96	46.1	3.14	0.860	58.8	35.9	0.2069	0.9765
3.02	44.8	3.33	0.860	58.8	35.9	0.0912	1.0287

[a] The maximum Monte Carlo relative uncertainty for all the Dakota:MOGA Stage 3 design optimization NU fission fraction results was 1%.

[b] The maximum Monte Carlo uncertainty for all the Dakota:MOGA Stage 3 design optimization k-eigenvalue results was 0.0007.

The optimum found at $k_{\text{eff}} = 0.9765$ was comparable to the result found during the MCNP grid search discussed in Section III.B.2. The fuel density was significantly higher than that in the grid search, but this difference may be simply due to the upper bound of the Dakota:MOGA search having been set 50% higher. The other parameter values selected by Dakota:MOGA were within 10% of the grid search optimal values.

III.C. STAGE 4 OPTIMIZATION STUDY

The primary optimization metrics for Stage 4 were the multiplication factor, k_{eff}, and the absolute fission rate per LBE neutron in the NU rods, but this was quickly changed to NU fission power after discovering that the NU fission rate per LBE neutron is not correlated with NU fission power. A sensitivity study was not performed during Stage 4 because of a large increase in the computational resources required to perform MCNP calculations on the more complex LBE dual-horn model (Fig. 3).

III.C.1. Machine-Determined Results

Although Dakota:MOGA's convergence tolerance was set to two standard deviations of the objective function with the highest uncertainty (the NU fission rate), the code still required the

full 1,000 MCNP evaluations to terminate successfully. The number of evaluations was decreased to compensate for a large increase in MCNP computational time required for the Stage 4 design model (see Fig. 14), which included two sophisticated LBE targets and their respective beamlines. Unfortunately, decreasing the number of evaluations increased the chance that the optimal values would be suboptimal, but this was tolerated because of time and computational resource constraints.

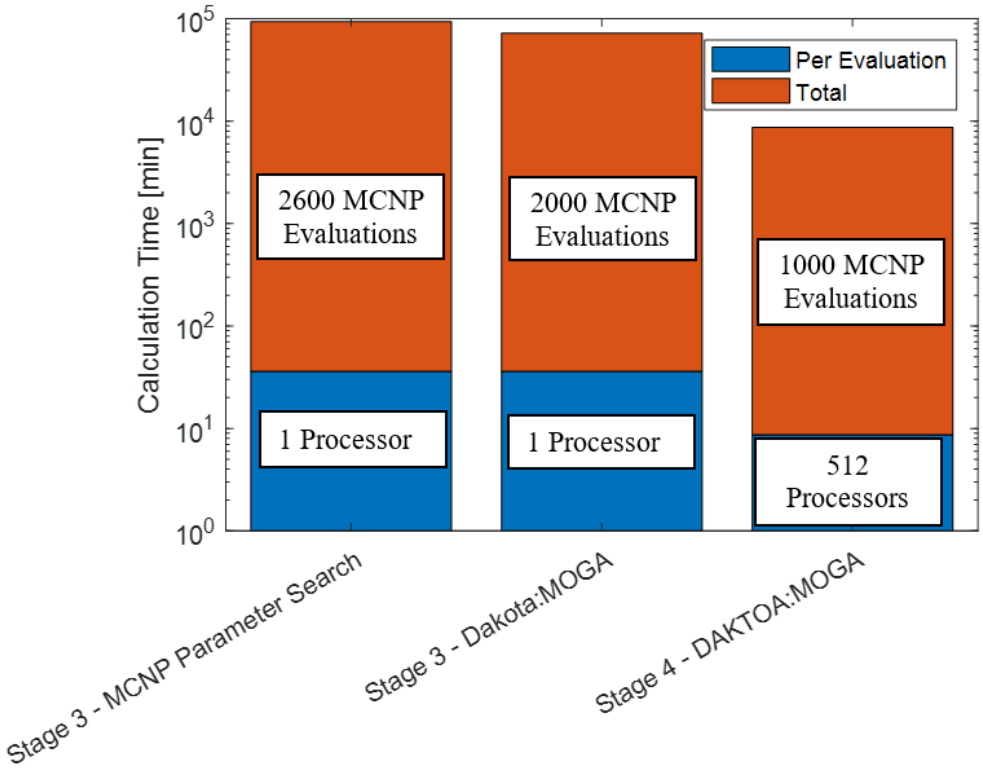

Fig. 14. Calculation times for UTA-3 Stage 3 and 4 optimizations.

