# An Assessment of Coupling Algorithms for Nuclear Reactor Core Physics Simulations[☆]

Steven Hamilton[a,*], Mark Berrill[a], Kevin Clarno[a], Roger Pawlowski[b]

[a] *Oak Ridge National Laboratory, 1 Bethel Valley Rd., Oak Ridge, TN 37831 U.S.A.*
[b] *Sandia National Laboratories, MS 0316, P.O. Box 5800, Albuquerque, NM 87185 U.S.A.*

**Abstract**

This paper evaluates the performance of multiphysics coupling algorithms on a light water nuclear reactor core simulation. The simulation couples the $k$-eigenvalue form of the neutron transport equation with heat conduction and subchannel flow equations. We compare Picard iteration (block Gauss-Seidel) with multiple variants of preconditioned Jacobian-free Newton-Krylov (JFNK). The performance of the methods are evaluated over a range of energy group sizes, boron concentrations and core power levels. A novel physics-based approximation to a Jacobian-vector product has been developed to mitigate the impact of expensive cross section processing steps. Numerical experiments demonstrating the efficiency of JFNK relative to standard Picard iteration are performed on a 3D model of a nuclear fuel assembly.

*Keywords:* multiphysics, Jacobian-free Newton-Krylov, nuclear reactor analaysis

[☆]Notice: This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

[*]Corresponding Author

*Email addresses:* hamiltonsp@ornl.gov (Steven Hamilton), berrillma@ornl.gov (Mark Berrill), clarnokt@ornl.gov (Kevin Clarno), rppawlo@sandia.gov (Roger Pawlowski)



CASL-U-2014-0149-000-b

## 1. Introduction

Determining the steady-state power and temperature distributions within an operating nuclear reactor is an important component of reactor design and analysis. This task requires simultaneously solving equations describing the distribution of neutrons throughout the reactor in addition to the transfer of heat through the fuel and structural materials and into fluid coolant regions. Current core analysis methods rely on the use of a Picard iteration, alternating between solving individual physics components. Although this approach offers a simple path to coupling different physics codes due to the minimial code interaction required, there are also significant drawbacks. Picard iteration lacks global convergence theory and, at best, a q-linear convergence rate [1]. Additionally, user-defined relaxation schemes are usually required to achieve convergence. Newton-based methods, however, are shown to be globally convergent with q-quadratic convergence rates. The downside to Newton-based methods is that the need for residual and sensitivity information requires more invasive access to application codes. While access to analytical Jacobian matrices is commonly infeasible, Jacobian-free Newton-Krylov (JFNK) methods [2] can be used to realize many of the benefits of Newton-based methods while only requiring evaluation of nonlinear functions. While JFNK methods have been successfully applied in many areas, to date little has been done to evaluate the viability of such methods for multiphysics reactor simulations.

In this study we investigate the use of a JFNK approach to solve multiphysics problems involving coupling between 3D discretizations of the radiation transport and heat transfer equations along with a simple subchannel flow model for modeling of pressurized water reactors. Because one of the dominant costs associated with the current model is the on-line generation of cross section data for use by the neutronics solver, the use of low-cost approximate function evaluations within the JFNK approach is considered. Both criticality ($k$-eigenvalue) and Boron search problems are considered.

The rest of the paper is organized as follows: Section 2 describes the physics models, Section 3 describes various coupling algorithms, Section 4 contains numerical results for a single PWR fuel assembly, and Section 5 presents conclusions and proposals for future ares of investigation.

## 2. Physics Models

In this document we consider solution of multiphysics problems involving coupling between neutron transport and heat transfer. In particular, we focus on the solution of problems involving light water reactors (LWRs). Most of the fundamental ideas described here are applicable to a wide range of reactor types, but certain aspects of the problem, such as geometric features, are particular to LWRs (and possibly pressurized water reactors in particular). For nuclear reactor problems, the standard formulation of the neutron transport equation is the $k$-eigenvalue problem

$$\begin{aligned}\hat{\Omega}\cdot\nabla\psi(\vec{r},E,\hat{\Omega})+\sigma(\vec{r},E,T)\psi(\vec{r},E,\hat{\Omega}) = \\ \int_0^\infty dE'\int_{4\pi}d\hat{\Omega}'\,\sigma_s(\vec{r},E'\rightarrow E,\hat{\Omega}'\rightarrow\hat{\Omega},T)\psi(\vec{r},E',\hat{\Omega}')+ \\ \frac{1}{k}\chi(\vec{r},E)\int_0^\infty dE'\int_{4\pi}d\hat{\Omega}'\,\nu\sigma_f(\vec{r},E',T)\psi(\vec{r},E',\hat{\Omega}')\,, \quad (1)\end{aligned}$$

where $\hat{\Omega}$ is the direction of particle travel, $E$ is the particle energy, $T$ is the temperature of the background material, $\sigma$ is the total cross section, $\sigma_s$ is the scattering cross section, $\nu\sigma_f$ is the neutron production cross section, and $\chi$ is the fission spectrum. The goal is to find the largest value of the eigenvalue $k$ and the corresponding eigenvector $\psi$. Because Eq. (1) represents an eigenvalue problem, the vector $\psi$ has no explicit magnitude. We choose a natural normalization by setting the global heat generation rate (due to nuclear fission occurring in the fuel) to a pre-defined value, i.e.

$$\int dV\int_0^\infty dE\int_{4\pi}d\hat{\Omega}\,\kappa\sigma_f\psi = P^*\,. \quad (2)$$

As noted in Eq. (1), the cross sections are dependent on the temperature of the media, $T$. Thus, for a reactor not operating at a constant temperature it is also necessary to solve a heat conduction equation within the solid fuel and clad regions with fission providing the thermal source, i.e.

$$-\nabla\cdot K(T)\nabla T = \int_0^\infty dE\int_{4\pi}d\hat{\Omega}\,\kappa\sigma_f(E)\psi(E,\hat{\Omega})\,, \quad (3)$$

where $K$ is the material thermal conductivity and $\kappa$ is the heat generated per fission event. Because no fission occurs in the clad regions, the source in

those locations is zero. The exterior surface of the clad is then coupled to the coolant through the subchannel model that solves equations describing the conservation of mass, momentum, and energy, i.e.

