炉物理の研究　第 73 号 (2021 年 3 月)

Experiences and Historical Perspectives on the Development and Application of Reactor Physics Methods for LWR Core Neutronic Analysis

Kord S. Smith

P.O Box 1082
Salmon, ID 83467
Email: kord@mit.edu

March 2021

Following my recent retirement from teaching nuclear engineering at MIT, Professor Akio Yamamoto asked me to write a short article for the Reactor Physics Division of the AESJ on my personal experiences and perspectives from my 45+ years of working on the development and application of nuclear reactor physics methods. It is my hope that this short note will provide useful perspectives to a target audience of graduate students and early researchers who seek to understand some of the history that led to the development of production LWR core neutronic analysis tools. It is my hope that these perspectives might prove useful to developers of future generations of nuclear reactor core analysis tools.

1) Computing Resources vs. Analysis Methods

Reactor analysis methods have been, and likely will continue to be, inexorably linked to the cost and capabilities of digital computers. As far back as 1945, the ENIAC computer was applied to calculations of early nuclear reactors (e.g., nuclear weapon cores), and early Monte Carlo particle simulation methods had been devised and tested. The primary technical challenges that stood between analysts and robust reactor analysis were the slow computational speed/small memories of early digital computers and the lack of experimental data for the energy dependence of neutron interaction probabilities (i.e., nuclear cross sections).

When I was first involved in reactor analysis in the mid 1970's there were two principle flavors of mainframe computers for tackling large reactor simulation problems: the IBM 360 or 370 and the CDC 6600 or 7600. IBM 370's had main memories of roughly 0.50 Mega-word (64 bit) and CDC 6600's had 0.1 Mega-word (60 bit) with roughly 2 MFLOPS of computational capability. Today, we have laptops with several Giga-words of memory and 10's of GFLOPS of computational capability. It is worth recognizing that 1970's reactor physics methods were constrained by **computers that had 10,000 times smaller memories and were 5,000 times slower** than modern laptops. It is also shocking today to realize how restricted and expensive computing was in the 1970's. During my graduate studies in the late 1970's, each student or researcher was allowed to submit only two computer jobs per day on MIT's IBM 370/168. Nevertheless, my EPRI-sponsored research budget for computing during the development of the QUANDRY nodal diffusion code at MIT was $25,000/year (~$107,000/year in 2021 dollars) – all for computing on machines probably less capable than a cheap cell phone of today!

My first full-core fast reactor Monte Carlo simulations on ANL's IBM 370/195 required overnight runs and could only track 100,000 total neutrons. One-sigma uncertainties in k-

effective were >0.20% delta-k, and tallies of local quantities were essentially useless. It was clear that the prospects of successfully using Monte Carlo methods for production reactor core analysis remained many decades in the future, and deterministic reactor analysis methods would dominate the 20th century landscape.

2) Nodal Diffusion

In the 1970's the PDQ finite-difference diffusion code from Bettis Atomic Power Laboratory (BAPL) was in widespread use for PWR analysis, and PDQ could handle 1-D, 2-D, and 3-D geometries in two or four energy groups, and its HARMONY model permitted nuclide depletion of reactor fuel materials. So how much memory was needed for modeling PWRs? In a standard PWR model (with homogenized pin-cell data, one mesh per pin-cell and 17x17pin fuel assemblies) a model has 17^2 mesh/assembly, $(15+2)^2$ assemblies/core, 2 groups/mesh that results in roughly 200,000 unknowns in 2-D. So just storing the 2-D pin-cell flux unknowns required half the memory of the IBM 370! Since one must also store the two-group cross section interaction matrix, spatial coupling terms, and old fission source, diffusion computations were effectively restricted to quarter-core 2-D models. With 3-D models needing several hundred axial meshes, it was impossible to accurately solve 3-D diffusion models of PWR cores.

Consequently, the task of the 1970's was to devise methods to accurately solve the 3-D 2-group LWR neutron diffusion equations using very coarse spatial meshes. Starting with the available CDC 6600 computer memory and assuming a **half-assembly radial mesh and 24 axial meshes**, one can deduce that the available storage per mesh in quarter-core geometry was $(0.100\text{x}10^6)/(24*17^2)=14$. It was known that first-order finite-difference methods with half-assembly (~10 cm) spatial mesh produced reactor core solutions with as much as 25% errors in assembly powers. Since target accuracies were 1-2%, higher-order methods would clearly be needed. So the daunting task of the 1970's was to devise more accurate coarse-mesh methods for accurately solving the 2-group neutron diffusion equation while storing less than 14 quantities per spatial mesh point.

