Artificial Intelligence Design of Nuclear Systems

Vladimir Sobes
Briana Hiscox
Emilian Popov
Marco Delchini
Richard Archibald
Cory Hauck
Paul Laiu
John Pevey
Wes Hines
Ben Betzler
Kurt Terrani

August 30, 2019

Transformational Challenge Reactor Program

ARTIFICIAL INTELLIGENCE DESIGN OF NUCLEAR SYSTEMS

Vladimir Sobes
Briana Hiscox
Emilian Popov
Marco Delchini
Richard Archibald
Cory Hauck
Paul Laiu
John Pevey
Wes Hines
Ben Betzler
Kurt Terrani

Date Published: August 30, 2019

CONTENTS

FIGURES

TABLES

ABSTRACT

Advanced manufacturing applications in the nuclear industry have the potential to revolutionize nuclear design. While current engineering design objectives may be consistent with objectives from the past, such as those intended to minimize the fuel cost or to limit the maximum fuel temperature for a given power level, the possible geometries which can be considered as candidate designs have changed. Optimizations applied to current designs in nuclear engineering are limited by traditional geometric shapes that are easy to manufacture with traditional processes: slabs as fuel plates, cylinders as fuel rods, spheres as fuel pebbles, axis-aligned coolant channels, and more. Advanced manufacturing allows for an approach to the nuclear design problem with arbitrary geometries that can enable previously unrealized optimizations in performance and cost. This work explores the application of arbitrary geometry to the nuclear engineering design process.

Without the usual constraints of conventionally manufacturable geometries, the design optimization process becomes less restrictive but more difficult. There are three major challenges to this change in the nuclear design paradigm.
1. The first challenge is based on the compact mathematical representation of arbitrary geometry. How do we represent the shape of the design while maintaining the freedom to make any changes to the shape and keeping the number of control parameters low?
2. The arbitrary geometry requires a high number of design parameters, making design optimization the second challenge.
3. The third challenge is to develop the technical capability to perform multiphysics simulation of nuclear designs with arbitrary geometry. This challenge is more practical in nature.

This work combines two software packages to perform multiphysics simulations of nuclear designs with arbitrary geometries. Two different approaches are used to compactly represent arbitrary geometry that can be manipulated and simulated. Finally, this work explores artificial intelligence (AI) and advanced methodologies to perform the optimization. The results are demonstrated on a series of five designs of increasing complexity.

If advanced manufacturing techniques become widely used in the nuclear industry, then AI design of nuclear systems will significantly impact the field. While AI will not replace the human designer, the use of AI algorithms will shift the designer's focus away from conjecturing good geometries to defining the design criteria and find optimal designs. In the five demonstration designs presented herein, the AI algorithm produced two results that initially seem counterintuitive to a human designer, but on further inspection, they have been deemed optimal. The exploration of these designs demonstrated how AI uses the physics of the problem to find optimal designs.

The implementation of AI as discussed herein is beginning to demonstrate applicability to the difficult nuclear design problem. Large investments in computational modeling and simulation research do not solve this engineering design problem: computational modeling and simulation are applied to determine engineering performance characteristics of a particular design, but additional optimization approaches are required to make a better design. The AI algorithms implemented and demonstrated here can be applied to design optimization based on a given set of criteria, assisting designers in the determination of which qualities improve a design to best meet the desired objectives within given constraints. Additional work is required to fully explore and demonstrate the application of AI and these tools to arbitrary reactor design, including to a commercial design concept, which would include real-world design constraints and objectives.

1. INTRODUCTION

1.1   OBJECTIVE

This project used artificial intelligence (AI) algorithms to explore how the freedom of geometry made possible with advanced manufacturing techniques changes the optimization of nuclear designs. In this initial effort, the goal was not to produce one optimal design for a given application, but rather to research and establish a framework for nuclear system design optimization that includes the types of arbitrary geometry enabled by advanced manufacturing.

This report presents the framework found to be the best for nuclear optimization with arbitrary geometry. Two approaches were selected to represent the geometry, both of which are very flexible and conservative in the number of characterizing parameters. Multiple optimization approaches and their associated trade-offs are discussed. Lastly, while it was not the explicit goal of this project to find optimal designs, the proposed framework was demonstrated on five design optimization scenarios of increasing complexity, as described herein.

The primary objective was to propose a path forward for nuclear systems design in an arbitrary geometry framework, and the secondary objective was to understand the possibilities and limitations of the proposed framework and how it will impact nuclear design.

1.2   LONG-TERM VISION AND IMPACT

The long-term vision for this project is to develop the initial framework into a very general nuclear engineering design tool. One of the powerful aspects of the framework is the modular nature of the relationship between the multiphysics calculations, the geometry representation, and the optimization software. Therefore, it will be relatively straightforward to incorporate more physics modules into the performance evaluation of one candidate design. While the current framework only considers the neutronics, heat transfer, and fluid-dynamics associated with the candidate designs, structural mechanics, thermophysical fuel performance, economics, fuel cycle capabilities, etc. can easily be added without requiring many changes in either the geometry representation module or the optimization module.

All of the modeling and simulation results evaluated relate specifically to the optimization module of the framework as either objectives or constraints. For example, the maximum fuel temperature may be a constraint on the design, while the mass of the fuel may be the minimization objective. Furthermore, any physical constraints from specific advanced manufacturing techniques are considered beneficial rather than detrimental, as they constrain the search space of possible solutions for the optimization module.

In the long-term, the AI-based nuclear design capability will create a different role for the human designer: to clearly and quantitatively determine the objectives and constraints for the design, as these impact the resulting optimized design. While the constraints are usually more straightforward, defining a single quantitative objective is difficult. Possible design objectives for the longer term are listed in Table 1, many of which are quantifiable with modern modeling and simulation. However, the designer's role will be to understand the balance between these objectives when defining a single quantitative objective function.

Table 1. Possible long-term optimization objectives

Economics	Fuel cycle	Reactor physics
Construction costs	Proliferation resistance	Enrichment
Operating costs	Activation products	Local peaking factors
Front end fuel costs	Waste	Burnable poison
Resource utilization	Cycle length	Thermal efficiency

These goals could effect a change in computational reactor physics. So far, computational reactor physics has been focused on evaluating the performance of a single design. This project's long-term objective is to develop solution methodology for finding optimal nuclear engineering designs given a set of performance characteristics.

1.3   DOCUMENT LAYOUT

This document is structured around the project's three critical components. Section 2 identifies the constraints on *Multiphysics Simulation*s chosen for initial demonstration, followed by implementation of the computational solvers for this problem. The near-term improvements to the multiphysics modeling are discussed, as well prospects for generalization beyond the initial proof of concept. Section 3 discusses the *Geometry Representation*. The important features of the geometry representation are presented, and the choices are considered, along with the advantages and disadvantages of the two selected representations. Section 4 presents the different *Optimization* approaches explored in this project, as well as their strengths and weaknesses. The roadmap for the follow-on optimization work is established. Figure 1 illustrates the connection between the three components and will be referenced further in the document.
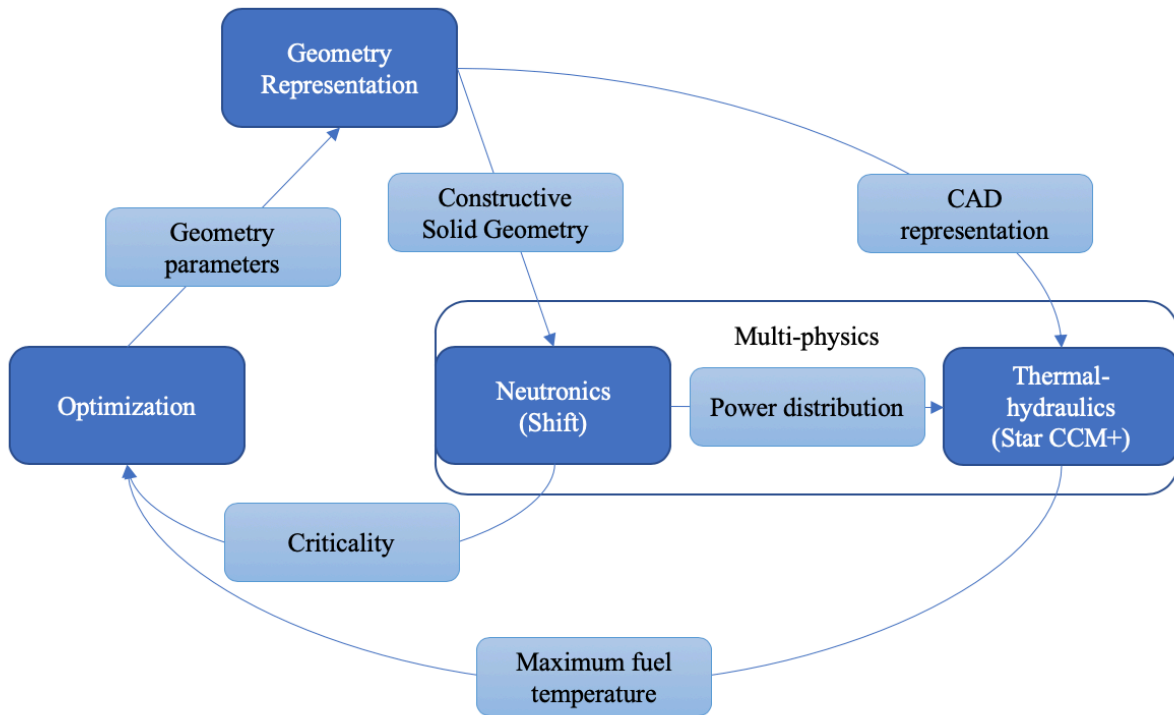
Figure 1. Framework flow connecting the multiphysics modeling capability (Section 2) to the geometry (Section 3) and optimization (Section 4) capabilities.

