

A Mixed-Integer Linear Programming Method for Optimal Orificing in Breed-and-Burn Cores

Chris Keckler,^{a*} Alper Atamtürk,^b Massimiliano Fratoni,^a Ehud Greenspan^a

^a*Department of Nuclear Engineering, University of California Berkeley, CA, USA*

^b*Department of Industrial Engineering and Operations Research, University of California Berkeley, CA, USA*

**keckler@berkeley.edu*

INTRODUCTION

An important aspect of any core design is to ensure that all assemblies can be adequately cooled during their residence time in the core. This can be challenging in the design of breed-and-burn (B&B) cores, where assembly powers may vary by orders of magnitude over their lifetime [1]. Assuming the flowrate within assemblies remains constant due to a fixed orifice at each assembly inlet, the orifice must be determined to cool the assembly at its peak power over the entire equilibrium cycle. Alternatively, the orifice could be placed at a fixed assembly location so that the flowrate for a given assembly can change as the assembly is shuffled (as is planned by TerraPower for their prototype reactor [1]). The former option may be required for experimental-purpose reactors where operational flexibility is of high importance, whereas the latter option may be more appropriate for commercial designs with well-defined cycles. Either orificing strategy leads to large over-cooling of assemblies at points during the cycle, which can have negative impacts on plant efficiency and may increase thermal stresses on core structures. Although it is, in principle, possible to design an orifice to be variable between, or even during, burn cycles, a feasibility study of variable orificing has yet to be published.

Determining the orificing scheme for any core is a trade-off between many simultaneous constraints occurring at various points in the burn cycle, and careful consideration must be given to the flow allocation if a viable core layout is to be achieved. In fast reactor design, earlier work has relied upon algorithmic approaches. SE2-ANL is a subchannel code for liquid metal reactors used at Argonne National Laboratory that has a built-in algorithm to determine assembly flow allocations [2]. The code relies upon the user to specify which assemblies should be grouped together, and does not account for variations in assembly power levels. These limitations introduce user bias into the problem and potentially limit its application to shorter fuel cycles with smaller power swings, making it difficult to know how far from optimal the provided solution may be. An approach suggested by Heidet et al. [3] is to begin with the highest powered assembly, calculate the flowrate required to remain within peak temperature constraints, and continue adding assemblies of incrementally lower power to that group until the average outlet temperature of the group is equal to the desired average reactor coolant outlet temperature. Once the average outlet temperature of a particular group of assemblies meets this value, a new group is started and the process is repeated. This process does not introduce user bias, as the grouping specification is determined by the algorithm.

Although these types of approaches have been adequate for most core types examined so far, they may not be adequate for B&B designs due to their inflexibility to account for large

power variations over the cycle and temperature constraints between adjacent assemblies. Typical homogeneous cores feature gently varying spatial power profiles due to long neutron mean free paths and small variation in fissile content across the core. B&B cores, however, can have sharp gradients in their power profiles as a result of their shuffling schemes, and many assemblies have extremely low power production due to them being loaded at either natural- or depleted-uranium enrichments. The trend for B&B cores to be large, low-leakage designs with many assemblies (often near 500) serves to further exacerbate the situation by making it difficult to group the assemblies into a limited number of orifice groups. Furthermore, because the residence time of fuel in B&B cores is necessarily long (often upwards of 15-20 years), the assembly powers in high burnup assemblies may be subject to significant uncertainty. Finally, shuffling schemes in B&B cores are often relatively fixed. There is less flexibility for modification of shuffling patterns due to other constraints, such as material damage and power peaking limitation, severely limiting the shuffling design space. All of these challenges contribute to making it much more difficult to design the orificing scheme for a B&B core as compared to the same task for a standard breeder or burner design.

This paper presents a new method for designing viable orificing strategies which is robust enough to handle the difficult case of B&B cores over their entire cycle length. This methodology is based upon mixed-integer linear programming [4] and can be solved efficiently in a computational setting. First, the relevant physics and constraints are given, followed by the model formulation and a discussion of the assumptions made. Finally, preliminary results and a comparison with a previous method are provided for an example B&B core.