The Stage 3 MCNP calculations were more efficient to parallelize by parameter set case rather than splitting up the particle histories within a single MCNP input. This was not the case for Stage 4. An equivalent single processor run for a Stage 4 parameter set would take approximately 100 times longer.

Originally, as in Stage 3, the two objective functions were optimized solely based on a k_{eff} calculation, whereas the true fission rate was calculated for a small subset of potential optima along the Pareto curve as a postprocessing step. Fig. 15 displays a Pareto curve that features the optimal points discovered in terms of k_{eff} and absolute fission rate.

Unfortunately, the postprocessing fixed-source calculations on three selected assembly con-

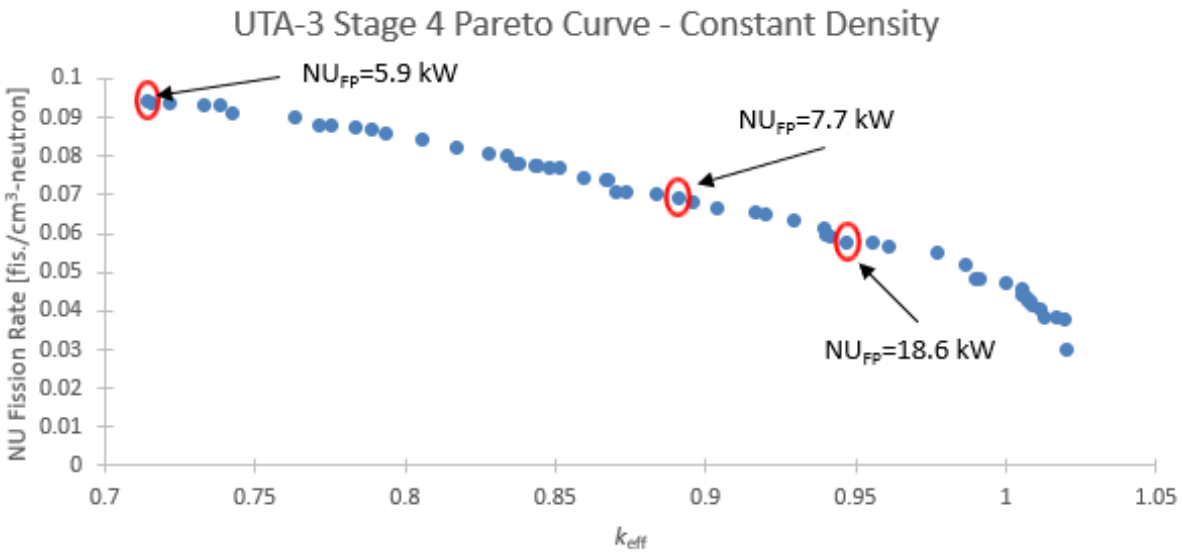

Fig. 15. Pareto front for NU fission rates per source neutron versus k_{eff}. Three points were selected for a fixed-source MCNP calculation to determine the final NU fission power.

figurations from the Pareto curve indicate that the NU power rates are not correlated with the NU fission rates per source neutron that were originally optimized. The assumption that the fission rates would scale linearly with the LBE neutron source rate was incorrect. The lack of correlation is believed to be attributable to three factors: (1) unlike in a traditional reactor, the neutron source distribution was not isotropic and symmetric, (2) the photoneutrons born outside the LBE targets (i.e., in fuel rods, water, and shielding) were not included, and (3) it was assumed that the system would be sufficiently close to a critical state for the k-eigenvalue flux distribution shape to be approximately equal to the fixed source distribution, but this may not have been true in several cases. Subsequently, fission power was difficult to predict without a fixed-source calculation, and it was logical to switch from optimizing the NU fission rate per source neutron to the fixed-source–problem fission power for the remainder of the study, adding an additional MCNP calculation to every parameter set evaluation.