In this initial report of a new beginning in AI design for nuclear systems, all three sections include subsections that discuss the future work and potential for the three components: multiphysics, geometry, and optimization.

Section 5 introduces the challenge problems of increasing difficulty that were used to study and demonstrate the proposed framework. Each challenge problem is centered around finding the optimal fuel geometry to meet a set of multiphysics constraints, with increasingly more relaxed geometrical constraints. As the geometrical constraints are relaxed, the complexity of the geometry representation increases, along with the complexity of the optimization space.

The geometry determined by AI optimization for two of the challenge problems seemed counterintuitive. Section 5 provides an deep dive into the physics of these two challenge problems to elucidate (1) how the AI algorithm was able to determine those geometries as optimal, and (2) why the counterintuitive geometries outperform the human designs.

The final section summarizes the proposed framework for AI optimization of nuclear systems divided into three timescales: near-, mid- and long-term outlooks.

2.   MULTIPHYSICS SIMULATIONS

2.1   PHYSICS PROBLEM SET UP

Early in the project, the multiphysics definition of the optimization problem was established to mimic the requirements for the Transformational Challenge Reactor (TCR) core design at the time (the design has evolved since then), with a few modifications for simplification. While the multiphysics definition of the problem determines the final *theoretical optimal* design, the exact definition is not viewed as essential to establishing the AI optimization framework. The term *theoretical optimal* refers to the global optimal point in the design space from the mathematics perspective. In practice, the global optimal design is hardly ever found in anything more complicated than the most trivial optimization problems. Rather, a stable, deep, local minimum is sought.

The exact multiphysics definition of the problem is not viewed as essential to the framework because the goal of this project was to develop a general framework which would define a method for performing AI optimization of nuclear systems that would apply to a wide range of objectives and constraints. Therefore, a set of multi-physics definitions was required for the study. Despite this, the prevailing goal was to develop a framework which would work just as well with a different set of multiphysics definitions.

Using this approach, the team worked to simplify the multiphysics definitions as much as possible to reduce the complexity and run-time of the multiphysics calculations while preserving some physical intuition. In the future, all the simplifying assumptions made here can easily be removed. However, for now, they would only cloud the construction of the framework and would take time away from the geometry and optimization work.

For the initial exploration of the framework and the challenge problems, the decision was made to optimize the geometry of a single piece of nuclear material which was modeled as suspended in the middle of a cylindrical shroud with forced coolant flowing from the bottom to the top (bottom-up). A schematic is provided in Figure 2. An important distinction is that the study was begun with the geometry optimization of the fuel piece only. This is computationally equivalent to determining the boundary between two materials: the fuel and the coolant. However, this differs from a three-material problem, such as one containing, fuel, moderator, and coolant. Future work to optimize the geometry of multiple materials is discussed in Section 3.

Figure 2. Schematic representation of the problem. Optimization is performed over the geometry of a single continuous piece of fuel suspended inside a shroud with bottom-up forced coolant flow.

The material definitions for the fuel and the coolant were fixed to the values in Table 2. All the fuel properties were taken to be temperature-independent for simplicity.

Table 2. Material properties for the fuel and the coolant used in the challenge problems.

Fuel: UO₂		Coolant: He	
Enrichment	19.75%	Inlet pressure	6 MPa
Density	10.8 g/cc	Inlet flow velocity	10 m/s
Thermal conductivity	4 W/mK	Inlet temperature	425 °C

The multiphysics constraints on the design were as follows:
1. neutron multiplication of 1.020 ± 500 pcm,
2. maximum fuel temperature of 600 °C $\pm$ 3%, and
3. design power of 10 kW $\pm$ 3%.

In this work, only the steady-state problem is solved. While 2000 pcm above critical is a lot, it should be viewed more as an arbitrary design goal. In this case we justify it by deferring reactivity control to later engineering solutions or optimization using more detailed multiphysics modeling. This is also a good example of a TCR-specific design objective that is arbitrary for this demonstration. For a real, commercial design in which the reactor would be generating power over an extended period of time, excess reactivity and its control would be a constraint and potential design objective e.g., longer cycles.

For this project, the design power was significantly reduced from the TCR design because the study was limited to a single piece of fuel instead of a fuel array.

The design objective was to minimize the fuel volume (i.e., mass, as the density was fixed). Nuclear engineering intuition dictates that the goal of minimizing the fuel volume with a constant component power will push any near-optimal design to the minimum allowed neutron multiplication and to the maximum fuel temperature. Therefore, constraints 1 and 2 were reformulated as a multiplication factor above 1.015 and a maximum fuel temperature below 618 °C, taking into account the allowable tolerances.

It was verified that in the multiphysics regime used for optimization in this effort, the fuel density changes are negligible with respect to both reactivity and fuel temperature. Furthermore, the temperature feedback coefficient for criticality is within the uncertainty of $\pm$ 500 pcm for temperatures of up to 925 °C. While we do not neglect to Doppler broaden the cross sections to the average operating temperatures, we do neglect the effect of Doppler reactivity feedback. That is, only solve the neutron transport problem once at mean temperatures. Similarly, all changes in thermophysical properties (e.g., thermal conductivity) due to irradiation are ignored due to the very short irradiation time assumed for TCR.

Photon transport was noted for heating calculations as well as neutron and photon heating in the coolant, but it was neglected as appropriate for the flow characteristics.

With approximations taken into account, only a pseudo-coupled neutronic/thermal-hydraulic problem was solved. The component power level to was fixed to 10 kW, regardless of the criticality or fuel geometry, the relative power distribution was passed to the thermal-hydraulics solver. This is considered pseudo-coupling because it is only one-way. For this set of challenge problems, only the impact of the power distribution on the temperature distribution is considered, with no feedback to the neutronics solve from the thermal-hydraulics. This is justified by the calculations of negligible temperature feedback in the operating regime.

2.2 CURRENT APPROACH FOR MULTIPHYSICS SIMULATIONS

State-of-the-art modeling and simulation codes used in the nuclear research community were used to model the physics described in the previous section. A coupling methodology was developed between the neutron transport code Shift (Pandya, et al., 2016) and Star CCM+ (Osher & Fedkiw, 2002; Sethian, 1999). However, to enable rapid development and testing of the framework, a low-fidelity surrogate model was developed which was found to be accurate to within several percent of the high-fidelity model, which is not sufficient for true design analysis but is sufficient to demonstrate these approaches. The surrogate model is below, followed by the full physics model.

2.2.1 Low Fidelity (Surrogate) Physics Model

The coupling of the neutron transport and the thermal-hydraulics in the surrogate model was accomplished the same way as in the high-fidelity model: the relative power density was passed from the neutron transport solve to the thermal-hydraulics solve. The solution for both the neutron transport and the thermal hydraulics was calculated only on the fuel for the surrogate model.

The neutron transport surrogate model is a one-group neutron diffusion model with zero flux boundary conditions. The one group constants were determined to match the infinite homogeneous medium multiplication factor of 1.51723 and a critical sphere radius of 35.78 cm. Zero boundary conditions for non–re-entrant boundaries were taken for simplicity, and the coolant in between the re-entrant surfaces was modeled with zero absorption and a large diffusion coefficient.

The thermal hydraulics surrogate model was a heat-diffusion model with the fuel thermal conductivity from Table 2. The total component power was normalized to 10 kW, and fixed temperature boundary conditions of 425 °C were assumed.

As crude as the physics approximations were for the surrogate model, the use of the surrogate model proved to be extremely valuable for rapid prototyping of the AI algorithms. For the multiphysics regime, which is defined in Section 2.1, the surrogate model was found to be easy to use and accurate to within a couple of percent compared to the full-fidelity physics solve described in the following subsection. This physics-informed surrogate model is expected to perform better than most surrogate models. This could be tested using machine-learning techniques. The surrogate model can be solved efficiently using standard solvers in commercial software such as MATLAB.

2.2.2 High Fidelity (Full) Physics Model

Python and java scripts have been written to automate the process shown in Figure 1. The optimization module will produce a set of parameters characterizing the geometry, and the geometry module will produce input files for the two physics simulators. On the order of d^2 number of candidate designs are evaluated simultaneously, where d is the number of parameters defining the geometry.