Finite-element methods (that directly expand the three-dimensional flux distribution within each 3-D spatial element) were explored in the 1970's. Recognizing that assembly power errors for coarse mesh finite-difference (CMFD) could be 25% and that FEM methods have second-order spatial convergence dictated that at least 4th order FEM expansions would be needed to reduce errors to 1-2%. Researchers soon found that finite-element methods consumed far too much memory to store the FEM coupling matrices and solution runtimes were unacceptably large.

The crucial breakthrough of the 1970's was the development of the "transverse integration method" for solving the homogenized Cartesian-geometry neutron diffusion equations that took place in Germany. This new method started by recognizing that one could construct a set of formally exact diffusion equations by integrating the three-dimensional diffusion equations over two directions (e.g. transverse to the x direction) and representing the x-dependence of the y-plus-z direction leakages by a formally exact source term (i.e. the RHS or transverse leakage):

$$D_g \frac{d^2}{dx^2}\bar{\phi}_{gx}(x) - \Sigma_{ag}\bar{\phi}_{gx}(x) + \frac{1}{k_{eff}}\chi_g \sum_{g'=1}^{G} \nu\Sigma_{fg'}\bar{\phi}_{g'x}(x) + \sum_{g'=1}^{G} \Sigma_{sg'g}\bar{\phi}_{g'x}(x) = \left[L_{gy}(x) + L_{gz}(x) \right].$$

If the transverse leakage shape as a function of x were known, the 1-D neutron diffusion equation could be solved for the x-direction flux shapes and the exact "coupling coefficients" between two adjacent nodes in the x direction. Similar equations could be solved for the y- and z-direction coupling equations, and the transverse leakage method reduced a 3-D diffusion problem into a set of three 1-D diffusion problems – one for each for each Cartesian mesh direction. The truly astonishing thing about this method was that intra-nodal transverse leakage shapes could be approximated by quadratic polynomials (constrained by fitting three neighboring node-averaged transverse leakages for a given direction), and the target accuracy of ~1% in assembly powers was obtained.

The initial method (the nodal expansion method, NEM) for solving the transverse-integrated diffusion equations used 4th- or 5th-order polynomials to represent the 1-D transverse-integrated intra-nodal flux shapes in each direction. Compared to finite-element methods, the transverse integration methods for 5th order flux representations required roughly a factor of 12 (e.g., 6^3/6*3) less unknown expansions coefficients, and the transverse integration method was far more computationally efficient than finite-element methods. I was in a group at MIT in 1976-1978 working on methods for solving the 3-D diffusion equation, and we developed the analytic nodal method (ANM) that did not use polynomial flux expansions to represent the intra-nodal transverse-integrated fluxes, but rather employed direct analytic solutions to the 1-D 2-group diffusion equations with quadratic source terms. The ANM method was later extended to four groups by researchers in South Africa and ultimately to general multi-group by Sten-Oran Lindahl of Studsvik. While mathematically more appealing, ANM was only slightly more accurate than NEM for standard UO_2-fuled cores because the primary source of inaccuracy in both methods is the quadratic transverse leakage shape. Later, the Semi-Analytic Nodal Method (SANM) was developed to use analytic solutions for the thermal group diffusion equation and polynomials in the fast group. SANM had the advantage of retaining the accuracy of ANM for steep flux gradients that occur at UO_2/MOX assembly interfaces, but eliminated much of the algebraic complexity of the group-coupled ANM. I also spent considerable time exploring alternative transverse leakage shape approximations with proper discontinuity between nodes and higher-order spatial representations, but I was never able to significantly improve on the simple quadratic fit developed for the original NEM formulation. When these transverse-integrated nodal methods had been developed and implemented, solutions times for full-core 3-D 2-group homogenized nodal diffusion core problems were approximately 30 seconds on the IBM 370 and CDC 7600.