Additionally, as the NU absolute fission rate per LBE source neutron was not correlated to fission power, it is highly likely that the NU fission fraction is not correlated to the NU fission power either. As a consequence, Stage 3's NU fission fraction optimization trends were not applicable to the Stage 4 design, which is why the same design parameters from Stage 3 were kept and augmented the LBE horn placement optimization in Stage 4. If nothing else, the Stage 3 NU fission

fraction results still serve as a good point of comparison between the MCNP parameter search and Dakota:MOGA methods.

Dakota:MOGA performed well in mapping the optimization trends for both k_{eff} and NU fission power (the newest optimization objective). Fig. 16 shows the effects of rod height, pitch, and diameter on k_{eff} and the NU fission rate.

The k_{eff} trends in Fig. 16(a), (c), and (e) show good agreement with the MCNP grid search, but interpreting the NU fission rate trends in Fig. 16 (b), (d), and (f) was more difficult because of sparser sampling—an unfortunate side effect of the chaotic sampling nature of genetic algorithms and of stochastic optimization methods in general. Again, the optimal pitch-to-diameter ratio was between 1.5 and 2. As shown in Fig. 16(a) and (b), the optimal range of pitch in terms of k_{eff} was between 2.75 and 4. The k_{eff} was still fairly insensitive to rod height, but an optimal range was clear between 30 and 75 cm. Predictably, k_{eff} was unaffected by the amount of NU mass but was significantly affected by the amount of LEU mass in the assembly (i.e., k_{eff} increased as the mass of LEU increased). Fig. 16(d)–(f) shows that a few discernible optimal spots are observable in the same regions of the NU fission rate. Trends in the LBE source position are not shown graphically. Because of the limited number of discrete value choices available to Dakota, large numbers of points are stacked on top of one another. The optimal positions are better understood by examining the best parameter configurations listed in Table V.