The ORNL-developed Shift Monte Carlo radiation transport code (Pandya, et al., 2016) is used for the neutronics solve to leverage its ideal scaling characteristics. The Shift code then produces the multiplication factor and a core fission tally distribution. These tallies are converted into power densities and normalized so that the total core power is 10 kW. The power density distribution is then read into a Java code which builds and meshes the geometry and runs the thermal hydraulics code Star CCM+ on each core. Star CCM+ returns the maximum temperature of each candidate design. Then the maximum temperature and multiplication factor (k_{eff}) is fed back into the optimization module. The optimization module proposes a new set of candidate designs. This process is repeated until a converged design is achieved.

In view of the diverse geometry configurations of the optimized designs, an accurate prediction of temperature distribution in the core component is necessary. It is important that the thermal-hydraulic model is capable of computing a conjugate (solid-fluid) heat transfer in arbitrary geometry shapes. To achieve this, a computational fluid dynamics (CFD) approach is taken. This method allows complex surfaces to be discretized with finite volume techniques, and it also allows the interface between the solid structure and the coolant (gas in this case) to be properly defined. On the solid side, a thermal diffusion of heat with a volumetric heat source is computed to determine the temperature distribution. The heat generated by nuclear fission is deposited in the core element according to the volumetric power distribution supplied by the reactor physics calculation.

On the fluid side, a Reynolds averaging of the velocity vector field is employed within the finite volume formulation. Since the assumed flow is highly turbulent, a two-equation model of turbulence is used (renormalized k-epsilon). The near-wall velocity field is resolved with the two-layer all wye (Y) plus method. All these models provide an accurate prediction of the wall heat transfer, which is critical for the proper resolution of core component temperature field.

To model the geometry variation necessary for running the suite of optimization codes, two approaches are taken. The first approach relies on using a geometry parametrization method. Within the CFD computation, the geometry is regenerated automatically every time a new combination of parameters is tried. The software that was used allows for modification of computational domain—both geometry and discretization—without user intervention. This reduces the runtime per single sample down to several (8– 10) minutes in parallel on 4–8 compute cores.

The second approach is based on importing a CAD file containing the geometry of the core element. In this case, the parametrization is not necessary, and the thermal computation relies on a predefined geometry configuration. This saves time for defining and setting up parameters, but it requires the candidate geometry to be composed of a single continuous surface. Both approaches are fully integrable in an autonomous workflow within the entire optimization suite.

2.3 FUTURE DEVELOPMENT FOR MULTIPHYSICS SIMULATIONS

The multiphysics modeling component of this project only defines the landscape in which the optimization is performed. Therefore, the general direction for future development of this project's multiphysics modeling component is to explore and understand how different physics assumptions, approximations, constraints, and goals change and challenge the design process. For clarity, a brief discussion is provided below describing how four different categories of changes to the multiphysics modeling could impact the optimization of design.

2.3.1 Problem Definition

It is worthwhile to explore and test the framework for AI optimization of nuclear design established in this document on a vastly different problem definition in nuclear design. This will evaluate the robustness of the approach. To distinguish from the following subsections, the problem definitions for all five of the challenge problems are summarized in Figure 2. This project optimizes one continuous piece of fuel in bottom-up forced coolant flow. A change in the problem definition could be to optimize a finite array of identical fuel pieces which interact in their neutronics and thermal hydraulics. Another alternative might be to optimize one component of a reactor design with respect to a fixed design for the rest of the reactor. In the latter problem definition, the remainder of the reactor design would provide the boundary conditions for the optimization.

One near-term change in the problem definition which will be explored is the optimization of nuclear design with multiple materials. This study only defines the boundary of the fuel, which is alternatively viewed as the boundary between the fuel and the coolant. Optimizing a critical nuclear design operating with a thermal neutron spectrum that depends on a fuel material, a moderator material, and a distinct coolant material (e.g., a gas-cooled, graphite-moderated reactor) will present new challenges to the way the geometry is represented. Another interesting example is to present the AI optimization with a choice of several coolant or moderating materials and to allow the algorithm to determine the combination of materials that optimizes performance.

2.3.2   Assumptions, Approximations and Physics Fidelity

Higher fidelity physics will be required to produce designs for which the simulated performance matches actual performance. The major stress on the proposed framework from an improvement in the fidelity of the multiphysics simulation will likely be the need for significantly longer run-time to evaluate candidate designs. This will require even better AI optimization algorithms to reduce the number of candidate designs that must be evaluated.

Upon removing some of the assumptions and approximations made to simplify the physics, the increasing run-time can be combated through *multiresolution physics modeling*. The idea behind *multiresolution physics modeling* is that full physics models are not initially required to get an idea of design performance. For the initial scoping study, the surrogate model (or some more advanced form of it) can be used. Furthermore, as the AI algorithm advances the design toward more optimal performance, low-fidelity versions of the full-physics codes can be used. Some examples are Monte Carlo neutron transport calculations with less particles and less-converged thermal-hydraulics solves. The trade-off between the speed and accuracy of the multiphysics calculations will be explored to determine the best approach.

Some of the geometry configurations explored, such as the truncated cone shown in Figure 3, produce complex velocity fields in the assumed cooling conditions. In Figure 3, this blunted cone (truncated cone with large frontal area) creates large flow detachment at the leading edge, resulting in an eddy that reflects from the shroud walls and affects the heat transfer on the side surfaces. For such rotating flows, the applied turbulence model may not be fully adequate, and the computed heat transfer coefficients may require reevaluation using higher order methods and/or model validation against experimental data. Increasing the boundary layer resolution and employing large eddy turbulence simulations are some of the conceivable methods. Case studies will be conducted to confirm and tailor the turbulence model coefficients. These methods will no likely be practical for production calculations due to their large consumption of computational resources, and their application will be limited to confirmatory analyses.
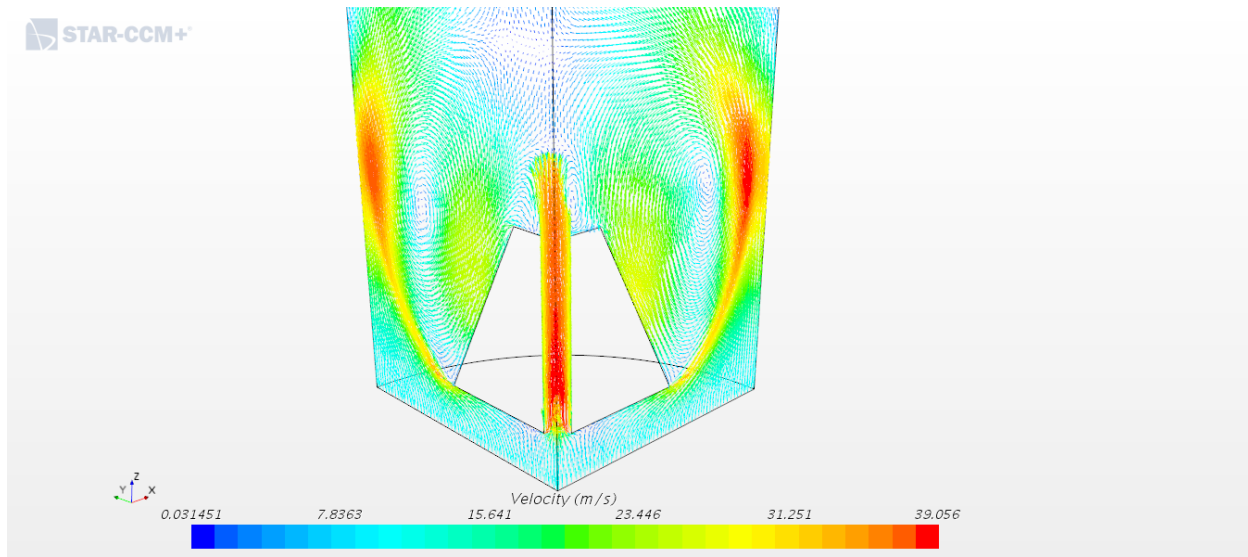
Figure 3. Velocity contours around a blunted truncated cone fuel element.

The CFD model developed for this project relies on numerical methods (temporal and spatial discretization) and physical models (turbulent flow, conjugate heat transfer, etc.) to predict the temperature and flow distribution in the geometry of interest. The inherent assumptions of these methods and models must be quantified to ensure correctness and accuracy of the CFD results. A common example in the CFD community is a grid refinement study which consists of performing the same CFD analysis on successive meshes of increasing density to investigate the induced changes in output parameters of interest (Eca & Hoekstra, 2014). To relate to the current study, a mesh sensitivity analysis could be performed on the maximum temperature in the fuel to show that the numerical uncertainty is within acceptable ranges. Further sensitivity analyses can also be performed on the turbulent model and the material properties, among others. Such sensitivity analyses will help quantify the thermal-hydraulics uncertainty of the CFD model and will provide reliable CFD results for optimization problems.