DESIGN CONSTRAINTS

Adequate Cooling Over the Cycle

The primary concern in the flow allocation problem is to ensure that all assemblies are cooled adequately during their time in the core. The most basic requirement is to keep the coolant outlet some threshold away from the boiling temperature. A more specific requirement for a metallic core may be to keep the peak inner-cladding temperature below the temperatures which lead to eutectic reactions with metal fuel. Alternatively for oxide cores, the peak centerline temperatures may be the limiting factor due to high uncertainty in fuel thermal conductivity at high burnup. Whatever the threshold value is, it is necessary to keep the operating temperatures below this value over the entire residence time, meaning that it is important to account for changes in power level as material is bred and assemblies are shuffled.

For an assembly of known power with single-phase coolant, the change in temperature from inlet to outlet at steady-state can be calculated approximately using a simple energy balance as in Equation 1, where \dot{Q} is the heat generated from all neutron and photon reactions in the assembly, \dot{m} is the mass flow rate of liquid coolant, and c_p is the thermal heat capacity of the coolant at constant pressure.

$$\Delta T = \frac{\dot{Q}}{\dot{m}c_p} = \frac{\alpha}{\dot{m}} \quad (1)$$

If c_p is taken to be a constant average value over the expected range of temperatures, $\alpha = \dot{Q}/c_p$ is a known constant for every assembly and the temperature gain is only a function of the allocated flow. Because the inlet temperature is assumed to be the same across all assemblies, the outlet temperature can be easily calculated. This simple model does not account for heat transfer between assemblies, which may be significant if the powers of adjacent assemblies are very different. However, heat transfer between adjacent assemblies will serve to smooth the temperature profile of the core, making it easier to conform to constraints on outlet temperatures. Therefore, a model without assembly-to-assembly heat transfer is expected to be conservative in determining an initial orificing design.

Plant Thermodynamic Efficiency

It is vitally important to keep the mixed outlet temperature near the value required by the balance-of-plant power conversion cycle. The bulk outlet temperature at steady-state can be calculated as a simple weighted average of all assembly outlet temperatures:

$$\bar{T}_{out} = \frac{\sum_i \dot{m}_i T_{out,i}}{\sum_i \dot{m}_i} \quad (2)$$

Maximum Coolant Flowrate

To meet the cooling demands of a particular assembly, it is necessary to match the coolant flowrate with the assembly power. However this cannot be done to arbitrarily high power levels due to constraints on the coolant velocity and pressure drop. To increase the coolant mass flowrate at a constant coolant density and flow area, the velocity must be increased. Excessively high flow velocity can lead to flow-induced vibrations of core components and accelerated corrosion [5]. The average flow velocity in an assembly can be related to the mass flowrate as shown in Equation 3, where ρ is the coolant density (assumed as an average constant value) and A_f is the coolant flow area—for typical B&B designs with wire-wrapped fuel pins, the flow area within a given assembly remains constant over the active fuel region.

$$v = \frac{\dot{m}}{\rho A_f} \quad (3)$$

An additional constraint on the coolant flowrate is induced by the pressure drop over the core, which is proportional to the square of the mass flowrate. The pressure drop is limited by the size of the primary coolant pumps. Although more or larger pumps can be added, typically an economic tradeoff

determines the optimal hydraulic power. The pressure drop over the core length can be calculated approximately as:

$$\Delta p = K \frac{\dot{m}^2}{2\rho A_f^2} + f \frac{L}{D} \frac{\dot{m}^2}{2\rho A_f^2} + \rho gh \quad (4)$$

The bundle friction loss can be calculated using a friction factor f as determined by the Novendstern model [6] for the average Reynolds number of the flow, and the form loss factor K is the same for each assembly.

Adjacent Assembly Outlet Temperatures

Due to the high thermal conductivity of liquid sodium, thermal striping of the above-core structures is a concern when the outlet temperatures between adjacent assemblies are significantly different. As coolant of different temperatures mixes in the outlet plenum, vortices at the interface of the two flows form. As these vortices impact the upper internals, the high thermal conductivity of sodium allows for high frequency temperature oscillations to be imparted to the structure, which can lead to rapid thermal expansion and contraction and eventual failure of the upper internals. This phenomenon can be minimized through a proper orificing strategy which limits the difference in outlet temperatures between adjacent assemblies, and is a key challenge for orificing in B&B cores. TerraPower aims to keep the maximum difference in outlet temperatures between adjacent assemblies below 50°C in the power production zone of their B&B reactor [1], although other designers may aim to keep this difference limited to below 30°F [7].