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{v}) = 0 \tag{4}$$

$$\frac{\partial \rho v_i}{\partial t} = -\nabla \cdot (\rho v_i \vec{v}) + (-\nabla p + \nabla \cdot \vec{\tau}) - \vec{g} \tag{5}$$

$$\frac{\partial U}{\partial t} + \nabla \cdot (U \vec{v}) = -p \nabla \cdot \vec{v} + \Phi + \nabla K(T) \nabla T + \dot{q}\,, \tag{6}$$

where $\rho$ is the mass density, $\vec{v}$ is the velocity, p is the pressure, $\vec{g}$ is the force exerted by gravity, $\vec{\tau}$ is the viscosity tensor, U is the internal energy density, $\Phi$ is the dissipation function, and $\dot{q}$ is the thermal source. Note that the internal energy density is related to the enthalpy density through $U = h - p$. We solve these equations using a two-equation approximation in which we assume that the coolant flow is only in the axial direction and neglect thermal diffusion between the channels. Assuming steady-state, this reduces to

$$\frac{\partial \rho v_z^2}{\partial z} + \frac{\partial p}{\partial z} - \frac{\partial \tau}{\partial z} = -g\,(-\nabla p + \nabla \cdot \tau) - g\,, \tag{7}$$

$$\frac{\partial h v_z}{\partial z} = -v_z \frac{\partial p}{\partial z} - \Phi + \frac{\partial}{\partial z} \nabla K(T) \frac{\partial T}{\partial z} + \dot{q}\,. \tag{8}$$

The thermal source consists primarily of convective heat transfer from the clad and is the primary coupling mechanism with the temperature of the clad/pellets.

For notational simplicity, in the remainder of this document we will allow $T$ to refer to not only the solution of the thermal diffusion equation in the pellets and clad but also the solution to the subchannel equations. Our goal is therefore to find distributions $\psi$, $T$ and a value $k$ such that Eqs. (1)–(3) are simultaneously satisfied. For ease of notation, we introduce an operator notation for discretized forms of the preceding equations:

$$\mathbf{A}(T)\phi = \lambda \mathbf{B}(T)\phi \tag{9}$$

$$\mathbf{R}_S \mathbf{R}_E \mathbf{B}(T)\phi = P^* \tag{10}$$

$$\mathbf{L}(T)T = \mathbf{R}_E \mathbf{B}(T)\phi\,, \tag{11}$$

where $\mathbf{R}_S$ and $\mathbf{R}_E$ are restriction operators in space and energy, respectively, and $\mathbf{L}(T)$ refers not only to the thermal diffusion equation (3) but also the

subchannel equations of (7)–(8). Here Eq. (9) has taken advantage of the fact that only the angle-integrated variable $\phi$ must generally be stored rather than the corresponding angle-dependent $\psi$. In addition we have written the eigenvalue as $\lambda \equiv \frac{1}{k}$.

### *2.1. Nuclear Data Evaluation*

Nuclear cross section data is generally tabulated at thousands of different energy values and a variety of temperatures for every isotope. Deterministic radiation transport solvers, however, are generally limited to using a small number (typically fewer than 100 and often less than ten) of energy "groups" representing average behavior over a range of energies [3]. In order for this collapsed multigroup data to accurately represent the original energy spectrum, it is necessary to compute the averages using a weighting function that matches the true solution as closely as possible. This weighting function is usually generated by performing several local 1-D (pincell) or 2-D (lattice) transport calculations to capture local nearest-neighbor contributions to the solution. Such a calculation would need to be performed for every unique type of fuel pin in the problem. In addition, because the cross section data is temperature-dependent, the variation in the collapsed data with temperature must be taken into account. One possible approach is to precompute multigroup data at several different temperatures and then interpolate the values to the temperature of interest. A more accurate approach, however, is to perform a separate cross section processing calculation at every different temperature under consideration. The latter approach of "on-line" generation of cross sections is the method considered in this paper.

### *2.2. Critical Boron Search*

A nuclear reactor operating under normal conditions is always maintained in a critical state, i.e. $k \equiv 1$. Because of this, rather than compute the value of $k$ for a given reactor configuration it is sometimes preferred to compute the value of some parameter which results in a critical configuration. In a pressurized water reactor (PWR) this parameter is often the concentration of soluble boron in the coolant. In this case, the eigenvalue is no longer an unknown and the matrix $\mathbf{A}$ becomes a function of the boron concentration, $C$. Thus Eq. (9) becomes

$$\mathbf{A}(T, C)\phi = \lambda_{\text{target}}\mathbf{B}(T)\phi \tag{12}$$

where $\lambda_{\text{target}}$ is the target eigenvalue for the search, which will typically be unity for full core calculations but may be another value for calculations involving only a portion of a reactor. Together with Eqs. (10) and (11), Eq. (12) represents the problem statement for a critical boron search.

We briefly note that other single-parameter critical searches (e.g. control rod height) can be written in a form similar to (12), with the matrices $\mathbf{A}$ and possibly $\mathbf{B}$ depending on the given parameter. Additional critical searches will not be considered further in this paper.

## 3. Solution Approaches

The radiation transport and heat transfer systems described in Eqs. (9)–(11) can be written as a single system of coupled equations as

$$f\begin{pmatrix}\phi\\ \lambda\\ T\end{pmatrix} = \begin{bmatrix} f_\phi(\phi,\lambda,T)\\ f_\lambda(\phi,T)\\ f_T(\phi,T)\end{bmatrix} = \begin{bmatrix}\mathbf{A}(T)\phi - \lambda\mathbf{B}(T)\phi\\ \mathbf{R}_S\mathbf{R}_E\mathbf{B}(T)\phi - P^*\\ \mathbf{L}(T)T - \mathbf{R}_E\mathbf{B}(T)\phi\end{bmatrix} = \mathbf{0}\,. \tag{13}$$

Alternatively, to solve the boron search problem described in Section 2.2, the system can be written as

$$f\begin{pmatrix}\phi\\ C\\ T\end{pmatrix} = \begin{bmatrix} f_\phi(\phi,C,T)\\ f_C(\phi,T)\\ f_T(\phi,T)\end{bmatrix} = \begin{bmatrix}\mathbf{A}(T,C)\phi - \lambda_{\text{target}}\mathbf{B}(T)\phi\\ \mathbf{R}_S\mathbf{R}_E\mathbf{B}(T)\phi - P^*\\ \mathbf{L}(T)T - \mathbf{R}_E\mathbf{B}(T)\phi\end{bmatrix} = \mathbf{0}\,. \tag{14}$$

In this section we explore strategies for solving these systems of equations.

### *3.1. Picard Iteration*

One straightforward approach to solving this system of equations is to alternate between solves of the individual physics in a block Gauss-Seidel approach. This Picard iteration can be written as shown in Alg. 1. It has been observed in several studies with various physics approximations and/or discretizations that this simple iteration scheme applied to light water reactor problems is prone to poor convergence and possibly divergence due to oscillations induced by certain error modes. The standard remedy for this issue is to introduce a damping parameter, $\omega$, such that Algorithm 1 is replaced by the modified version given in Alg. 2. Optimal values for the damping parameter are typically between 0.3 and 0.6 [4, 5, 6]. Note that it

**Algorithm 1** Picard Iteration (Gauss-Seidel)

Given $T_0$
**for** $m = 0, 1, \ldots$ until converged **do**
  Solve $\mathbf{A}(T^m)\hat{\phi} = \lambda \mathbf{B}(T^m)\hat{\phi}$ for $\hat{\phi}$, $\lambda$
  Set $\phi^{m+1} = \frac{P^*}{\mathbf{R}_S \mathbf{R}_E \mathbf{B}(T^m)\hat{\phi}}\hat{\phi}$
  Solve $\mathbf{L}(T^{m+1})T^{m+1} = \mathbf{R}_E \mathbf{B}(T^m)\phi^{m+1}$ for $T^{m+1}$
**end for**