2.3.3 Constraints

In one respect, additional constraints on the design make the optimization problem easier to solve, as they can reduce the size of the search space. On the other hand, the challenge is in determining which regions of the search space are cut off by the additional constraints. Two high-priority constraints to add are limitations coming from particular advanced manufacturing techniques and limitations from structural mechanics. Limitations from both of these sources will cut off regions of the search space where candidate designs propose pieces that are too thin/narrow. While the need for lower maximum fuel temperature drives the designs to look like traditional heat transfer "fins," limitations from manufacturing and load bearing will push back on exceedingly unmanufacturable geometries.

2.3.4 Objectives

The objective established for the challenge problems is to minimize the fuel volume for the given problem definition and multiphysics constraints. An alternative objective might be to maximize the component power for a given fuel volume.

Overall, this framework was focused on single objective optimization. The optimization framework that can respond to a general category of quantitative objective functions is considered a valuable contribution. However, definition of the quantitative objective function will then become the major task

for the human engineer. Two of the challenge problems have shown how the AI algorithm finds the optimal solution exactly for the objective function defined, raising the question of whether the solution is for the problem that the team wanted to solve.
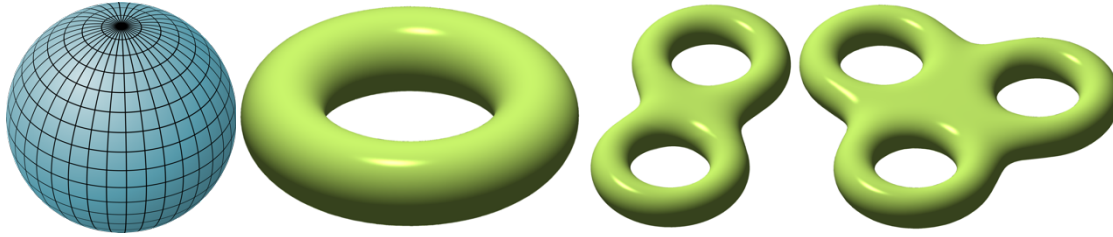
3. GEOMETRY REPRESENTATION

3.1 GEOMETRY REPRESENTATION CONSIDERATIONS

There are two major challenges with geometry representation for nuclear systems design. The first is preserving the flexibility to represent a wide variety of geometrical configurations, referred to here as *arbitrary geometry*. The second major challenge (and the opposing force) is to minimize the number of parameters which control the geometry to reduce the search space for the optimization algorithm. Furthermore, the changes in the geometry should be smooth with respect to changes in the controlling parameters in order to simplify the optimization.

One challenge which appears with geometry optimization for the multiphysics problem defined here is the formation of holes. From the mathematical subject of topology, there is a strict distinction between objects with different numbers of holes (technically referred to as *orientable surfaces with genus n*, where *n* represents the number of holes). Figure 4 provides several examples. The key distinction between objects of different genera is that they cannot be transformed into one another through a continuous deformation of the surface. For example, the surface of a sphere can be continuously deformed into a cube, but it cannot be deformed into a torus without introducing a hole. In the context of some more familiar nuclear designs, the nuclear fuel pellet is a continuous deformation of the sphere, while the cladding without end caps is a continuous deformation of the torus with one hole. A water moderator around a single fuel rod can also be viewed as a continuous deformation of the torus with one hole. A fuel spacer grid in a pressurized water reactor can be viewed as continuous deformation of a torus with $17 \times 17 = 289$ holes. In a simplified manner, all real components in nuclear engineering map directly to a unique object in the series started as shown in Figure 4 through a continuous deformation of the surface.

In the subject application, which has bottom-up coolant flow, introducing holes into the geometry drastically changes the cooling characteristics and the maximum fuel temperature of the design. The introduction of holes will likely produce discontinuity in the multiphysics solutions for a wide range of nuclear applications. Therefore, this section addresses continuous surface deformation and the introduction of holes into the prospective design.

Figure 4. Topological objects of different genera. https://en.wikipedia.org/wiki/Genus_(mathematics)

The 2013 article by Sigmund and Maute (Sigmund & Maute, 2013) outlines the most promising approaches to geometry representation in the field of topology optimization in applied mathematics. After careful consideration of each approach, a surface-control-point approach for continuous surface deformation and was selected, along with level-sets for opening and closing of holes. These two approaches and their application to nuclear design will be discussed in the next subsection.

The density approach was discarded, as it basically pixelates the geometry of the design and searches for the pixels that should contain fuel material and those that would contain the coolant material. The main reason for this decision was that the team could not find a way to restrict the number of fuel pieces in this representation. For example, a pixel of fuel could be generated away from the main geometry in the middle of the coolant field. Furthermore, in this representation, the team could not determine how to ensure that hole formation would happen continuously throughout the geometry and open up a new coolant channel as opposed to forming enclosed voids in the middle of the fuel without access for the coolant flow.

Furthermore, it is important to note that, unlike applications to topology optimization in other engineering disciplines, there will be no access to surface derivatives. For example, there is no way to directly (other than by finite-difference) calculate the change in the maximum fuel temperature as an infinitesimal perturbation on the geometry of the fuel is made. Looking forward to creating a framework for a wide variety of nuclear engineering applications, the team designed this algorithm based on the assumption that such derivatives will not generally be available.

3.2   GEOMETRY DEFINITION

One standard in design optimization is to define curves by interpolation through a set number of control points (Gero, 1985). By manipulating the number and position of the control points, many possible surfaces can be described by a relatively small number of degrees of freedom. In this study, both linear and spline interpolation were used to describe a 2D surface that could be extended to 3D by assuming either z-axis symmetry to produce a rod type geometry or rotational symmetry to produce a torus type geometry. Control point geometry represents the next advancement step, as compared to geometric parameter optimization, where a set of parameters control the surface shape.

Level sets methods are a class of tools for locating and tracking surfaces. In the context of the current project, these surfaces form the interface between different materials, including voids. The key component of the method is the level set function f. Interfaces are the surfaces on which f is a fixed constant. These interfaces are often chosen to be zero by convention. For example, in two dimensions, interfaces may be defined as the set of all points (x,y) such that $f(x,y) = 0$. In the context of shape optimization, a level set function is sought that optimizes a certain objective.

To approximate level sets numerically, one must introduce a parameterization of the function f via the introduction of a computational grid or by an approximation in terms of global basis functions such as

polynomials, Gaussians, or trigonometric functions. In either case, the parameters of the representation form the space over which an optimization algorithm must be implemented.

3.3 GEOMETRY REPRESENTATION ADVANTAGES/DISADVANTAGES

The main advantage of control point optimization is that the range of possible shapes that can be described by control points is large compared to the number of control points used. Therefore, it is possible to effectively search over a large range of surface shapes by sampling over a few degrees of freedom. However, with the current implementation in 2D, only extensions along some axis of symmetry can produce 3D surfaces. In this implementation, shapes are either rod- or torus-like in one axis of symmetry. More complicated shapes cannot be represented with the structural ease possible with using level-sets to define surface meshes.

Level sets are relatively easy to implement, can be generalized to arbitrary dimensions, and naturally incorporate the separation and merging of interfaces. However, the equation that defines the evolution of the level-set function f depends on the specification of the interface velocity. In the context of shape optimization, this velocity is tied to the gradient of the objective in the shape optimization problem. For the current project, estimating this gradient requires many samples of an expensive, high-fidelity code. An alternative approach currently being pursued is to approximate (see Subsection 3.2 above) the level set function with a parameterized expansion of global basis functions and optimize over the parameters. Such an approach requires fewer high-fidelity samples and can be used in conjunction with less expensive surrogate models. While this approach can represent very complex interfaces, it cannot represent arbitrary fine scale structures like a grid without a large number of parameters. In addition, some logic must be introduced in the optimization procedure that forbids parameterization that leads to non-physical or non-practical structures, such as freely suspended materials or material that is locally too thin to be structurally stable.

3.4 FUTURE DEVELOPMENT FOR GEOMETRY REPRESENTATION

In general, control points should be used as an intermediary method between geometric shape optimization and level-set optimization. The degrees of freedom, and hence the total cost of optimization, are an order of magnitude smaller than level-sets. Future development of 3D control points can shorten the bridge between level-set and control point optimization. The first step would be to parameterize the symmetry axis for the 2D control point method and to allow some feature (such as parameter size or rotation angle) to vary throughout the original symmetry axis. With a few additions to the size of the degrees of freedom occurring from the parameterization of the symmetry axis, completely different surfaces that could optimize fluid flow can be tested. Similarly, a set number of channels can be introduced to linearly increase the degrees of freedom while exploring porous candidate designs. Modest adjustments to the current control point method can be extended to bridge a significant region of design space that exist between geometric shape optimization and level-set optimization. (It is possible to generate a quick picture of a 3D control point assembly with a few internal cooling channels that vary as a function of the coolant-flow-axis).

In the level set context, multiple materials can be handled using multiple level set functions. Specifically, m level set functions can account for as many as 2^m different materials. However, for structural reasons such as cladding around a fuel, the relative geometries of different materials may need to be coupled together in a more direct fashion. The major challenge of level sets in three dimensions is the increased computational complexity relative to two-dimensional problems. However, structural restrictions may also be more complicated to express and enforce in three dimensions. In addition, restrictions on the geometry may be required that are not present in two dimensions.