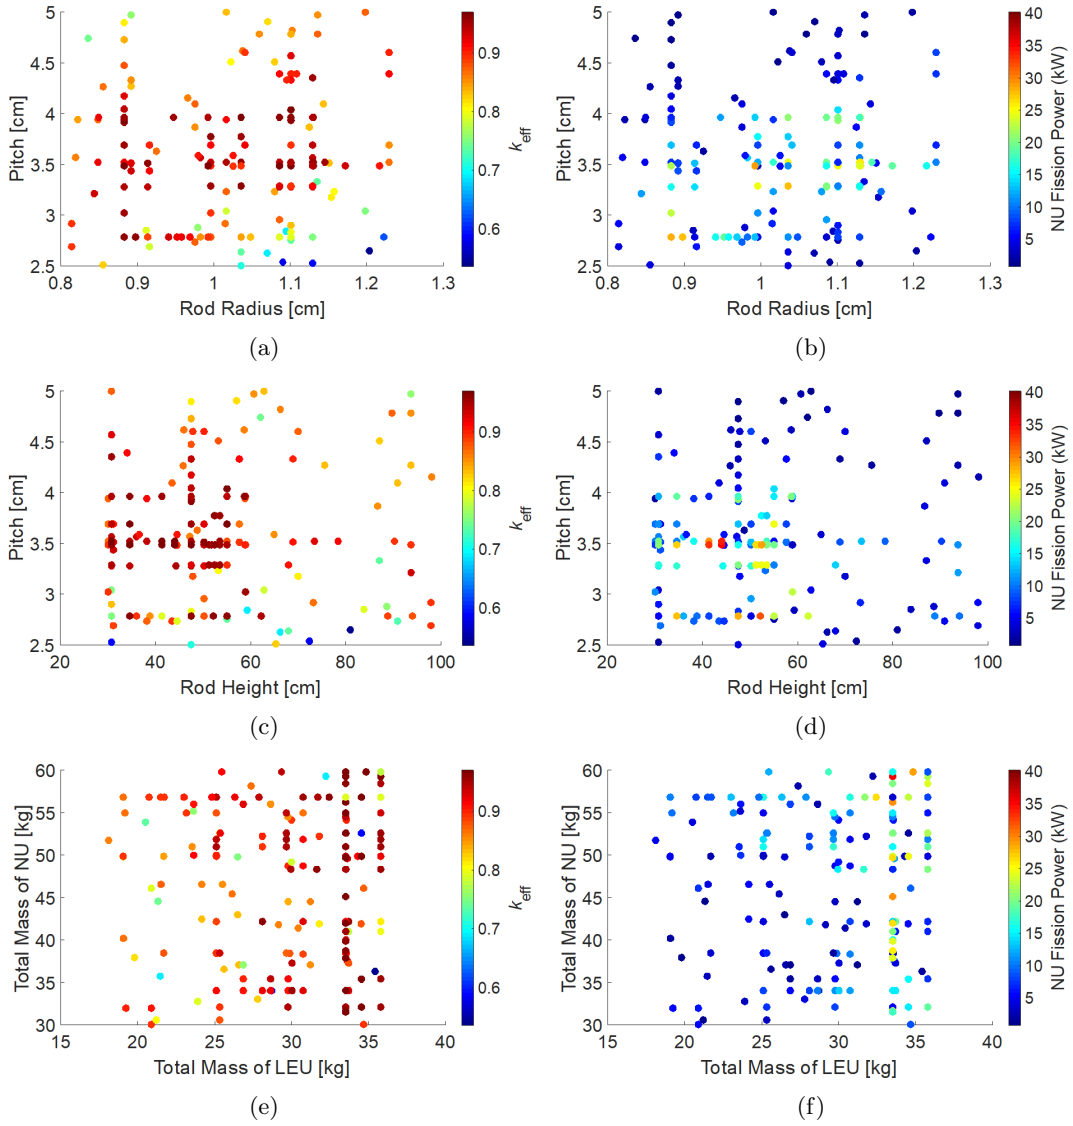

Fig. 16. Calculated k_{eff} (left) and NU fission power (right) as functions of various parameters, including fuel rod height, radius, pitch, LEU mass, and NU mass.

TABLE V
UTA-3 Stage 4 design parameter sets that were predicted by Dakota:MOGA to produce the minimum and maximum NU fission powers and the three best parameter sets near the target k-eigenvalue of 0.95

Rod Height (cm)	Pitch (cm)	Rod Radius (cm)	Mass of NU (kg)	Mass of LEU (kg)	Left LBE Pos. #	Right LBE Pos. #	NU Fission[a] Power (kW)	k_{eff}[b]
Minimum NU Fission Power								
93.7	4.78	1.061	37.1	31.2	8	10	0.789	0.8522
Highest NU Fission Power Near $k_{\text{eff}} \approx 0.95$								
47.5	2.78	0.883	56.8	35.8	2	4	27.7	0.9460
47.5	3.49	1.101	58.4	35.8	2	6	29.2	0.9556
41.3	2.78	0.883	58.4	35.8	2	6	34.5	0.9589
Top Three NU Fission Power Cases								
47.5	2.78	0.883	56.8	33.5	6	6	37.9	0.9664
47.5	2.78	0.883	59.8	33.5	6	6	39.6	0.9667
47.5	2.78	0.883	59.2	33.5	6	6	40.2	0.9668

[a] The maximum Monte Carlo relative uncertainty for all the Dakota:MOGA Stage 4 design optimization NU fission power results was 6%.

[b] The maximum Monte Carlo uncertainty for all the Dakota:MOGA Stage 4 design optimization k-eigenvalue results was 0.0007.

Although Dakota:MOGA was constrained to k_{eff} less than 0.97, values closer to 0.95 were the final design target. For k_{eff} between 0.945 and 0.96, optimal NU fission powers were found to range from 27 to 35 kW. Overall, the results indicate that the system is not too sensitive to the LBE position, but the system appears to be optimal for at least one of the targets to be near position 6, which typically happens to lie just inside the last ring of LEU fuel rods (for a fully loaded core with 36 kg of LEU).

With the updated optimization scheme involving a fixed-source and k_{eff} MCNP model evaluation, Dakota:MOGA's performance significantly improved. Fig. 17 shows a Pareto front of the optimal configurations between k_{eff} and NU fission power.