**Algorithm 2** Damped Picard Iteration

Given $T_0$
**for** $m = 0, 1, \ldots$ until converged **do**
  Solve $\mathbf{A}(T^m)\hat{\phi} = \lambda \mathbf{B}(T^m)\hat{\phi}$ for $\hat{\phi}$, $\lambda$
  Set $\phi^{m+1} = \frac{P^*}{\mathbf{R}_S \mathbf{R}_E \mathbf{B}(T^m)\hat{\phi}}\hat{\phi}$
  Solve $\mathbf{L}(\widehat{T})\widehat{T} = \mathbf{R}_E \mathbf{B}(T^m)\phi^{m+1}$ for $\widehat{T}$
  Update $T^{m+1} = \omega\widehat{T} + (1-\omega)T^m$
**end for**

is also possible to perform damping on the scalar flux (or power) instead of temperature.

Because Eq. (12) does not represent a linear eigenvalue problem the way that Eq. (9) does, solving the coupled system of equations using Picard iteration is slightly more complicated. One approach is to use a traditional solution approach (e.g. power iteration) to solve the original eigenvalue problem, but modify the algorithm so that the boron concentration is updated instead of the eigenvalue [7, 8]. This approach, however, presents two problems. First, subspace eigenvalue solvers (e.g. Arnoldi's method or a generalized Davidson method) that are commonly far more efficient than power iteration can no longer be used. Second, every time the boron concentration is modified the problem cross sections must be updated which can be a significant computational expense if cross sections are being processed on-line. A more common approach is to solve the boron search problem indirectly by solving Eqs. (9)–(11), but periodically updating the boron concentration such that the eigenvalue at convergence is unity, resulting in Alg. 3. In this approach, an estimate for the derivative of the reactivity with respect to the boron concentration, $\frac{\partial \rho}{\partial C}$, must be provided.

**Algorithm 3** Damped Picard Iteration with Boron Search

Given $T^0$, $C^0$, $\frac{\partial \rho}{\partial C}$
**for** $m = 0, 1, \ldots$ until converged **do**
  Solve $\mathbf{A}(T^m, C^m)\hat{\phi} = \lambda \mathbf{B}(T^m)\hat{\phi}$ for $\hat{\phi}$, $\lambda$
  Set $\phi^{m+1} = \frac{P^*}{\mathbf{R}_S \mathbf{R}_E \mathbf{B}(T^m)\hat{\phi}}\hat{\phi}$
  Set $C^{m+1} = C^m - \frac{1}{\frac{\partial \rho}{\partial C}}(\lambda_{\text{target}} - \lambda)$
  Solve $\mathbf{L}(\widehat{T})\widehat{T} = \mathbf{R}_E \mathbf{B}(T^m)\phi^{m+1}$ for $\widehat{T}$
  Update $T^{m+1} = \omega \widehat{T} + (1 - \omega)T^m$
**end for**

**Algorithm 4** Newton's Method

Given $x_0$
**for** $m = 0, 1, \ldots$ until converged **do**
  Solve (approximately) $\mathbf{J}(x^m)\delta x = -f(x^m)$ for $\delta x$
  Set $x^{m+1} = \delta x + x^m$
**end for**

### *3.2. Jacobian-Free Newton-Krylov*

One approach for improving on the behavior of Picard iteration for solving Eq. (13) is to use Newton's method (Algorithm 4). Computation of the Jacobian matrix, $\mathbf{J}$, is generally not possible because the temperature dependence of nuclear data is not available in closed form, but is rather the result of solver numerous local transport problems as described in Section 2.1. The availability of a nonlinear function but difficulty in formulating the corresponding Jacobian matrix suggests that the use of Jacobian-free Newton Krylov (JFNK) methods might be appropriate [9, 2]. JFNK methods are based on two primary ideas. First, if the Newton correction equation is solved using a Krylov subspace method, then access to the full Jacobian matrix is not necessary, only the action of the Jacobian applied to a vector is required. The second idea is that the product of the Jacobian and a given vector can be approximated using a finite difference approach, e.g.

$$\mathbf{J}(u^m)v \approx \frac{f(u^m + \epsilon v) - f(u^m)}{\epsilon}. \tag{15}$$

Note that second (or higher) order approximations are possible, but the form in Eq. (15) is far more common because it requires only a single function

evaluation per Jacobian-vector product because the value of $f(u^m)$ can be computed and stored once per Newton iteration so that each subsequent Jacobian-vector product requires only the evaluation of $f(u^m + \epsilon v)$. The JFNK method therefore allows an approximate Newton method to be performed using only evaluations of the nonlinear function. Discussions on the selection of $\epsilon$ can be found in Refs. [2] and [10]. JFNK methods have been used previously to solve the $k$-eigenvalue problem by itself [11, 12] as well as time-dependent multiphysics problems involving radiation transport coupled with thermal-hydraulics [13], though to date we are unaware of any studies considering steady-state coupling of radiation transport with thermal-hydraulics capabilities.

*3.2.1. JFNK Preconditioning*

In order for JFNK to be competitive with other solution approaches, it is necessary to efficiently solve the Newton correction equation. Because a Krylov method is used to solve the linear system involving the Jacobian matrix, the development of effective preconditioning strategies is vital. Although the use of JFNK eliminates the need to explicitly form or store the Jacobian matrix, some knowledge about the Jacobian is still beneficial in the construction of a preconditioning. The true Jacobian corresponding to the function defined by (13) can be written as

$$\mathbf{J}\begin{pmatrix}\phi\\ \lambda\\ T\end{pmatrix} = \begin{bmatrix} \mathbf{A}(T) - \lambda\mathbf{B}(T) & -\mathbf{B}(T)\phi & \dfrac{\partial(\mathbf{A}(T)\phi - \lambda\mathbf{B}(T)\phi)}{\partial T} \\ \mathbf{R}_S\mathbf{R}_E\mathbf{B}(T) & 0 & \dfrac{\partial(\mathbf{R}_S\mathbf{R}_E\mathbf{B}(T)\phi)}{\partial T} \\ -\mathbf{R}_E\mathbf{B}(T) & 0 & \dfrac{\partial(\mathbf{L}(T)T - \mathbf{R}_E\mathbf{B}(T)\phi)}{\partial T} \end{bmatrix}, \quad (16)$$

where the entries in the last column of this block matrix are left as simply partial derivatives with respect to temperature to indicate that these terms are generally not available in closed form due to the dependence on material properties. Similarly, the Jacobian matrix corresponding to (14) can be

written as

$$\mathbf{J}\begin{pmatrix}\phi\\ C\\ T\end{pmatrix} = \begin{bmatrix} \mathbf{A}(T) - \lambda_{\text{target}}\mathbf{B}(T) & \dfrac{\partial\left(\mathbf{A}(T)\phi\right)}{\partial C} & \dfrac{\partial(\mathbf{A}(T)\phi - \lambda\mathbf{B}(T)\phi)}{\partial T} \\ \mathbf{R}_S\mathbf{R}_E\mathbf{B}(T) & 0 & \dfrac{\partial(\mathbf{R}_S\mathbf{R}_E\mathbf{B}(T)\phi)}{\partial T} \\ -\mathbf{R}_E\mathbf{B}(T) & 0 & \dfrac{\partial(\mathbf{L}(T)T - \mathbf{R}_E\mathbf{B}(T)\phi)}{\partial T} \end{bmatrix}, \tag{17}$$