4. OPTIMIZATION

4.1 OPTIMIZATION APPROACHES

The statistics community has expended considerable effort developing emulators that can approximate functional forms given uncertain data (Higdon, Kennedy, Cavendish, Cafeo, & Ryne, 2004) (Kenned & O'Hagan, 2001). Gaussian processes (GPs) are the preferred methods for approximating both the mean field and covariance processes from arbitrary pointwise values under the assumption that the values are independently identically distributed $N(0,\sigma^2)$ random variables (Rasmussen & Williams, 2005). These GP approximation methods can be used with the optimization data to approximate the global optimization surface given a limited sampling of the full optimization surface. The team has developed a method that statistically determines the next set of samples to be simulated. This is likely to minimize a given loss function and should maintain any set of constraints given all the available simulations that have been run. This method grows in accuracy as more samples are simulated. For the case of optimization of the annulus challenge problem (Subsection 5.3), GP optimization was able to determine a near optimal solution in a couple of hundred samples in total.

The pattern search (also known as *direct search*) method (Kolda, Lewis, & Torczon, 2008) (Audet & Dennis, 2002) was applied for solving this topological optimization problem. Pattern search is a class of derivative-free optimization methods that determines the search direction by probing objective function values at sample points in the optimization domain. To incorporate the multiphysics constraints, they are built into the objective function as penalty terms. In the optimization procedure, the GP approximation is used to estimate the criticality and maximum temperature at a given point in the optimization domain; this avoids the prohibitive computational cost of extensive multiphysics simulations.

The pattern search method used here consists of two stages: a global random search and a local refined search. The global search stage randomly samples search points that are further away from the current location to explore the optimization domain, while the local search procedure probes adjacent points to guarantee convergence to local minimums. Full multiphysics simulations are then run on the convergent points, and the GP approximation is updated with the simulation results. Then a new iteration of pattern search is started.

4.1.1 Genetic Algorithms

Genetic algorithms are optimization techniques that borrow ideas from the natural world, including mutation, breeding (crossover), and survival of the fittest. Like in the natural world, in a genetic algorithm, unique individuals are compared according to a set of criteria and are combined to create new, unique individuals. These are then added to the population, and the best are again selected and combined. While the general outline of the algorithm is simple, the execution can be complex. This is due to having to explore potentially complex ideas such as how individuals should be judged, how individuals should be created and mutated, and how heuristics can be imposed to reduce unnecessary computation. This has been explored by starting with simple geometries of solid spheres and cylinders and solved for neutron multiplication using the relatively quick but least accurate 1-dimensional diffusion equations. By using simple, quick methods to evaluate the fitness of individuals, other important parameters related to the task were explored, such as mutation rates and varying crossover methods. These studies were built upon further by switching to a more accurate discrete ordinates method for solving for criticality and temperature and adding the annulus geometry type.

4.2 FUTURE DEVELOPMENT FOR OPTIMIZATION

Future developments will focus on three topics: improved optimization algorithms (Subsection 4.2.1), faster evaluation of the function to be optimized (Subsection 4.2.2), and genetic algorithms (Subsection 4.2.3).

4.2.1 Optimization Algorithms

Currently, the optimization is based on a pattern search algorithm (Torczon, 1997). The reason for this choice is that most topology optimization problems are based on objectives that require expensive simulations, making gradient approximations difficult. In the current problem, the objective is the volume of the fuel, which is easily computable; the expensive simulation comes from computation of the constraints based on temperature and criticality. These constraints can be approximated using surrogate models based on a relatively small number of samples of a simulation code, allowing for a local gradient calculation.

4.2.2 Function Evaluation

Optimization in high dimensions invariably requires a large number of samples. Thus, in order to reduce the total sampling cost, cheaper surrogate models for function approximation will be developed; this will be done specifically for the constraints which are determined by neutronic and thermal hydraulic simulations. Currently, two surrogates are used: the Gaussian process model is a data-based model, and the reduced-order simulation is a physics-based surrogate. Other data models (such as neural networks) will be explored, as well as other physics-based surrogates and numerical surrogates which reduce cost using coarse grid resolution. In addition, these approaches will be combined into multi-level (Giles, 2008) and multi-fidelity approximations (Peherstorfer, Willcox, & Gunzberger, 2018) to maximize overall model fidelity for a given computational budget.

4.2.3 Genetic Algorithms

Work is ongoing to increase the robustness of the genetic algorithm code so that more complex geometries can be studied. A drawback to genetic algorithms is that as the individuals being optimized become more complex as the number of individuals evaluated increases. This increase in complexity can be countered with more effective mutation, crossover, and fitness operators that are more specifically tailored to the problem, including the use of more advanced machine learning techniques such as reinforcement and learning.

5. OPTIMAL GEOMETRY RESULTS

5.1 INTRODUCTION TO THE CHALLENGE PROBLEMS

The challenge problems described in the following subsections were designed to be progressively more difficult to gradually challenge the team's ability to:
1. represent increasingly more arbitrary geometry,
2. perform reliable multiphysics calculations, and
3. perform optimization over an increasing number of free parameters.

The other considerations in the design of the challenge problems were reasonable expectations of the existence of a global optimum (or a deep and stable local optimum) and the possibility of cross-validation between the challenge problems.

In establishing an AI framework for nuclear systems design, for the human engineer, the burden will shift towards defining the objective for the optimization well. It was important to use engineering judgement to create the challenge problems to eliminate the possibility of trivial solutions being optimal. For example, the case of simply asking to minimize the mass of a critical cylinder of fuel has a well-known solution from neutron diffusion theory. However, the AI solution to a slight modification of this simple problem was unexpected, as described in the following subsection. Another so-called "trivial" solution which had to be avoided was the optimal solution becoming fractal-like in nature to optimize cooling. A heat-transfer, fin-like solution becomes optimal for minimizing peak fuel temperatures. For many applications, this is not practical because of structural or manufacturing constraints.

In the following subsections, the critical sphere radius for the material definitions is 0.358 m, which corresponds to a volume of 0.192 m^3 and a surface area of 1.6 m^2. The infinite homogeneous multiplication factor is 1.51, and the maximum allowable fuel temperature was 618 °C.
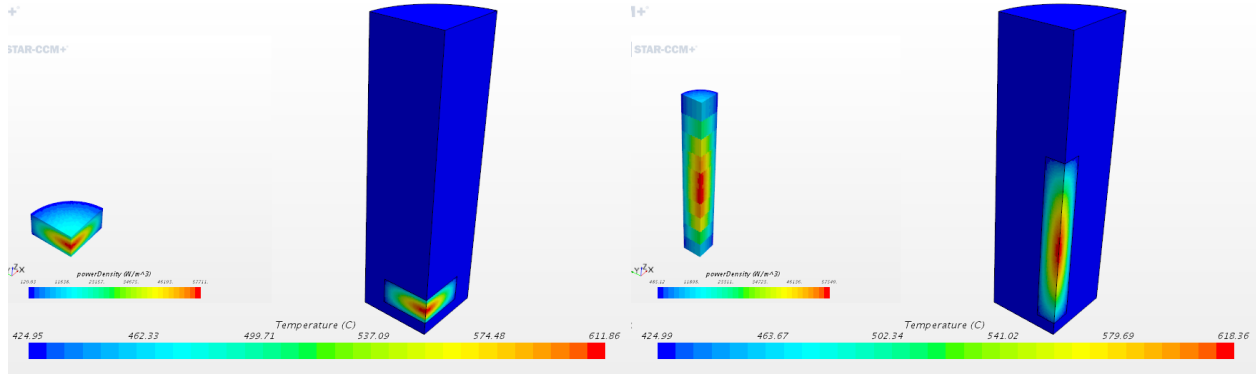
5.2 CYLINDER CORE

The first step was to restrict the geometry to a right-cylinder. This geometry is far from arbitrary and is described by two free parameters; radius (R) and height (H). As for all the challenge problems, the goal was to minimize the fuel volume while maintaining a neutron multiplication factor above 1.015 and a maximum fuel temperature below 618 °C. The cylinder is positioned in a shroud with bottom-up coolant flow, as described above.

The goal of this challenge problem was to start with a simple problem, which allowed the team to:
- Exercise each of the multiphysics modeling modules,
- Establish the neutronics/thermal-hydraulics coupling,
- Compare the surrogate physics model to the full-physics model, and
- Get started on the optimization module with only two free parameters.

Furthermore, the cylinder core set up a reference for cross-validation of the multiphysics modeling of the Star Core in Subsection 5.5 and the Level-Sets Core in Subsection 5.55.6.

The AI solution to the cylinder core also provided the first non-intuitive solution. By restricting the geometry to a right-cylinder, one of three possible solutions was expected to be optimal: (1) a flat disk like a coin, (2) a compact cylinder with an R/H ratio near 0.5, like an LWR fuel pellet, or (3) a long, thin cylinder like a fuel rod. The AI optimization found two near-equivalent solutions to the posed problem that were disjoint. The two solutions (within two significant figures) were combinations of R = 0.67 m, H = 0.37 m and R = 0.27 m, H = 2.3 m. These are the only two solutions, and no solution exists in between. Figure 5 gives a visualization of the two solutions.