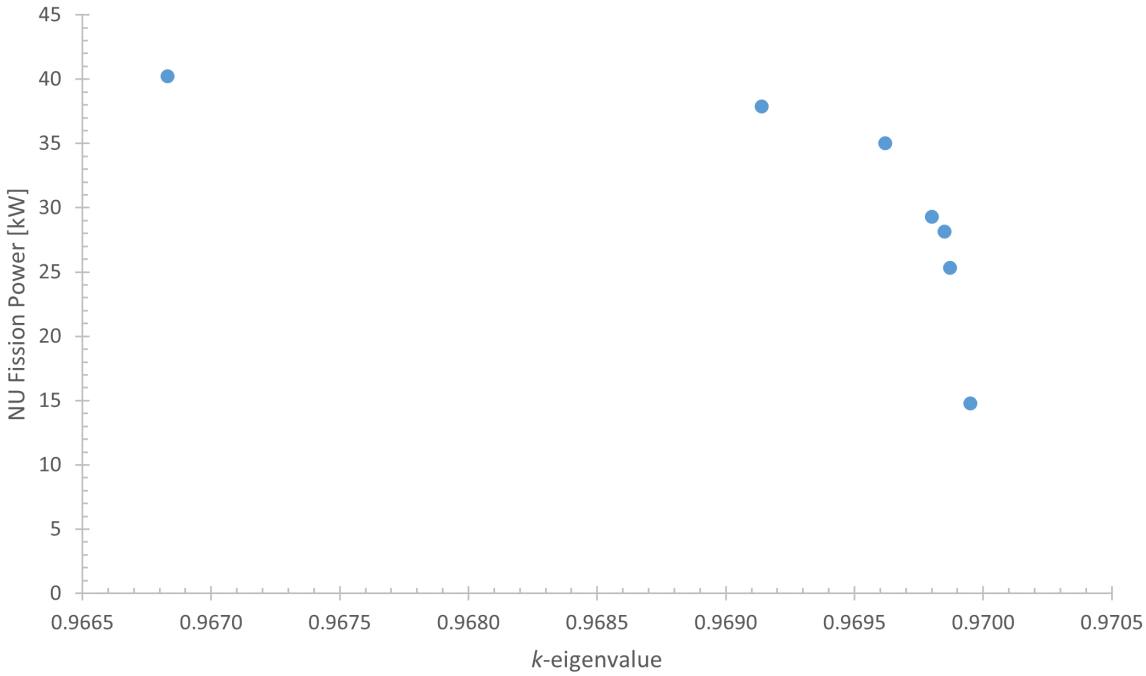

Fig. 17. Pareto front for NU fission power versus k_{eff}.

Unfortunately, Dakota:MOGA produced a much smaller Pareto front than usual, perhaps in part because only 1,000 MCNP evaluations were sampled as a result of time and computational resource constraints, or perhaps because Dakota:MOGA merely became caught in a local minimum. Additionally, the Pareto front stretched over a very limited range of k_{eff}: 0.9665 to 0.97. The points between $0.9695 \leq k_{\text{eff}} \leq 0.9700$ were within one standard deviation of the constraint. For points with a mean k_{eff} of 0.9693, 16% of the data inadvertently crossed the constraint and was eliminated.

This percentage only increased as the mean increased closer to the constraint, so the points closest to $k_{\text{eff}} = 0.97$ may not have represented the true shape of the Pareto front. Determining trends from a larger sample of cases would be better—as shown in Fig. 18, for example.