A simple approach to preconditioning a linear system involving (16) or (17) is a block diagonal approach in which preconditioners for each physics are applied independently, i.e.

$$\mathbf{P}_D = \begin{bmatrix} \widehat{\mathbf{A}} - \lambda\widehat{\mathbf{B}} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \widehat{\mathbf{L}} \end{bmatrix}, \tag{18}$$

where $\widehat{(\cdot)}$ indicates that some approximation of the operator is used in the preconditioner. Possible choices for this approximation include incomplete factorizations, algebraic multigrid methods, or physics-based approximations. This preconditioner selection has the advantages of being relatively simple to construct and inexpensive to apply (assuming the approximations to the individual physics operators are inexpensive). Additionally, the block diagonal structure of this preconditioner means that the application of the different physics components can be done independently and simultaneously, allowing the possibility of treatment of the physics domains in parallel. Neglecting all terms in the Jacobian that correspond to coupling between different physics components may be lead to a reduction in the effectiveness of the preconditioner. Therefore it may be beneficial to capture some of the off-diagonal terms from the true Jacobian while maintaining enough structure to allow efficient application of the preconditioner. Possible approaches include a block lower triangular preconditioner given by

$$\mathbf{P}_L = \begin{bmatrix} \widehat{\mathbf{A}} - \lambda\widehat{\mathbf{B}} & 0 & 0 \\ 0 & 1 & 0 \\ -\mathbf{R}_E\mathbf{B} & 0 & \widehat{\mathbf{L}} \end{bmatrix}, \tag{19}$$

or a block upper triangular approach given by

$$\mathbf{P}_U = \begin{bmatrix} \widehat{\mathbf{A}} - \lambda\widehat{\mathbf{B}} & -\mathbf{B}\phi & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \widehat{\mathbf{L}} \end{bmatrix} . \tag{20}$$

Note that the upper triangular preconditioner of (20) applies only to the Jacobian (16). Development of preconditioning strategies that account for coupling behavior between physics components has the potential to significantly improve the convergence behavior of the JFNK approach and the development of such techniques is an attractive area of research for future investigations.

*3.2.2. Approximate Jacobian-Vector Products*

As noted in Section 2, the cross section data appearing in Eq. (1) is dependent on the material temperature, resulting in the temperature dependence of the matrices **A** and **B**. Generating data suitable for use in a deterministic radiation transport solver generally involves performing a large number of small (1-D or 2-D) radiation transport calculations involving many energy groups. The results of these calculations are then used to compute effective data involving fewer energy groups and potentially averaged spatially. Because the tabulated data used in these small calculations is temperature-dependent, a new calculation must be performed for every region in the reactor for which a distinct temperature is defined. Furthermore, every time the temperature of a given region is modified, a new calculation must be performed. Although each individual calculation represents a relatively small computational burden, the large number of calculations that may be required in a given simulation may result in a large portion of the overall computational effort being spent in this cross section processing. Using the JFNK method, evaluation of the function in Eq. (13) at a perturbed state point appears as

$$\widehat{f}\begin{pmatrix} \phi + \Delta\phi \\ \lambda + \Delta\lambda \\ T + \Delta T \end{pmatrix} = \begin{bmatrix} \mathbf{A}(T+\Delta T)(\phi+\Delta\phi) - (\lambda+\Delta\lambda)\mathbf{B}(T+\Delta T)(\phi+\Delta\phi) \\ \mathbf{R}_S\mathbf{R}_E\mathbf{B}(T+\Delta T)(\phi+\Delta\phi) - P^* \\ \mathbf{L}(T+\Delta T)(T+\Delta T) - \mathbf{R}_E\mathbf{B}(T+\Delta T)(\phi+\Delta\phi) \end{bmatrix} , \tag{21}$$

and similarly evaluation of Eq. (14) results in

$$\widehat{f}\begin{pmatrix}\phi+\Delta\phi\\ C+\Delta C\\ T+\Delta T\end{pmatrix}=\begin{bmatrix}\mathbf{A}(T+\Delta T,C+\Delta C)(\phi+\Delta\phi)-\lambda_{\text{target}}\mathbf{B}(T+\Delta T)(\phi+\Delta\phi)\\ \mathbf{R}_S\mathbf{R}_E\mathbf{B}(T+\Delta T)(\phi+\Delta\phi)-P^*\\ \mathbf{L}(T+\Delta T)(T+\Delta T)-\mathbf{R}_E\mathbf{B}(T+\Delta T)(\phi+\Delta\phi)\end{bmatrix}, \tag{22}$$

where the appearance of $\mathbf{A}(T+\Delta T)$ and $\mathbf{B}(T+\Delta T)$ indicate the need to recompute cross sections based on the perturbed temperature distribution and boron concentration at every linear (Krylov) iteration. This is in contrast to Picard iteration, where cross sections need only be updated once every outer iteration. With the expectation that several Newton iterations will be required to converge a given problem and potentially dozens of linear iterations will be necessary for every Newton iteration, the time spent simply updating cross sections is likely to be prohibitive with a straightforward application of JFNK.

In order to circumvent this potential bottleneck, we propose using an approximate Newton update equation,

$$\widehat{\mathbf{J}}(u^m)\delta^m=-f(u^m)\,, \tag{23}$$

where $\widehat{\mathbf{J}}$ indicates that an approximation to the Jacobian is used. In contrast to inexact Newton methods which involve solving the true Newton correction equation in an approximate manner [14], the current approach is more accurately described as a preconditioned nonlinear Richardson iteration [15]. In the interest of maintaining the attractive matrix-free nature of JFNK, we can determine a corresponding approximate function evaluation, $\widehat{f}$, such that a finite difference operation approximates a product with $\widehat{\mathbf{J}}$ rather than the full Jacobian, i.e.

$$\widehat{\mathbf{J}}(u^m)v\approx\frac{\widehat{f}(u^m+\epsilon v)-f(u^m)}{\epsilon}\,, \tag{24}$$

where $\widehat{f}$ is an approximation to Eq. (13) that does not require recalculation of cross section data. This approach is consistent with strategies studied in Ref. [16]. It should be noted that the full nonlinear function, including processing of all cross sections, must still be evaluated once every nonlinear iteration and is used to evaluate convergence of the nonlinear iterations.