The second breakthrough for nodal methods was one I made shortly after leaving graduate school. In a short ANS Transactions paper, I summarized a very simple method for reformulating any linear transverse-integrated nodal method into a nonlinear form that nearly resembled CMFD. This technique is called Nonlinear Diffusion Acceleration (NDA), as the method converts the complex nodal coupling equations for the interface net current into a simple CMDF-like form where the interface-averaged net currents are related to the node-averaged flux of two neighboring nodes by

$$J_g^{i+1/2} = -\frac{\vec{D}_g}{\Delta x}\left(\bar{\phi}_g^{i+1}\right) + \frac{\overleftarrow{D}_g}{\Delta x}\left(\bar{\phi}_g^{i}\right),$$

where the two diffusion-coefficient-like terms are computed analytically using the formal definition of discontinuity factors (discussed in Section 4 of this note) for any selected transverse-integrated nodal method. In later years it became common to cast the net current relation in terms of the conventional CMFD geometrical-mean diffusion coefficients for the two neighboring nodes (D_g) and D-tilde terms that introduce unknown nonlinear coefficients that multiplies the sum of the two neighboring node-averaged group fluxes.

$$J_g^{i+1/2} = -\frac{D_g}{\Delta x}\left(\bar{\phi}_g^{i+1} - \bar{\phi}_g^{i}\right) + \frac{\tilde{D}_g}{\Delta x}\left(\bar{\phi}_g^{i+1} + \bar{\phi}_g^{i}\right)$$

This formulation, while slightly different from the previous equation, has similar performance but its relationship to coarse mesh finite-difference is easier to understand. If a CMFD model is accurate, the value of D-tilde is identically zero. When CMFD is inaccurate, D-tilde provides a correction term. If one worries about the non-linear D-tilde term being too large, one can use the previous equation and an ANM nodal model, so that none of the 1-D "spatial truncation error" of CMFD is included in the nonlinear term. In more recent years, Y. A. Chao has termed this representation as ACMFD (Analytic CMFD), and it is identical to that originally used in my 1979 M.S. thesis.

The global nodal solution procedure (often referred to as CMFD acceleration) is simple:

1. One solves the conventional CMFD nodal balance equations (an eigenvalue problem) to partially converge a global core flux solution

2. Using the partially converged node-averaged fluxes, one solves the nodal equations (using NEM, ANM, or SANM, etc.) for the net current at each nodal interface

3. Assuming that the net current and nodal fluxes are converged, one backs out the value of D-tilde that preserves the nodal method interface net current using the proceeding equation and the partially converged fluxes

4. One iteratively solves the new CMFD-like nodal balance equations with the updated values of D-tilde to partially converge an updated global core flux solution

5. One returns to step 2 until the global core flux solution is completely converged

If this non-linear iteration converges, it is guaranteed to converge to the same solution that would be obtained by directly solving the linear nodal equations. The beauty of NDA is that all iterative equations are identical in form as conventional CMFD equations (seven-stripe block diagonal matrix in 3-D Cartesian geometry). The total storage required is only 13 quantities per node (in two groups) where direct ANM requires storage of 128 quantities per node. This trick of converting linear equations into an equivalent set of nonlinear equations not only significantly reduced the computing memory required, but it also eliminated many of the mathematical operations needed to iteratively invert the flux matrix at each fission source iteration. The development of NDA was crucial for enabling solutions of large 3-D diffusion problems on

1970's era computers, and it reduced both computer storage and computational effort. The memory reduction of NDA was later to prove critical to nodal method adaptation to engineering workstations and IBM PCs (with only 640 kilobytes of memory) in the early 1980's. It is worth noting that even though one could run nodal core models on early PCs, execution times to compute a full-core core power distribution were on the order of 24 hours – thousands of times slower than on the mainframe computers of the day.

While NDA has found a home in most production advanced nodal diffusion codes, it was not universally accepted when first introduced. Mathematicians naturally objected to converting linear equations into nonlinear equations, without rigorous proof of convergence. Engineers, like myself, were not so easily offended by the nonlinearity because production core analysis tools involve nonlinear iterative coupling to thermal-hydraulic models for Doppler and coolant density feedback – which are far more nonlinear effects than the D-tilde nonlinearity.

Once the transverse leakage shape has been incorporated as a source in the 1-D nodal equations, it becomes straightforward to add additional source terms to the RHS to allow for intra-nodal spatial distributions of cross sections to treat such things as: fuel nuclide density variation from depletion, fuel temperatures, coolant densities, and xenon number densities. All of these terms are important when one desires to maintain 1-2% accuracy on assembly powers in the presence of real-world feedback and fuel depletion effects. It should be noted that the very nature of the transverse integration procedure limits the spatial representation of feedback terms to essentially products of one-dimensional polynomials, where it is possible (but complicated) in a FEM method to represent directly the non-separable nature of the feedback shape – when those terms are known from other physics.