Figure 5. Visualization of two equivalent solutions for the cylinder core. Full geometry (quarter symmetry) is shown on left of each respective solution, and the CFD solution is on the right. The power density distribution is inserted in the CFD plot on the left. Left: R = 0.67 m, H = 0.37 m. Right: R = 0.27 m, H = 2.3 m

Although it was not intuitively anticipated at first, the reasoning behind the two AI solutions was soon apparent. In this case, a multiphysics problem involving criticality and heat transfer is being solved. Criticality can be viewed, in a simplified manner, as the balance between fuel volume, which increases criticality, and neutron leakage (non-reentrant surface area) which decreases criticality. For heat transfer, the same two meta-parameters play a role, but in the opposing direction (as this study is trying to minimize the maximum fuel temperature). Fuel volume increases the maximum fuel temperature, and surface area decreases the maximum fuel temperature. Both of the solutions found by the AI had nearly the same volume of 0.52 m^3 and the same surface area of 4.4 m^2. On the other hand, in the subsequent challenge problems, no degenerate optimal solutions were found, even for geometries with equivalent surface-to-volume ratios as more complex neutronics, thermo-dynamic, and fluid flow considerations come into play.
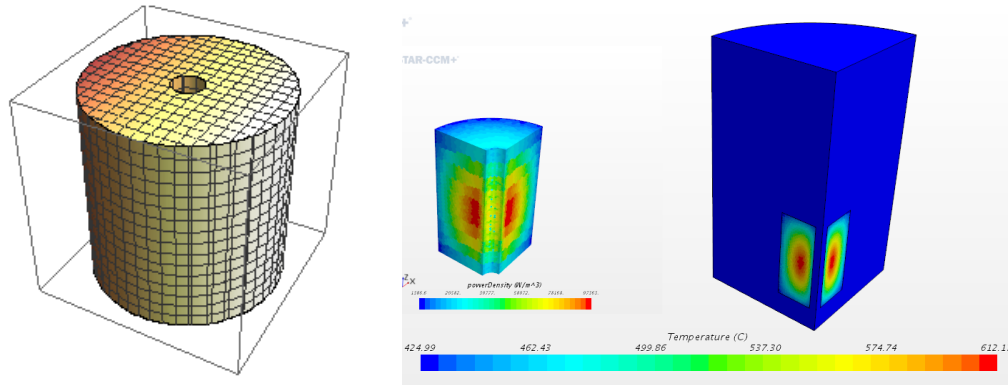
5.3   ANNULUS CORE

The annulus core addressed the same problem formulation as before, but it allowed for the geometry to be parameterized by three free parameters: radius of the inner cylindrical flow channel, the outer radius of the fuel annulus, and the height of the fuel. The annulus core presented several new challenges:
- In both the surrogate model and the full multiphysics simulations, the reentrant geometry of the inner, cylindrical cooling channel had to be addressed.
- The optimization space increased to three degrees of freedom.
- The two radii parameters had to be restricted so that the inner radius was strictly smaller than the outer radius.
- This allowed for experimenting with the dramatic change in performance that results from introducing a hole in the design

With respect to the discussion in Subsection 3.1, the cylinder core has no holes and can be continuously deformed into a sphere, while an annulus core has one hole and can be continuously deformed into a torus.

The AI optimization found the optimal design to be an annulus with an inner cooling channel radius of 0.057 m, an outer fuel radius of 0.34, m and a fuel height of 0.70 m. In this case, the volume was 0.25 m^3, a reduction of over 50% from the cylinder core (due to reduced neutron leakage) and only 30% above the minimum volume to maintain criticality at any maximum fuel temperature. The surface area for this design was 2.5 m^2. Figure 6 shows the relative proportions of the optimal annulus core.

Figure 6. Relative proportions of the optimal annulus core: geometry (left), and CFD results (right).

5.4 CONE CORE

The cone core is the first design in the series of challenge problems to diverge from traditionally manufacturable geometries and to seek out improved performance by considering simple shapes which could easily be made with advanced manufacturing but are difficult to manufacture using traditional methods. The cone core was a variant of the annulus core, and the inlet and outlet radii for both the inner cooling channel and the outer fuel radius could vary independently. The cone core presented a new set of challenges:

- The optimization space has increased to 5 free parameters: inlet inner radius, inlet outer radius, outlet inner radius, outlet outer radius, and height. The respective radii also had to be restricted.
- This was the first design to really challenge the fluid dynamics aspect of the calculations. As the cone core either became tapered into the flow or blunted, the fluid dynamics changed.
- As the surrogate model did not include a fluid dynamics approximation, it could not be used for the scoping studies.

Some possible candidate geometries are presented in Figure 7. Figure 7 visually presents only the effect on the outside of the geometry. The cone core is a variant of the annulus core where the outside of the fuel had the freedom to be tapered into the bottom-up coolant flow or blunted. The internal flow channel also had the freedom to either constrict the cooling channel or expand from the bottom to the top.

18
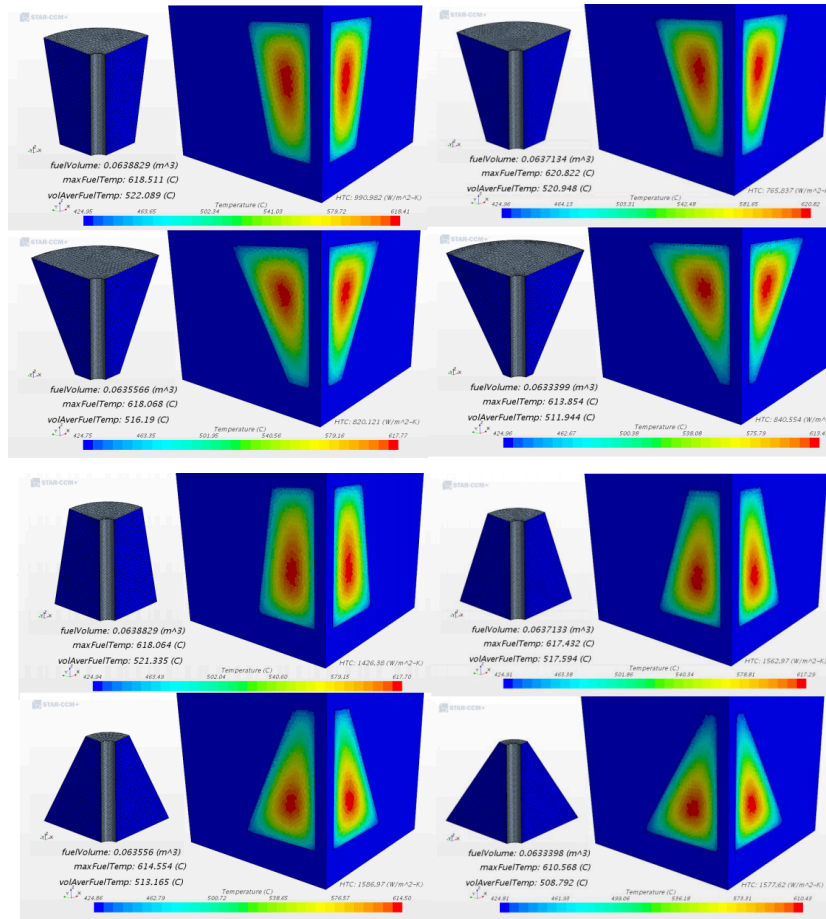
Figure 7: Candidate designs for the cone core. Showing the tapering of the outside of the geometry (top 4 figures) and the blunting of the outside of the geometry (bottom 4 figures).

The AI optimization found that the optimal design had an inlet inner radius of 0.07 m, an inlet outer radius of 0.37 m, an outlet inner radius of 0.05 m, an outlet outer radius of 0.35 m, and a height of 0.64 m.

The cone core was the second unintuitive result that the AI algorithm provided. As this was the first experience trying to optimize the flow around the fuel piece, an answer was expected which would taper the nose (inlet) of the cone core to smooth out the flow. This expectation was the opposite of what the AI found to be optimal.