One possibility for approximating the Jacobian is to simply neglect the

temperature variation of cross sections during an approximate Jacobian-vector product, corresponding to an approximate Jacobian of

$$\widehat{\mathbf{J}} \begin{pmatrix} \phi \\ \lambda \\ T \end{pmatrix} = \begin{bmatrix} \mathbf{A}(T) - \lambda \mathbf{B}(T) & -\mathbf{B}(T)\phi & 0 \\ \mathbf{R}_S \mathbf{R}_E \mathbf{B}(T) & 0 & 0 \\ -\mathbf{R}_E \mathbf{B}(T) & 0 & \dfrac{\partial(\mathbf{L}(T)T)}{\partial T} \end{bmatrix}, \tag{25}$$

or equivalently an approximate function evaluation of

$$\widehat{f} \begin{pmatrix} \phi + \Delta\phi \\ \lambda + \Delta\lambda \\ T + \Delta T \end{pmatrix} = \begin{bmatrix} \mathbf{A}(T)(\phi + \Delta\phi) - (\lambda + \Delta\lambda)\mathbf{B}(T)(\phi + \Delta\phi) \\ \mathbf{R}_S \mathbf{R}_E \mathbf{B}(T)(\phi + \Delta\phi) - P^* \\ \mathbf{L}(T + \Delta T)(T + \Delta T) - \mathbf{R}_E \mathbf{B}(T)(\phi + \Delta\phi) \end{bmatrix}. \tag{26}$$

This selection, however, results in no updated temperature information being communicated from the heat transfer solver to the neutronics domain during a given nonlinear iteration. This lack of information is expected to have a detrimental effect on the convergence behavior of the nonlinear solver.

It is possible to include some information about the temperature feedback effect on cross sections without performing a full cross section processing step at each function evaluation. One such approach can be accomplished by noting that the strongest temperature feedback effect is due to an increase in absorption with increasing temperature (largely due to Doppler broading in $^{238}$U). As shown in Fig. 1, the temperature dependence of absorption cross sections is approximately linear over a wide range of temperatures. This suggests a modification to the Jacobian approximation of Eq. (25) which uses a linear approximation to the temperature dependence of the absorption cross section and neglects the temperature dependence of all other cross sections:

$$\widehat{\mathbf{J}} \begin{pmatrix} \phi \\ \lambda \\ T \end{pmatrix} = \begin{bmatrix} \mathbf{A}(T) - \lambda \mathbf{B}(T) & -\mathbf{B}(T)\phi & \dfrac{\partial \sigma_a}{\partial T}\phi \\ \mathbf{R}_S \mathbf{R}_E \mathbf{B}(T) & 0 & 0 \\ -\mathbf{R}_E \mathbf{B}(T) & 0 & \dfrac{\partial(\mathbf{L}(T)T)}{\partial T} \end{bmatrix}. \tag{27}$$

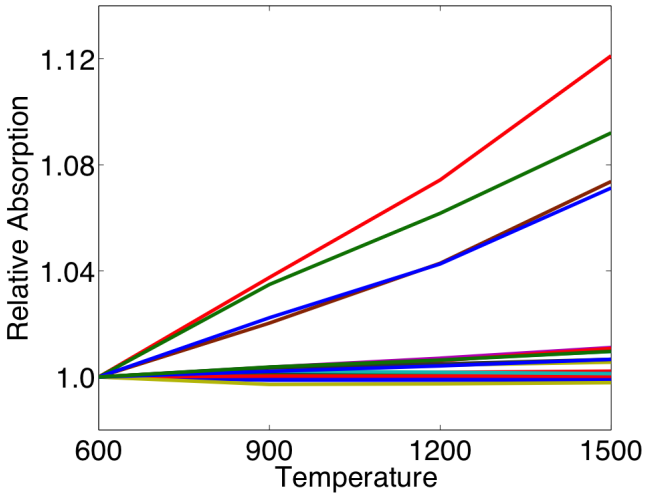


Figure 1: Temperature dependence of 23 group homogenized absorption cross sections for a 3.1% enriched PWR fuel pin. Each curve represents the absorption cross section in a different energy group relative to the group value at $T = 600K$.

The function evaluation corresponding to Eq. (27) is

$$\widehat{f}\begin{pmatrix}\phi+\Delta\phi\\ \lambda+\Delta\lambda\\ T+\Delta T\end{pmatrix} = \begin{bmatrix}\mathbf{A}(T)(\phi+\Delta\phi)-(\lambda+\Delta\lambda)\mathbf{B}(T)(\phi+\Delta\phi)+\dfrac{\partial\sigma_a}{\partial T}\phi\Delta T\\ \mathbf{R}_S\mathbf{R}_E\mathbf{B}(T)(\phi+\Delta\phi)-P^*\\ \mathbf{L}(T+\Delta T)(T+\Delta T)-\mathbf{R}_E\mathbf{B}(T)(\phi+\Delta\phi)\end{bmatrix}. \tag{28}$$

## 4. Results

For the numerical experiments in this paper, the radiation transport equation is approximated using the simplified $P_N$ ($SP_N$) angular approximation [17, 18, 19]. These equations are discretized spatially using a finite volume approach as implemented in the Denovo package [20]. The simple nature of the $SP_N$ equations offers significant advantages in the current study: the operators $\mathbf{A}$ and $\mathbf{B}$ from Eq. (9) can be explicitly constructed as sparse matrices and therefore algebraic preconditioners (algebraic multigrid, for instance) can be easily applied. Aside from the development of appropriate preconditioners, it is expected that similar behavior would be observed for other

transport formulations (such as discrete ordinates). The cross sections used by the $SP_N$ equations are generated by the XSProc module of the SCALE package [21]. The heat transfer and subchannel equations are solved using the Advanced Multiphysics (AMP) package [22]. The heat transfer equation is discretized using standard trilinear continuous Galerkin finite elements and the subchannel equations employ a finite difference approximation [23]. Validation of the pin heat transfer models in AMP has been performed through comparison to data from several experiments [24, 24]. Preliminary investigation into multiphysics coupling involving AMP and Denovo was performed in Ref. [25].

To test the behavior of different nonlinear solvers on a realistic problem, we consider the solution of CASL AMA Progression Problem 6 [26]. This problem consists of a single $17 \times 17$ PWR fuel assembly with 264 fuel pins containing 3.1% enriched $UO_2$, 24 guide tubes and a single central instrumentation tube. Eight spacer grids are located along the axial length of the assembly, as well as upper and lower assembly nozzles. A full description of the problem, including detailed material and geometric details, is contained in Ref. [26].