Number of Orifice Groups

From the standpoint of operational and manufacturing simplicity, it is desirable to limit the number of unique orifices. Although hard limits are not placed on the number of unique orifices by physics constraints, standard fast reactors typically employ as few as practical, usually around 5-15. TerraPower plans to have 20 orifice groups in their prototype B&B, although this prototype will employ enriched uranium assemblies and a limited number of natural uranium assemblies, making the overall core power profile flatter than a breed-and-burn system fed with natural- or depleted-uranium, and thus making the orificing potentially easier. Therefore, it is perhaps reasonable that a B&B core will have somewhat more than 20 orifice groups.

MODEL FORMULATION

This study proposes to use mixed-integer linear programming to determine an orificing strategy while accounting for the constraints outlined above. This allows for the problem to be stated in such a way that algorithmic biases are avoided and the extensive theory behind mixed-integer linear programming can be applied, achieving a provably optimal solution, within the assumptions of the model. For this formulation, it is assumed that the radial power profile is known on an assembly-level from neutronics analysis. The orificing problem can be stated as below, where the objective is to minimize the number of orifice groups.

min p :

$$\begin{aligned}
 \textcircled{1} \quad & \Delta T_{max} \geq \frac{\alpha_{i,k}}{\dot{m}_i} && \forall k, \forall i \\
 \textcircled{2} \quad & \bar{T}_{out} = \frac{\sum_i \dot{m}_i T_{out,i,k}}{\sum_i \dot{m}_i} && \forall k \\
 \textcircled{3} \quad & v_{max} \geq \frac{\dot{m}_i}{\rho A_f} && \forall i \\
 \textcircled{4} \quad & |T_{out,i,k} - T_{out,i',k}| \leq \xi && \forall i, \forall i' : i \perp i' \\
 \textcircled{5} \quad & \dot{m}_i = \sum_{j=1}^p \delta_{i,j} x_j && \forall i \\
 \textcircled{6} \quad & \sum_{j=1}^p \delta_{i,j} = 1 && \forall i \\
 \textcircled{7} \quad & \delta_{i,j} \in \{0, 1\} && \forall i, \forall j \\
 \textcircled{8} \quad & \dot{m}_i \geq 0 && \forall i
 \end{aligned}$$

In this formulation, i spans the number of assemblies, j spans the number of orifice groups, and k spans the number of burn steps of interest (i.e. to account for power variations). \bar{T}_{out} is the desired bulk outlet temperature, ΔT_{max} is the maximum allowable temperature gain, v_{max} is the maximum allowable flow velocity, and ξ is the maximum allowable difference in outlet temperature between adjacent assemblies (where \perp is meant to indicate “adjacent to”). The user may specify these parameters to adhere to the specific restrictions on maximum temperatures and flowrates that are relevant to their system. Constraints $\textcircled{5}$, $\textcircled{6}$, and $\textcircled{7}$ force the assembly flowrates to be selected out of a finite pool of continuous variables whose cardinality is to be minimized. Constraint $\textcircled{8}$ forces all determined flowrates to be nonnegative.

This formulation is not linear, as constraints $\textcircled{4}$ and $\textcircled{5}$ are both in fact non-convex. Additionally, constraints $\textcircled{5}$ and $\textcircled{6}$ contain summations from 1 to p , where p is the variable that is being minimized. In order to construct a mixed-integer linear formulation, we discretize x to be a data vector of possible flowrates, specified by the user. By providing x as data, constraint $\textcircled{5}$ becomes linear, where the summation is now taken to be over the length of x . A new constraint can be introduced to count how many unique flowrates are selected, and this variable can be put into the new objective function. Constraint $\textcircled{8}$ also becomes redundant, as long as the values specified in x are physically realistic (i.e. non-negative).

Even with this linearization, constraint $\textcircled{4}$ remains non-convex. A similar approach can be used, however, to fully linearize the formulation. Since T_{in} is the same for all assemblies, $T_{out,i} - T_{out,i'} = \Delta T_i - \Delta T_{i'}$. Furthermore, the quantity $\Delta T = \alpha/\dot{m}$ can only take on a discrete number of values, as x has been discretized. This means that ΔT can be specified as a 3-dimensional table, with values tabulated for each assembly at all possible flowrates at every burn step. The difference between outlet temperatures can thus be written as the difference between sums of data multiplied by δ_{ij} , which is linear.