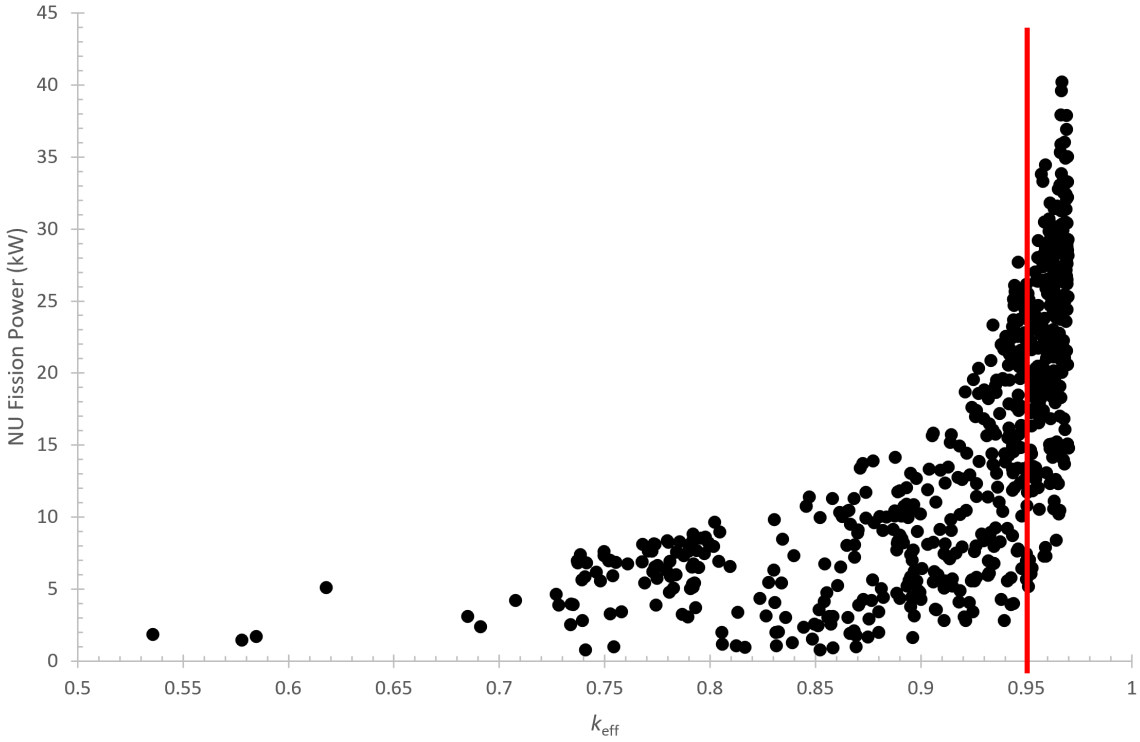

Fig. 18. Calculated NU fission power versus k_{eff} for 1,000 configurations of the UTA-3 Stage 4 design model sampled by Dakota:MOGA.

Dakota:MOGA's broader NU fission power versus k_{eff} trends align better with expectations, once sufficiently far from the $k_{\text{eff}} \leq 0.97$ constraint. Unlike the previous per-source-fission neutron rates, maximum NU fission power was positively correlated with k_{eff}. Higher k_{eff} reinforced a higher potential max NU fission power, which follows from higher neutron multiplication within the LEU fuel. The best NU fission power obtained was approximately 30 kW for a UTA-3 configuration near $k_{\text{eff}} = 0.95$.

Although these results are favorable, the trade-offs involved in using Dakota:MOGA to optimize the UTA-3 assembly have become increasingly clear. Once the model became sufficiently complex, either more MCNP function evaluations (samples) were required to determine a good set of optima and a well-defined Pareto front or more human assistance was required. In this work,

the latter was selected because of time and funding constraints. However, another option was discovered: the development of a reduced-order model through Gaussian process regression [24]. This method showed more promise in terms of reduced human intervention and computational resource requirements.

IV. CONCLUSIONS

In this work, the performances of Dakota's automated single-objective genetic algorithm (SOGA) and MOGA were compared with that of a traditional MCNP parameter search (optimized by a human). The objective of the study was to optimize an ADSA design for the production of ^{99}Mo with LEU in accordance with national and international goals. The optimization of ^{99}Mo production was represented functionally by maximizing the assembly's k_{eff} in early stages of the design and maximizing the k_{eff} and the fission rate in the NU fuel rods in later stages. MCNP models of the UTA-3 assembly were created for each parameter configuration and used to calculate the objective quantities for both optimization methods.

The overall optimal parameter trends for the final design evaluated (Stage 4) were as follows. The k_{eff} trends were easier to interpret than the sparser NU fission power trends. Sparse sampling is an unfortunate side effect of using genetic algorithms as stochastic optimizers. The optimal pitch-to-diameter ratio was found to be between 1.5 and 2. The k_{eff} was fairly insensitive to rod height, but an optimal range was clear between 30 and 75 cm. As expected, a near-maximal mass of LEU was required in a central cylindrical pattern to produce configurations with high k_{eff} and NU fission power within the UTA-3 assembly. The k_{eff} was less sensitive to LBE target placement. However, placing one or both targets near the boundary between the last ring of LEU and the first ring of NU rods was more optimal.