For the base configuration, we model the assembly at 17.67 MW and 1300 ppm dissolved boron; the effect of power level and boron concentration on solver convergence will be studied later in this section. The base configuration uses a 56 energy group cross section library and uses the XSProc module of the SCALE package to collapse these cross sections to 23 groups for Denovo $SP_N$ calculations. Distinct cross sections are used for each fuel pin and for each of 49 axial levels. An $SP_3$ angular order (containing 2 angular moments) is used for all calculations, along with $P_1$ scattering. A $2 \times 2$ spatial mesh per pin cell is used in the x–y plane, with a maximum axial mesh size of 2 cm, resulting in 290,156 mesh cells for the full assembly. The AMP heat transfer problem contains 15,504 mesh cells per fuel pin (over both the fuel and clad meshes), resulting in approximately 4.1 million total cells. A linear continuous finite element discretization of the heat transfer problem is used. All problems in this study are executed in parallel on 289 processing cores on the OLCF EOS cluster, resulting in 1 fuel pin per core.

Five different solver approaches are considered. First is a damped Picard iteration with the damping applied to the temperature component of the solution as described in Section 3. Except where otherwise noted, all calculations use a damping factor of $\omega = 0.45$ which appears to produce nearly optimal convergence behavior for a wide range of problems. Within Picard iteration

the $k$-eigenvalue problem is solved with a generalized Davidson eigensolver, which has been shown to be highly efficient for solving the $k$-eigenvalue problem [27, 28], and the heat transfer and subchannel flow equations are solved simultaneously using a JFNK approach. The next solver consists of a full JFNK solver using the nonlinear function evaluation described by Eq. (13). Recall that this selection requires performing full cross section processing at every *linear* iteration. The final two solvers are modified JFNK approaches based on the function evaluations in Section 3.2.2. The approach denoted by MJFNK1 corresponds to the function evaluation of Eq. (26) in which the temperature dependence of the cross sections is entirely neglected during a nonlinear iteration. The approach denoted by MJFNK2 corresponds to Eq. (28) in which a linear approximation to the temperature dependence of the absorption cross section is used during each nonlinear iteration. This linear approximation is determined by performing stand-alone XSProc calculations at 600K and 1500K to compute pin-homogenized absorption cross sections. In both modified JFNK approaches, only a single cross section processing step is performed per nonlinear iteration.

A brief mention should be made of the memory usage of the various solution approaches. Picard iteration itself uses virtually no memory on its own (only a single extra copy of the solution vector to determine convergence), and therefore the memory usage is determined by the individual solvers. In this study, the JFNK solver used by the heat transfer solver internally uses a GMRES linear solver which required a subspace containing up to 20 vectors containing the thermal solution. The generalized Davidson solver used in the $\mathrm{SP}_N$ calculations requires multiple subspaces, which combined required the storage of up to 75 vectors of the length of the $\mathrm{SP}_N$ solution vector. Memory requirements for Anderson acceleration are equivalent to Picard except that a small number of addition vectors (in this study, 2) are used by the solver. For the JFNK-based methods, the primary memory requirement is through the GMRES linear solver used. In this study, no restarting was employed resulting in a subspace size of as many as 80 vectors, each having the combined length of the thermal and $\mathrm{SP}_N$ solutions. Note that all of the GMRES linear solvers as well as the generalized Davidson offer the possibility of using more aggressive restarting capabilities to limit the size of the subspaces that are used, typically at a cost of performing a small number of additional iterations to reach the same convergence criteria. In addition to memory associated with solver subspaces, there is an additional memory cost associated with forming the problem operators and corresponding pre-

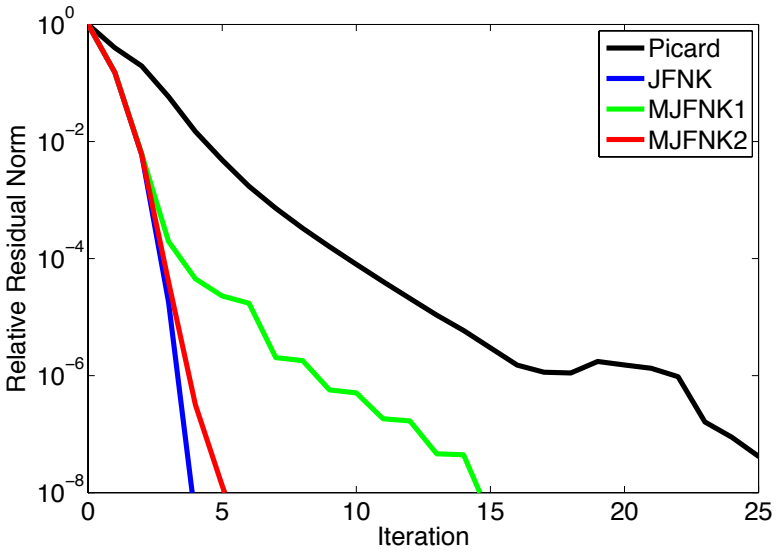


Figure 2: Nonlinear convergence behavior for Picard and JFNK methods.

conditioners. These costs, however, are consistent across every solver option because consistent parameters and preconditioning options were used.

Figure 2 shows the convergence behavior for Picard iteration and each of the JFNK variants. Picard iteration demonstrates the expected linear convergence rate and convergence for the full JFNK is quadratic. The first modified JFNK method follows the convergence behavior of the full JFNK for the first few iterations before separating and ultimately converging linearly with a rate similar to that of Picard iteration. The second modified JFNK follows the convergence rate of full JFNK even more closely and although it eventually deviates from quadratic convergence it does result in linear convergence at a very rapid rate. It should be noted that the convergence criteria used to generate this plot were several orders of magnitude tighter than what is typically applied to such multiphysics problems. This was done to emphasize the pertinent convergence features of each method.

The convergence behavior of Picard iteration as a function of damping parameter at several different power levels is shown in Figure 3. These curves have the same general shape as the corresponding plot from Ref. [25] and consistent with the behavior observed in Refs. [4, 5, 6], all of which reported using damping factors between 0.3 and 0.6. An interesting feature of this curve is that the convergence behavior depends on the power level, with high

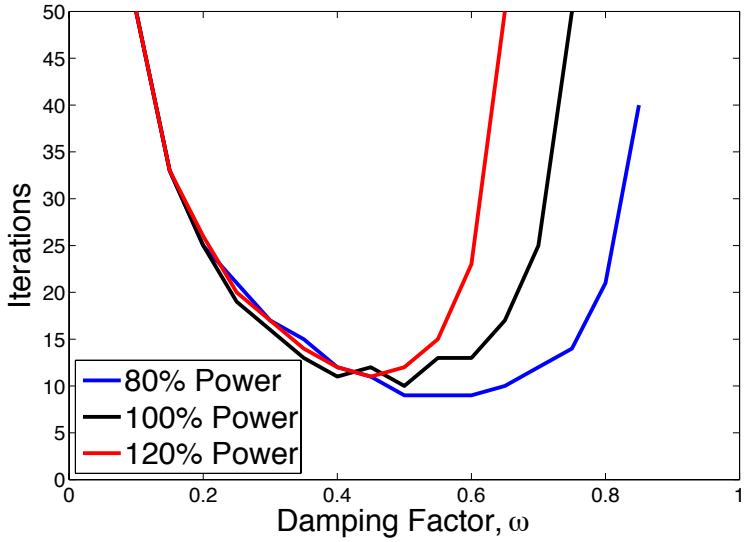


Figure 3: Picard convergence vs. damping parameter for different power levels.

power levels requiring smaller damping factors for optimal convergence.