In addition to these linearizations, a relaxation of $\textcircled{2}$ is added to make the problem more realistic. Instead of requiring the bulk outlet temperature to be exactly equal to the desired value throughout the cycle, the bulk outlet temperature is required to be within a user-specified tolerance centered around the desired value, allowing for a larger feasible set.

The mixed-integer linear orificing formulation updated with these modifications is given below, where γ is the tolerance on bulk outlet temperature, $\Omega_{i,j,k}$ is the discretized ΔT for assembly i with flow j at step k , and q is the length of x :

$$\begin{aligned}
 \min_{\delta_{ij}} \sum_{j=1}^q \beta_j : \\
 \textcircled{1} \quad & \Delta T_{max} \geq \frac{\alpha_{i,k}}{\dot{m}_i} && \forall k, \forall i \\
 \textcircled{2a} \quad & \sum_{i=1}^n \dot{m}_i \geq \frac{\sum_i \alpha_{i,k}}{\bar{T}_{out,k} + \gamma - T_{in}} && \forall k \\
 \textcircled{2b} \quad & \sum_{i=1}^n \dot{m}_i \leq \frac{\sum_i \alpha_{i,k}}{\bar{T}_{out,k} - \gamma - T_{in}} && \forall k \\
 \textcircled{3} \quad & v_{max} \geq \frac{\dot{m}_i}{\rho A_f} && \forall i \\
 \textcircled{4a} \quad & \sum_{j=1}^q (\delta_{i,j} \Omega_{i,j,k}) - \sum_{j=1}^q (\delta_{i',j} \Omega_{i',j,k}) \leq \xi && \forall i, \forall i' : i \perp i' \\
 \textcircled{4b} \quad & \sum_{j=1}^q (\delta_{i',j} \Omega_{i',j,k}) - \sum_{j=1}^q (\delta_{i,j} \Omega_{i,j,k}) \leq \xi && \forall i, \forall i' : i \perp i' \\
 \textcircled{5} \quad & \dot{m}_i = \sum_{j=1}^q \delta_{i,j} x_j && \forall i \\
 \textcircled{6} \quad & \sum_{j=1}^q \delta_{i,j} = 1 && \forall i \\
 \textcircled{7} \quad & \delta_{i,j} \in \{0, 1\} && \forall i, \forall j \\
 \textcircled{8} \quad & \beta_j \geq \delta_{i,j} && \forall i, \forall j
 \end{aligned}$$

SAMPLE RESULTS

An instance of the model developed above is built in AMPL [8] for the case of the B&B core with 2D shuffling designed by Qvist et al. [5]. A Monte Carlo neutronics model of the core is built with hex-z geometry and smeared assembly compositions, where all assemblies of a given batch have the same smeared composition as outlined in [9]. Coupled neutron-photon transport is used to obtain assembly powers, as photon energy deposition is especially important in low-powered blanket assemblies towards the core periphery. Because all assemblies within a batch have the same composition, the power gradients at the batch interfaces are generally sharper in this model than they would be if the composition of each assembly were explicitly accounted for. This makes it more difficult to limit the difference between outlet temperatures of adjacent assemblies which span the batch interfaces. Therefore, the scenario considered here is likely more challenging than the true equilibrium cycle power profile. The

flowrates in this instance are discretized over a range from 0.05 to 100 kg/s to account for the dramatic differences between the lowest- and peak-assembly powers. The parameters specified for the optimization are: $\xi = 40^\circ\text{C}$, $\Delta T_{max} = 210^\circ\text{C}$, $T_{in} = 355^\circ\text{C}$, $\bar{T}_{out} = 510^\circ\text{C}$, $v_{max} = 12\text{ m/s}$, and $\gamma = 5^\circ\text{C}$. This instance is solved with the CPLEX mixed-integer solver with a branch-and-bound method for a wall-clock of 24 hours. The results are presented in Table I along with the results from the algorithm suggested in [3], where the peak assembly powers over the cycle are used to determine the flow allocated to each assembly.

TABLE I. Selected parameters from the orificing schemes determined through the previously suggested algorithm in [3] and the proposed mixed-integer linear programming approach.