Dakota:MOGA performed well in the early phases of design and found optimal parameter sets that were slightly better than human-informed optimization when the number of parameters was limited to 20 or fewer. However, when hundreds to thousands of parameters needed to be optimized (such as in the case of optimal fuel rod pattern development), Dakota:SOGA did not perform as well for small sample sizes (1,000–10,000 function evaluations). Dakota:SOGA found only random rod patterns that were suboptimal, even when a quarter-core symmetry assumption was implemented. Therefore, a proportional increase in objective function evaluations (hundreds of

thousands to millions) would be required to obtain a good result, but this number of evaluations was beyond the scope of this study.

As the UTA-3 assembly model became more complex in the later stages of design, the MCNP calculations became significantly more computationally expensive . This factor, combined with a genetic algorithm's need for a larger number of functional evaluations to fully converge in general, meant that more human intervention was required to interpret trends and identify optimal regions beyond a few local minima identified in the Pareto fronts. Even so, Dakota:MOGA discovered more optimal parameter sets than MCNP alone in Stage 3 and explored the optimization space reasonably well in Stage 4. Future work will incorporate a reduced-order model to make automated optimization of a complex nuclear system more feasible with moderate computational power.

ACKNOWLEDGMENTS

The authors acknowledge the staff at Niowave Inc. for their support and collaboration throughout the design process. Financial support for this research was provided by the US Department of Energy's (DOE's) National Nuclear Security Administration Office of Material Management and Minimization's Molybdenum-99 Program. The report was authored by UT-Battelle LLC under contract no. DE-AC05-00OR22725 with the DOE. The publisher acknowledges the US government license to provide public access under the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).

REFERENCES

[1] NNSA, "NNSA's Molybdenum-99 Program: Establishing a Reliable Supply of Mo-99 Produced without Highly Enriched Uranium," (2019)URL https://www.energy.gov/nnsa/nnsa-s-molybdenum-99-program-establishing-reliable-supply-mo-99-produced-without-highly.

[2] NNSA, "NNSA Helps Global Health Industry Achieve Major Nuclear Nonproliferation Milestone," (2023)URL https://www.energy.gov/nnsa/articles/nnsa-helps-global-health-industry-achieve-major-nuclear-nonproliferation-milestone.

[3] M. Bencomo, "Parametric Study of Mo-99 Production Using a Subcritical Low Enriched Uranium Assembly Design Proposed by Niowave Inc." Master's Thesis, Texas A&M University (2016).

[4] C. D. Bowman, "Once-Through Thermal-Spectrum Accelerator- Driven Light Water Reactor Waste Destruction Without Reprocessing," *Nuclear Technology*, **132**, *1*, 66 (2000); 10.13182/NT00-1.

[5] H. Shahbunder, C. H. Pyeon, T. Misawa, and S. Shiroya, "Experimental Analysis for Neutron Multiplication by Using Reaction Rate Distribution in Accelerator-Driven System," *Annals of Nuclear Energy*, **37**, 592 (2010); 10.1016/j.anucene.2009.12.022.

[6] J. Engelen, H. A. Abderrahim, P. Baeten, D. De Bruyn, and P. Leysen, "MYRRHA: Preliminary Front-End Engineering Design," *International Journal of Hydrogen Energy*, **40**, 15137 (2015); 10.1016/j.ijhydene.2015.03.096.

[7] E. Van Abel, G. Piefer, and T. Radel, "Design of a Commercial Scale Accelerator Driven Subcritical Aqueous Assembly," *Transs. Am. Nucl. Soc.*, **114**, 798 (2016).

[8] A. K. Grimm, C. H. Boulware, T. L. Grimm, W. A. Peters, and M. A. Zamiara, "Niowave's Domestic Radioisotope Production from Uranium and Radium," *Transs. Am. Nucl. Soc.*, **120**, 229 (2019).