Table 1 provides the time required to achieve a relative convergence tolerance of $10^{-4}$ for each solution strategy at four different power levels. The time required for Picard iteration to converge is approximately constant for most power levels, showing a slight upward trend at higher power levels indicative of the stronger coupling present at high power. Despite the fact that JFNK exhibits quadratic convergence behavior, the large cost per iteration due to performing cross section processing at every linear iteration results in large run times that are not competitive with other approaches. The first modified JFNK method performs very well at low power levels, significantly reducing runtimes relative to Picard iteration. At high power levels, however, the time required for convergence greatly increases, reaching over twice the runtime of Picard at 120% power. This behavior is easily understood by noting that the low heat generation rate at low power results in relatively small changes in temperature and therefore the effect of neglecting the temperature dependence of cross sections is not too large, but high power levels produce large temperature variations and a corresponding degradation due to not capturing the effects of changes in cross sections. The second modified JFNK approach performs very well across all power levels, resulting in the fastest time to solution for all cases and displaying very little variation with

| Method | Power Level | | | |
|---|---|---|---|---|
| | 60% | 80% | 100% | 120% |
| Picard | 1296 | 1150 | 1353 | 1523 |
| JFNK | 6168 | 5958 | 8908 | 8844 |
| MJFNK1 | 734 | 1035 | 1693 | 3164 |
| MJFNK2 | 695 | 985 | 971 | 961 |

Table 1: Solver timing in seconds as function of power level.

| Method | Boron Concentration | | | |
|---|---|---|---|---|
| | 0 ppm | 600 ppm | 1300 ppm | 2000 ppm |
| Picard | 1370 | 1604 | 1353 | 1445 |
| JFNK | 8079 | 9200 | 8908 | 8391 |
| MJFNK1 | 2743 | 3858 | 1693 | 1704 |
| MJFNK2 | 1325 | 1008 | 971 | 1002 |

Table 2: Solver timing in seconds as function of boron concentration.

power level from 80–100%.

Table 2 shows the time to reach convergence for each solver approach at several different boron concentrations. The convergence behavior does not appear to be very strongly dependent on the boron concentration, the first three solvers produced the longest runtime at 600 ppm boron while MFJNK2 had the longest runtime at 0 ppm. For MJFNK1, there is over a factor of two difference in the time to solution at 600 ppm versus 2000 ppm; the other solvers do not experience such strong dependence. It should be noted that for MJFNK2 a single linear approximation to the absorption cross section is used for all boron concentrations, indicating a high level of robustness of the approach to this parameter.

Table 3 shows the overall timing results for three different combinations of cross section processing and transport energy group structures, effectively modifying the cost of cross section processing from relatively cheap to very expensive. When eight group cross section processing is performed, the computational time for most methods is dominated by the time spent in the thermal (and subchannel) portions of the calculation. In this case the full JFNK approach is still slower than Picard iteration, though it is easy to en-

| Method | XSProc Groups/$SP_N$ Groups | | |
|---|---|---|---|
| | 8/8 | 56/23 | 252/23 |
| Picard | 980 | 1353 | 3868 |
| JFNK | 1685 | 8908 | 51727 |
| MJFNK1 | 1070 | 1693 | 3793 |
| MJFNK2 | 636 | 971 | 2013 |

Table 3: Solver timing in seconds for different energy group structures.

vision that improvements in preconditioning could lead to such an approach offering an advantage relative to Picard. However, this coarse energy group structure does not produce results with appropriate accuracy and therefore finer energy group stuctures are needed. As the number of energy groups is increased, the time spent in cross section processing quickly dominates the overall runtime. This is especially dramatic in the case of 252 group cross section processing in which the time required for the full JFNK approach is more than an order of magnitude greater than for Picard iteration. Thus, for the large numbers of energy groups typically required for accurate reactor physics calculations it is apparent that naïve application of JFNK will result in far worse behavior than Picard iteration. By appealing to modified JFNK approaches that avoid frequent cross section processing, however, significant reduction in computational time relative to Picard is possible, with around a factor of two reduction evident for MJFNK2 in the 252 group XSProc case.

In order to more thoroughly assess the merits of the different solvers, we now examine the amount of time spent in each physics component as well as the number of times each physics component was applied. Tables 4, 5, and 6 provide this detailed information for the same three different energy group structures; Note that solver setup, mesh transfer, and related operations are not included in any component and therefore the sum of the time spent in the various physics components will sum to less than the total time.

Comparing the individual columns of these tables, we can discern exactly where the computational gains of the modified JFNK approach are realized. In all cases, the amount of time spent in the thermal operator and cross section processing are greatly reduced relative to Picard but the time spent in the $SP_N$ operator is reduced only very little or possibly increased. The reason for this is that Picard iteration is able to utilize a generalized Davidson eigensolver that has been demonstrated to be highly efficient for stand-alone

| Method | Total | $SP_N$ | Thermal | XSProc |
|---|---|---|---|---|
| Picard | 980 (11) | 32 (161) | 851 (302) | 65 (11) |
| JFNK | 1685 (5) | 34 (203) | 444 (203) | 1075 (203) |
| MJFNK1 | 1070 (8) | 61 (371) | 763 (371) | 40 (8) |
| MJFNK2 | 636 (5) | 34 (208) | 427 (208) | 30 (5) |

Table 4: Timing in seconds (operator applies) by component for 8 group XSProc, 8 group $SP_N$.

| Method | Total | $SP_N$ | Thermal | XSProc |
|---|---|---|---|---|
| Picard | 1353 (11) | 178 (273) | 832 (298) | 312 (11) |
| JFNK | 8908 (5) | 160 (258) | 649 (258) | 7912 (258) |
| MJFNK1 | 1693 (8) | 259 (434) | 898 (434) | 257 (7) |
| MJFNK2 | 971 (5) | 137 (228) | 470 (228) | 171 (5) |

Table 5: Timing in seconds (operator applies) by component for 56 group XSProc, 23 group $SP_N$.

$k$-eigenvalue calculations. Multiphysics coupling involving radiation transport solvers that do not have optimized eigensolvers or using discretizations that do not easily lend themselves to a generalized Davidson approach would be expected to experience more significant timing reduction in the transport portion of the calculation.