The optimal cone-like core configuration is in the geometry shape of a truncated annular cone, with the inlet surface being larger than the outlet (blunt body). A more detailed analysis was performed to confirm the correctness of the optimization algorithm by analyzing the cooling efficiency of the optimal geometry. Two options are considered when the element geometry varies in a way that leads to tapering (option one) or blunting (option two) the element.
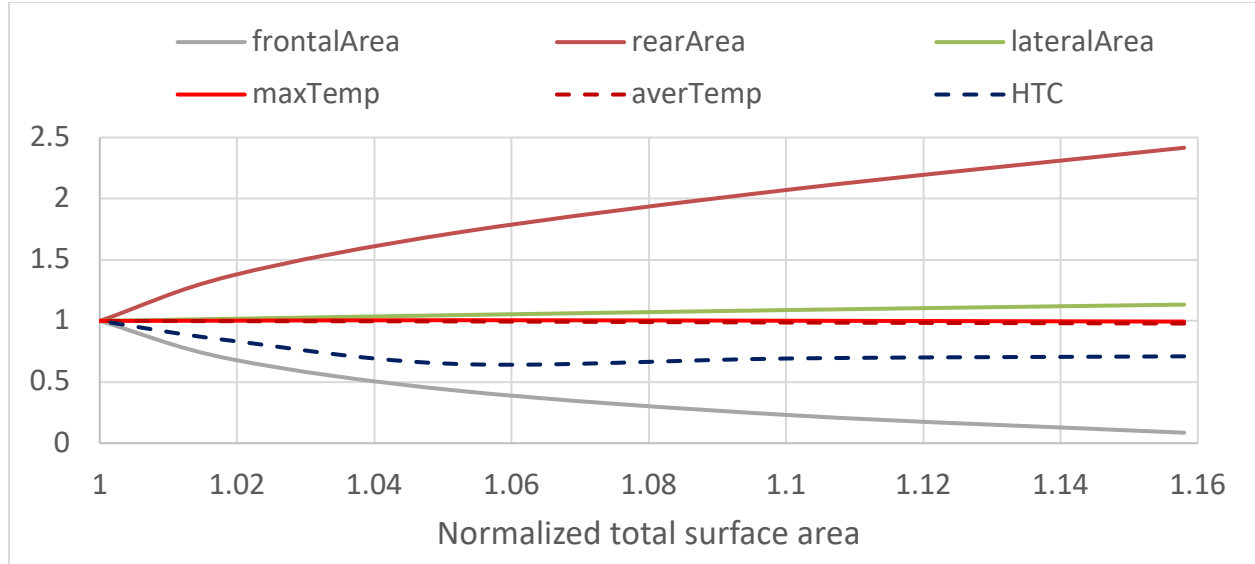
Figure 8 Variation of cone geometry side surfaces and overall heat transfer coefficient (HTC) in function of total surface area for the case of tapering the core element.
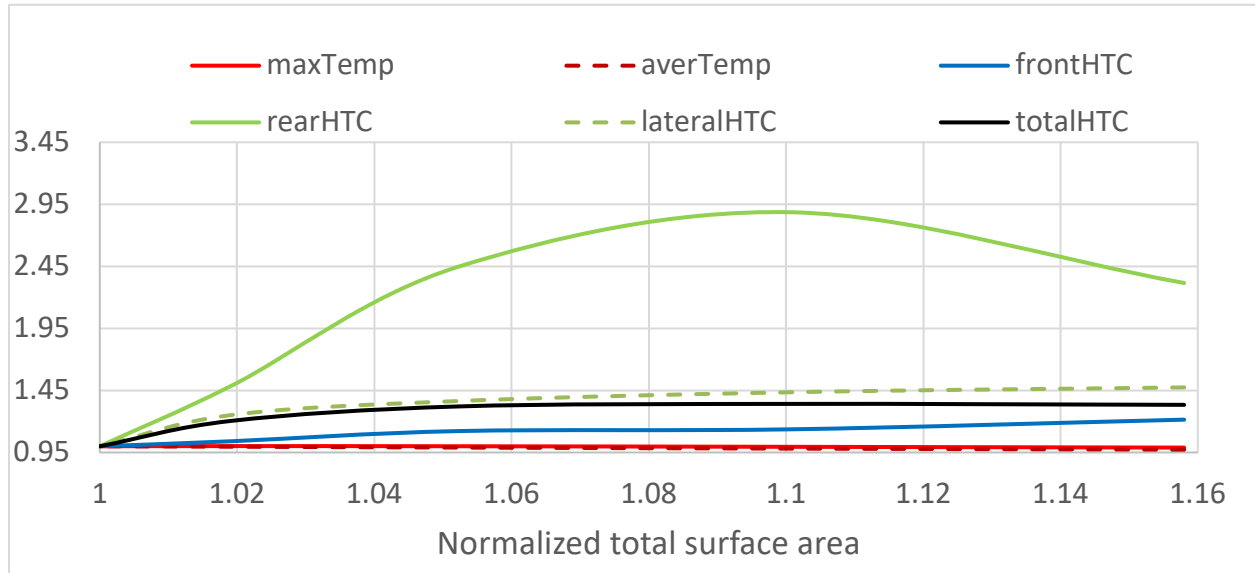
Figure 9 By surface heat transfer coefficients (HTC) for the blunt core configuration showing the improvement of cooling efficiency.

Results are summarized in Figure 8 and Figure 9. Figure 8 shows the variation of the surface areas and the overall increase of the total area, when, during the optimization search, the element shape changes from a regular annulus to a truncated cone. Figure 9 demonstrates the outperformance of the blunted configuration, thus confirming the correctness of the optimization algorithm. The total heat transfer coefficient (HTC) increases by about 30%; this is mainly caused by the enhancement of the frontal and side surface cooling efficiency. Future analyses will continue to confirm and improve the model accuracy for proper heat transfer prediction.

5.5 STAR CORE

The star core is the first arbitrary geometry core that is no longer confined to circles perpendicular to the direction of flow. In general, only two restrictions have been placed on the geometry of the star core. The first is that the star core is uniform in cross section in the axial (flow) direction. The second is that the star core is not allowed to have any holes or internal coolant channels. These two restrictions make it necessary to optimize the perimeter of the cross section of the fuel piece perpendicular to the direction of the coolant flow. The height of the fuel can also be varied. A candidate design is presented in Figure 10.
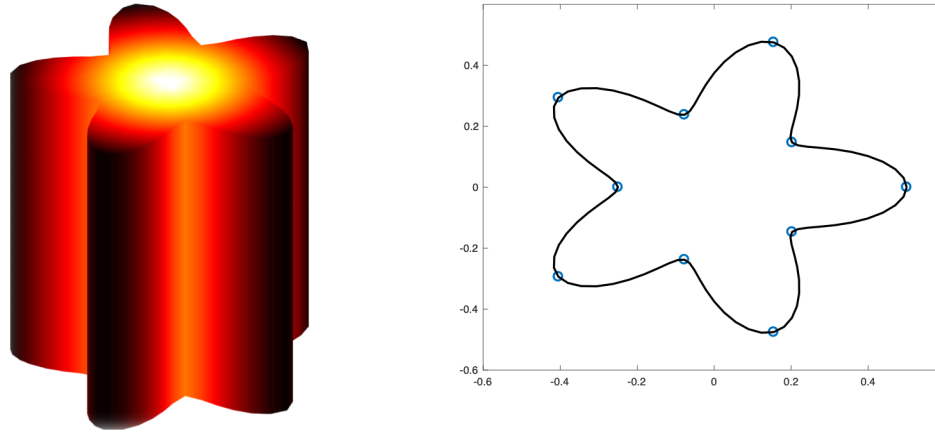
Figure 10: A candidate design for the Star Core.

The star core is a very challenging geometry because it is so free in what it can do. In fact, engineering intuition holds that the optimal geometry for the restrictions on the star core would be some fractal pattern which would present infinite cooling surface area to reduce the maximum fuel temperature but would curl the geometry so that the surface area would present the possibilities for neutrons to re-enter the fuel geometry and maintain criticality. This should be pictured as some cylinder-like shape made up of infinitely thin sheets that allow coolant flow to pass through. Such a solution is obviously not practical once manufacturing and structural considerations are added into the multiphysics constraints.

To maintain the current multiphysics modeling capabilities relevant, the fixed surface area constraint has been added. This constraint translates into a very interesting, practical question in thermal-hydraulics: given a fixed surface area, what is the optimal geometry for cooling: With unrestricted an surface area, the sphere (smallest surface-to-volume ratio) is known to be the least optimal geometry for cooling, and the fin/fractal solution is the most optimal (infinite surface to volume ratio). However, the geometries that lie in between are unknown. Therefore, the solution is sought for the star core with the minimal fuel volume for a fixed surface area of 6 m^2.

With the surface area of the fuel fixed, it was anticipated that the optimization would also be convergent with an increasing number of control points. That is, the shape of the cross section with a smaller number of control points would be a rough approximation of the shape with a larger number of control points. This convergence was expected to occur only after some (to be determined) number of control points. Increasing the number of control points beyond this amount would yield diminishing returns.

5.6   LEVEL-SETS CORE

The level sets core was the most flexible geometry solution. In this challenge problem, the only restriction imposed was that of the axially uniform cross section, along the direction of the coolant flow and a fixed surface area. The goal of the level-sets core was to find the optimal fuel cross section geometry—that is, the minimal fuel volume meeting the multiphysics constraints—with a fixed surface area. The height of the fuel piece was an additional free parameter in the optimization. The level-sets core allowed for this new geometry representation to be used to cross-validate the cylinder core, the annulus core, and the star core solutions.

The major new development of the level-sets core was the ability to transition between geometries of different genus number, or to create and destroy holes or internal cooling channels. In the level-sets core, the AI optimization had the ability to determine the appropriate number of internal cooling channels and how to best distribute the limited surface area between fin-like structures on the outer surface and internal cooling channels.