[9] N. Nelson, M. B. R. Smith, Z. Karriem, J. Navarro, C. P. Denbrock, R. N. Wahlen, and T. L. Grimm, "Radiation Shielding Analysis of Niowave's Uranium Target Assembly 2

(UTA-2) Facility for Molybdenum-99 Production," ORNL/TM-2021/2269, Oak Ridge National Laboratory (2022).

[10] R. B. Wilkerson, C. Denbrock, T. Grimm, J. Nash, J. Navarro, and R. Wahlen, "Radiation Dose Modeling for Niowave's Accelerator-Driven Uranium Target Assembly 3," ORNL/TM-2024/3646, Oak Ridge National Laboratory (2024).

[11] S. B. Sadineni, "Benchmarking Photoneutron Production of MCNPX Simulations with Experimental Results," Master's Thesis, University of Nevada (2002); 10.25669/mcev-7nhw.

[12] N. Nelson, A. Conant, J. Navarro, C. Denbrock, R. Wahlen, and T. Grimm, "A Comparison of the Performance of Three Nongradient-Based Optimization Algorithms during the Design of an Accelerator-Driven Subcritical Assembly for Mo-99 Production," *Proc.Int. Conf. on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2023)*, Niagara Falls, Canada, Aug. 13–17, American Nuclear Society (2023).

[13] A. Conant, N. Nelson, J. Navarro, T. Grimm, and R. Wahlen, "MCNP Neutronic Design Optimization of an Accelerator-DrivenSubcritical Assembly for Mo-99 Production," *Proc.Int. Conf. on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2023)*, Niagara Falls, Ontario, Canada, Aug. 13–17, American Nuclear Society (2023).

[14] B. Adams, W. Bohnhoff, K. Dalbey, M. Ebeida, J. Eddy, M. Eldred, and R. Hooper, "Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.15 User's Manual," SAND2021-14253, Sandia National Laboratories (2021-11); 10.2172/1829573.

[15] C. J. Werner, "MCNP Users Manual - Code Version 6.2," LA-UR-17-29981, Los Alamos National Laboratory (2017).

[16] J. E. Morel, L. J. Lorence Jr., R. P. Kensek, J. A. Halbleib, and D. P. Sloan, "A Hybrid Multigroup/Continuous-Energy Monte Carlo Method for Solving the Boltzmann-Fokker-Planck Equation," *Nucl. Sci. Eng.*, **124**, 369 (1996); 10.13182/NSE124-369.

[17] A. Wang, A. Maslowski, T. Wareing, J. Star-Lack, and T. Schmidt, "A Fast, Linear Boltzmann Transport Equation Solver for Computed Tomography Dose Calculation (Acuros CTD)," *Medical Physics*, **46**, *2*, 925 (2018); https://doi.org/10.1002/mp.13305.

[18] E. E. Lewis and W. F. Miller Jr., *Computational Methods of Neutron Transport*, American Nuclear Society, La Grange Park, IL (1993).

[19] X-5 Monte Carlo Team, "MCNP — A General Monte Carlo N-Particle Transport Code, Version 5," LA-UR-03-1987, Los Alamos National Laboratory (2008).

[20] T. Back, *Evolutionary Algorithms in Theory and Practice : Evolution Strategies, Evolutionary Programming, Genetic Algorithms*, Oxford University Press, Inc., New York (1996).

[21] A. Konak, D. W. Coit, and A. E. Smith, "Multi-Objective Optimization Using Genetic Algorithms: A Tutorial," *Reliability Engineering and System Safety*, **91**, 992 (2006); doi:10.1016/j.ress.2005.11.018.

[22] A. et al, "Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.15 Reference Manual," , Sandia National Laboratories (2021).

[23] R. S. Bivand and D. W. S. Wong, "Comparing Implementations of Global and Local Indicators of Spatial Association," *TEST*, **27**, 716 (2018); https://doi.org/10.1007/s11749-018-0599-x.

[24] C. Ball, J. Navarro, A. Conant, C. Denbrock, R. Wahlen, and T. Grimm, "UTA–3 Subcritical Linear Accelerator–Driven System Data Analysis and Optimization," ORNL/TM-2024/3456, Oak Ridge National Laboratory (2024).