## 5. Conclusions

In this study we have provided an assessment of several different nonlinear solvers for use in problems involving coupled neutronics and thermal

| Method | Total | $SP_N$ | Thermal | XSProc |
|---|---|---|---|---|
| Picard | 3868 (13) | 157 (244) | 939 (332) | 2739 (13) |
| JFNK | 51727 (5) | 139 (228) | 591 (228) | 50817 (228) |
| MJFNK1 | 3793 (10) | 313 (523) | 1103 (523) | 2049 (10) |
| MJFNK2 | 2013 (5) | 133 (223) | 466 (223) | 1221 (5) |

Table 6: Timing in seconds (operator applies) by component for 252 group XSProc, 23 group $SP_N$.

hydraulics. In particular, comparisons of damped Picard iteration, Anderson acceleration, and Jacobian-free Newton-Krylov have been performed. Because a naïve implementation of JFNK results in performing a very large number of cross section processing steps, two modified variants of JFNK have been introduced which only require processing cross sections at each nonlinear iteration. Numerical results on CASL AMA Problem 6 indicate that:

- if on-line generation of cross sections represents a large portion of the runtime of a calculation then a direct JFNK implementation results in a prohibitively large number of cross section processing steps,
- approximations to the nonlinear operator that avoid cross section processing can largely preserve the fast convergence rate of JFNK without the overhead,
- computational savings of 1.4-2 relative to Picard iteration with Newton-like methods are possible, and
- Newton-like methods offer robust convergence behavior that does not depend on the selection and optimization of a damping parameter.

[1] C. T. Kelley, Iterative Methods for Linear and Nonlinear Equations, Vol. 16 of Frontiers in Applied Mathematics, SIAM, Philadelphia, PA, 1995.

[2] D. Knoll, D. Keyes, Jacobian-free Newton-Krylov methods: a survey of approaches and applications, Journal of Computational Physics 193 (2004) 357–397.

[3] E. E. Lewis, W. F. Miller, Jr., Computational Methods of Neutron Transport, American Nuclear Society, Inc., La Grange Park, Illinois, USA, 1993.

[4] X. Liu, X. Cheng, Thermal-hydraulic and neutron-physical characteristics of a new SCWR fuel assembly, Annals of Nuclear Energy 36 (2009) 28–36.

[5] L. Monti, T. Schulenberg, Coupled ERANOS/TRACE system for HP-WLR 3 pass core analyses, in: International Conference on Mathematics, Computational Methods & Reactor Physics, Saratoga Springs, NY, 2009.

[6] J. Yan, et al., Coupled computational fluid dynamics and MOC neutronic simulations of westinghouse PWR fuel assemblies with grid spacers, in: 14th International Topical Meeting on Nuclear Reactor Thermal Hydraulics, Toronto, ON, 2011.

[7] D. R. Vondy, T. B. Fowler, G. W. Cunningham, VENTURE: A code block for solving multigroup neutronics problems applying the finite-difference diffusion-theory approximation to neutron transport, version II, Tech. Rep. ORNL-5062/R1, Oak Ridge National Laboratory (1977).

[8] A. Dall'Osso, A neutron balance approach in critical parameter determination, Annals of Nuclear Energy 35 (9) (2008) 1686–1694.

[9] P. N. Brown, Y. Saad, Hybrid Krylov methods for nonlinear systems of equations, SIAM Journal on Scientific and Statistical Computing 11 (1990) 450–481.

[10] R. P. Pawlowski, J. N. Shadid, J. P. Simonis, H. F. Walker, Globalization techniques for Newton-Krylov methods and applications to the fully coupled solution of the Navier-Stokes equations, SIAM Review 48 (4) (2006) 700–721.

[11] J. Knoll, H. Park, C. Newman, Acceleration of $k$-eigenvalue/criticality calculations using the Jacobian-free Newton-Krylov method, Nuclear Science and Engineering 167 (2) (2011) 133–140.

[12] D. Gill, Y. Azmy, Newton's method for solving $k$-eigenvalue problems in neutron diffusion theory, Nuclear Science and Engineering 167 (2) (2011) 141–153.

[13] J. C. Ragusa, V. S. Mahadevan, Consistent and accurate schemes for coupled neutronics thermal-hydraulics reactor analysis, Nuclear Engineering and Design 239 (3) (2009) 566–579.

[14] R. S. Dembo, S. C. Eisenstat, T. Steihaug, Inexact Newton methods, SIAM Journal on Numerical Analysis 19 (2) (1982) 400–408.

[15] C. T. Kelley, Iterative Methods for Linear and Nonlinear Equations, SIAM, 1995.

[16] P. N. Brown, H. F. Walker, R. Wasyk, C. S. Woodward, On using approximate finite differences in matrix-free Newton-Krylov methods, SIAM Journal on Numerical Analysis 46 (4) (2008) 1892–1922.

[17] E. W. Larsen, J. E. Morel, J. M. McGhee, Asymptotic derivation of the multigroup P1 and simplified PN equations with anisotropic scattering, Nuclear Science and Engineering 123 (1996) 328–342.

[18] P. S. Brantley, E. W. Larsen, The simplified P3 approximation, Nuclear Science and Engineering 134 (2000) 1–21.

[19] E. W. Larsen, G. Thömmes, A. Klar, T. Götz, Simplified PN approximations to the equations of radiative heat transfer and applications, Journal of Computational Physics 183 (2) (2002) 652–675.

[20] T. Evans, A. Stafford, R. Slaybaugh, K. Clarno, DENOVO: A new three-dimensional parallel discrete ordinates code in SCALE, Nuclear Technology 171 (2010) 171–200.

[21] B. T. Rearden, et al., Modernization enhancements in SCALE 6.2, in: PHYSOR 2014 – The Role of Reactor Physics Toward a Sustainable Future, 2014.

[22] K. T. Clarno, et al., The AMP (Advanced MultiPhysics) nuclear fuel performance code, Nuclear Engineering and Design 252 (1) (2012) 108–120.

[23] J. E. Hansel, J. C. Ragusa, S. Allu, M. A. Berrill, K. T. Clarno, Analysis of physics-based preconditioning for single-phase subchannel equations, in: 2013 International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2013), 2013.

[24] A. M. Phillippe, K. T. Clarno, J. E. Banfield, L. J. Ott, B. Philip, M. A. Berrill, R. S. Sampath, S. Allu, S. P. Hamilton, A validation study of pin heat transfer for MOX fuel based on the IFA-597 experiments, Nuclear Science and Engineering 178.

[25] S. Hamilton, K. Clarno, M. Berrill, T. Evans, G. Davidson, R. Lefebvre, R. Sampath, Multiphysics simulations for LWR analysis, in: International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, Sun Valley, ID, 2013.

[26] S. Palmtag, Coupled single assembly solution with VERA (Problem 6), Tech. Rep. CASL-U-2013-0150-000, Consortium for Advanced Simulation of LWRs (2013).

[27] S. Hamilton, Numerical solution of the $k$-eigenvalue problem, Ph.D. thesis, Emory University, Atlanta, GA (2011).

[28] S. Hamilton, M. Benzi, A Davidson method for the $k$-eigenvalue problem, Transactions of the American Nuclear Society 105 (2011) 432–434.