	Number of groups	Max. adjacent outlet temp. difference ($^\circ\text{C}$)	Min. bulk outlet temp. ($^\circ\text{C}$)
Algorithm from [3]	9	143	497
MILP	34	40	505

It is seen that the mixed-integer programming approach requires nearly 4 times as many orifice groups as the algorithmic approach, however the latter does not adhere to the required constraints. Because the algorithmic approach does not provide any means for accounting for the maximum outlet temperature difference between adjacent assemblies, it should not be expected that the approach provide a solution that is feasible in this regard. Additionally, the bulk outlet temperature ends up well below the desired value because the algorithm consistently forces all groups to have an average outlet temperature less than or equal to the desired bulk temperature. The mixed-integer programming approach, however, is not skewed by algorithmic biases, and simply will return an orificing scheme with the lowest number of groups that adheres to all constraints. Therefore, the resulting 34 groups returned by the linear program should be viewed as the absolute smallest number of orifice groups possible for the specified data. In order to reduce the number of groups further, either the constraints must be relaxed or the core design must be changed.

The flowrates determined between the two methods vary substantially, with some assemblies having differences of 56% between the methods. The relative 2-norm difference between the flowrate vectors is 27%, indicating that overall the mixed-integer programming approach tunes the assembly flows much tighter than the algorithmic approach.

Although the algorithmic approach does not adhere to all constraints, it returns a solution within fractions of a second. Conversely, the MILP formulation may take days to solve to provable optimality due to the exponential complexity of branch-and-bound methods. The longer run time for the MILP approach is perhaps outweighed by the simple process to achieve an optimal solution, as compared to the algorithmic approach which might take many design iterations to achieve even a feasible but suboptimal solution. Furthermore, the MILP approach has the advantage of giving an estimate of the degree of suboptimality, whereas the algorithmic approach has no such equivalent.

FUTURE WORK

To make the constraints on pressure drop more explicit, it is planned to incorporate additional constraints using average friction factors. Because pressure drop is typically proportional to the square of the mass flow rate, this additional constraint can be incorporated while keeping the problem convex, although additional computational expense is expected due to the deviation from purely linear constraints. Due to the sharp gradients in power profile, assembly-to-assembly heat transfer could be potentially be significant. It is planned to incorporate a simple heat transfer model between adjacent assemblies so that this phenomenon can be accounted for. In the case that it is, the orificing task will be relaxed because the thermal-stripping constraints will be less restrictive, potentially allowing for less orifice groups for a given core. Additional methods to achieve a smaller number of orifice groups by altering the physical core and plant design will be explored. Finally, methods for accelerating the solution will be implemented, including modifications to the branch-and-bound search and addition of constraints [10] to strengthen the model.

ACKNOWLEDGMENTS

This material is based upon work supported under a Department of Energy Nuclear Energy University Programs Graduate Fellowship.

REFERENCES

1. P. HEJZLAR and ET AL, "Terrapower LLC Traveling Wave Reactor Development Program Overview," *Nuclear Engineering and Technology*, **45**, 6 (2013).
2. M. SMITH (December 2017), personal Communication.
3. F. HEIDET, T. K. KIM, and T. TAIWO, "Orificing Strategy for Sodium-Cooled Fast Reactors," *Transactions of the American Nuclear Society*, **110** (2014).
4. A. ATAMTÜRK and M. W. P. SAVELSBERGH, "Integer-Programming Software Systems," *Annals of Operations Research*, **140**, 1, 67–124 (Nov 2005).
5. S. QVIST, J. HOU, and E. GREENSPAN, "Design and performance of 2D and 3D-shuffled breed-and-burn cores," *Annals of Nuclear Energy*, **85**, 93–114 (2015).
6. A. E. WALTAR, D. R. TODD, and P. V. TSVETKOV, editors, *Fast Spectrum Reactors*, Springer (2012).
7. C. GRANDY (November 2017), Personal Communication.
8. R. FOURER, D. M. GAY, and B. W. KERNIGHAN, *AMPL: A Modeling Language for Mathematical Programming*, Brooks/Cole Publishing (2003).
9. C. KECKLER, M. FRATONI, and E. GREENSPAN, "On the Importance of Control Assembly Layouts in Standing Wave B&B Cores," *Proceedings of PHYSOR*, Cancun, MX, American Nuclear Society (2018).
10. A. ATAMTÜRK and V. NARAYANAN, "Cuts for Conic Mixed-Integer Programming," in M. FISCHETTI and D. P. WILLIAMSON, editors, *Proceedings of 12th International Conference on Integer Programming and Combinatorial Optimization*, Springer, 16–29 (2007).