5.7   TORUS CORE

The torus core is the final challenge problem in the immediate future. The shape is some continuous deformation of the torus—a shape with one hole—but the cross section of the *tube* of the torus is not a trivial geometry. The annulus core can be viewed as one possible candidate for the torus core in which the cross section of the *tube* of the torus is a rectangle. The cone core is a torus core with a *tube* cross section of a trapezoid. The torus pictured in Figure 4 has the cross section of a circle. For the torus core, the cross section shape of the *tube* is permitted to be any 2-dimensional geometrical figure. In other words, the optimal solution to the torus core can be described as the optimal geometry with one internal cooling channel with rotational symmetry around the coolant flow direction. A representative visualization is provided in Figure 11.
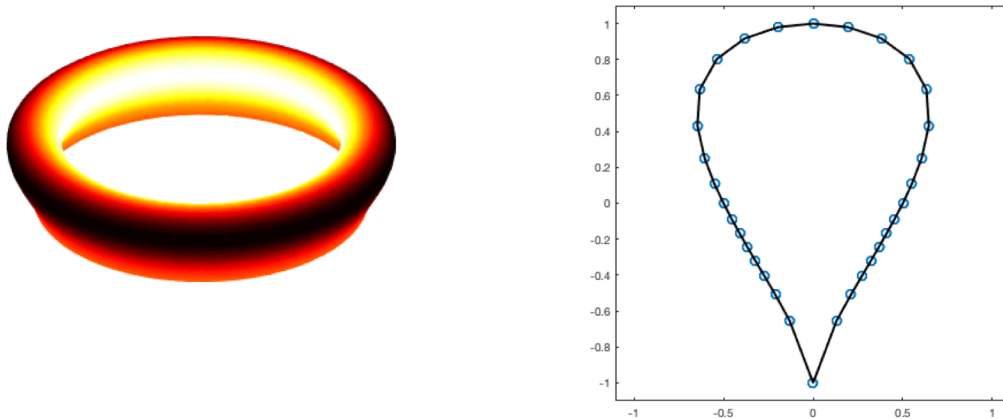
Figure 11. A representative visualization of a design candidate for the torus core challenge problem (left)), and a cross section view of the tube (right).

The torus core presents the unique challenge of solving for fuel shape, which optimizes the coolant flow dynamics. It was expected that the torus core would greatly optimize the coolant fluid flow around the fuel piece. Furthermore, the behavior of the optimization code is expected to be convergent as the number of control points was increased. A solution with 10 control points on the surface of the *tube* should look like a rough approximation of a solution with 20 control points. Increasing the number of control points beyond a certain (believed to be small) number would yield diminishing returns in performance.

5.8 NEAR-FUTURE CHALLENGE PROBLEMS

In the challenge problems above, a wide variety of constrained and semi-arbitrary geometries (2-dimensional cross sections) were explored. the near-future challenge lies in exploring fully arbitrary geometries in three dimensions. Keeping the current problem definition, objective, and multiphysics constraints, the artificial surface area constraint will likely need to be enforced to prevent the divergence of the optimal solution into infinitely thin, fractal, structures. However, determining the best, truly arbitrary geometry in three dimensions with a fixed surface area (e.g., 6 m^2) is extremely significant.

In this challenge problem, two new issues must be addressed:

1. The initialization/scoping study becomes particularly difficult as efforts become focused on a large search space of arbitrary geometries in 3 dimensions.
2. Even the local optimization becomes computationally intensive as the transition is made to hundreds or thousands of free parameters which control the geometry. This is motivating us to look at intelligent uses of supercomputers to aid in the optimization of design.

6. FUTURE WORK DIRECTION SUMMARY

6.1 NEAR-TERM DEVELOPMENT

The last significant item remaining to close out this preliminary investigation of the framework for AI optimization of nuclear design is to solve the arbitrary geometry problem fully in three dimensions with restricted surface area. For the problem definition, objective, and multiphysics constraints chosen, this is the ultimate challenge. Once it is possible to solve for the optimal 3D geometry for a given surface area, then it will be possible to solve for a slight increase and decrease in the surface area. The most educational part for nuclear engineering design will come from a continuous visualization of how the optimal geometry changes as the amount of available surface area is slowly increased. At some minimal surface area, the solution will be a single piece of fuel and will either be a sphere or some small modification of a sphere and a fairly large volume. As the available surface area increases, the volume of the optimal design will rapidly decrease, and the shape of the design will adapt to maintain criticality and improve fuel cooling. With increasing surface area, internal cooling channels will begin to form. The shape and number of those cooling channels will provide insights into nuclear systems design. Upon further increasing the available surface area, the AI algorithm will be capable of breaking apart the fuel into several pieces. At this point, the engineering trade-offs between internal cooling channels and independent fuel pieces can be studied. In the long-term development, the AI optimization framework will allow for the study of things like optimal geometries for cooling, shape and number of internal cooling channels, and the trade-offs of breaking off independent fuel pieces as the problem definition, objectives, and multiphysics constraints are modified.

6.2 MID-TERM DEVELOPMENT

The goal for mid-term development is to connect the research with advanced manufacturing techniques and to develop a design to be built and experimentally tested. This connection is essential for demonstrating the relevance of the AI optimization framework for nuclear engineering applications. Multiphysics modeling capabilities must be expanded to include thermophysical fuel performance, structural mechanics, as well as constraints from advanced manufacturing. The goal of this project will be to select a specific nuclear component and quantitatively express the multiphysics constraints on the design and the objective of the optimization. Then, the AI optimization framework will be used to develop a design. The design will be manufactured and experimentally tested against the predicted performance from the simulation software. The performance of the design will be judged in context of the

performance of the same component of different design and will be produced with traditional manufacturing techniques.

6.3 LONG-TERM DEVELOPMENT

The long-term vision for the AI optimization framework of nuclear design is two-fold. First, there is the obvious goal of creating a very general framework which to provide reliable designs for particular nuclear engineering problems. To achieve this end, the framework must be tested on a variety of different nuclear engineering problems using multiphysics simulation of many different phenomena. The second long-term goal of this project is for the computational power and AI algorithms developed here to be used to learn the meaning of optimal design in nuclear engineering and to learn how designs can be improved to be more optimal for particular objective functions with certain multiphysics constraints. Efforts in high-fidelity computational nuclear engineering have focused on determining the performance of set designs; the objective of this work is to harness the power of modern modeling and simulation software to understand how to *improve* those designs.

7. BIBLIOGRAPHY

Audet, C., & Dennis, J. E. (2002). Analysis of generalized pattern searches. *SIAM Journal on optimizaiton, 13*(3), 889-903.

Eca, L., & Hoekstra, M. (2014). A procedure for the estiamtion of the numerical uncertainty of CFD calculations based on grid refinement studies. *Journal of Computational Physics*, 104-130.

Gero, J. (1985). Design Optimization. *Academic Press*.

Giles, M. B. (2008). Multilevel Monte Carlo path simulation. *Operations Research, 56*(3), 607-617.

Higdon, D., Kennedy, M., Cavendish, J. C., Cafeo, J. A., & Ryne, R. D. (2004). Combining field data and computer simulations for calibration and prediction. *SIAM Journal of Scientific Computing, 26*, 448-446.

Kenned, M. C., & O'Hagan, A. (2001). Bayesian calibration of computer models (with Discussion). *Journal of the Royal Statistical Society - Series B, 63*, 425-464.

Kolda, T. G., Lewis, R. M., & Torczon, V. (2008). Optimization by direct search: New perspectives on some classical and modern methods. *SIAM reivew, 45*(3), 385-482.

Osher, S. J., & Fedkiw, R. P. (2002). *Level Set Methods and Dynamic Implicit Surfaces.* Springer-Verlag.

Osher, S., & Sethian, J. A. (1988). Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Compuational Physics*, 12-49.

Pandya, T. M., Johnson, S. R., Evans, T. M., Davidson, G. G., Hamilton, S. P., & Godfrey, A. T. (2016). Implementation, capabilities, and benchmarking of Shift, a massively parallel Monte Carlo radiation transport code. *Journal of Computational Physics, 308*, 239-272.

Peherstorfer, B., Willcox, K., & Gunzberger, M. (2018). Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *Siam Review, 60*(3), 550-591.

Rasmussen, C. E., & Williams, C. K. (2005). Gaussian Processes for Machine Learning. *Adaptive Computation and Machine Learning*.

Sethian, J. A. (1999). *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science.* Cambridge: University Press.

Siemens. (n.d.). *STAR-CCM+ documentaion and code description*. Retrieved from https://www.plm.automation.siemens.com/global/en/products/simcenter/STAR-CCM.html

Sigmund, O., & Maute, K. (2013). Topology optimization approaches, A compartive review. *Structural Multidisciplinary Optimization, 48*, 1031-1055.

Torczon, V. (1997). On the convergence of pattern search algorithms. *SIAM Journal on optimization, 7*(1), 1-25.