3) Lattice Physics

Nodal methods start from the assumption that fuel assembly cross sections have been homogenized into appropriate assembly-averaged cross sections and there is no spatial variation of cross section within each node (e.g., element in FEM parlance). Early PWR analysis methods use 1-D cylindrical fuel pin/clad/coolant neutron transport calculations to compute spatial flux distributions in ~100 energy groups, and these fluxes were collapsed in both space and energy to obtain 2-group pin-cell-averaged cross sections that could be used in Cartesian-geometry PDQ core diffusion models. Assembly-averaged cross sections for PWRs were sometimes obtained from similar 1-D calculations with additional water added to treat the guide tube volumes within a fuel assembly. Other times, assembly diffusion calculations were performed with few-group pin-cell-averaged cross section data, and assembly-averaged cross sections were obtained by preserving reaction rate integrals over the fuel assembly. The later two-step approach was preferred when burnable absorber pins were introduced into PWR assemblies, or when multiple enrichment fuel pins were present – as in BWR assemblies.

Nodal methods benefited tremendously from the late 1970's development of lattice transport methods. Rather than relying on the pin-cell fine-energy transport and coarse-energy **assembly diffusion** factorization approach, researchers focused on development of pin-cell fine-energy transport and coarse-energy **pin-cell-homogenized assembly transport** methods. This led to the development of codes such as WIMS, PHOENIX, CASMO, DIT, APOLLO-2, and

PARAGON. While there are many differences in the approaches taken in these codes, they each provided fully integrated resonance self-shielding (to compute multi-group cross sections), fine-group pin-cell transport, pin-cell-homogenized coarse-group assembly transport, and integrated nuclide depletion – all in an integrated code. These codes produce assembly-averaged few-group cross sections as a function of assembly nuclide depletion, coolant density, fuel temperature, etc., by modeling each fuel assembly type with reflective boundary conditions. From the nodal code perspective, pre-built "libraries" of assembly-averaged cross sections for each unique assembly type are produced as a function of all conditions encountered within a reactor core – so assembly-homogenized cross section data is assumed to be available to downstream nodal codes.

In the 1990's and 2000's lattice physics methods were extended to directly solve the fine-energy **heterogeneous assembly transport** problem in codes such as CASMO-4, HELIOS, APOLLO-3. These advancements focused almost exclusively on the Method of Characteristics (MOC) transport approach rather than the widely available discrete ordinates (S_N) transport methods of the national laboratories. MOC has a distinct advantage over S_N because it explicitly represents curvilinear geometries and its' transport angular flux sweep is independent of all other angles' angular fluxes – leading to a rather substantial advantage in computational requirements for the finely discretized spatial regions of LWR lattices.

The S_N community had extensively developed diffusion synthetic acceleration methods to reduce the computational burden of converging transport solutions, but these methods proved to be far inferior to simply superimposing a pin-cell Cartesian acceleration mesh on top of the lattice geometry and using NDA to formulate a consistent CMFD-like multigroup diffusion problem. The CMFD-like equations are used to solve the eigenvalue problem, and the resulting pin-cell averaged fluxes are used to prolong the detailed intra-pin-cell fluxes (and sources) before starting the next transport sweep. This application of NDA is very similar to the multi-grid acceleration methods developed in the 1980's – but it has an added advantage of precisely preserving the fluxes on the acceleration mesh. NDA, like multi-grid, is easily applied in several stacked levels. In the early 2000's, NDA-accelerated MOC transport was extended to full-core 2-D problems with one level of acceleration for the pin-cell mesh and a second level for the assembly mesh, and the resulting Cartesian diffusion problems are readily solved by a variety of iterative methods. NDA-accelerated MOC LWR core problems are typically converged using only 10-20 angular flux sweeps – rather than 1000's typically needed for direct transport sweep solutions of high dominance ratio (e.g., >0.99) LWR cores.

Multi-assembly or full-core heterogeneous transport solutions are valuable tools for assessing the accuracy of nodal methods when "homogenization errors" are investigated, as explained in Section 4 of this note. Because of computing limitations of the 1970's, homogenization errors were first assessed using 2-D multi-assembly or full-core heterogeneous pin-cell-homogenized diffusion solutions (e.g. with PDQ). The ultimate determination of nodal method accuracy was made indirectly by comparison with measured PWR and BWR fission detector spatial distributions from operating power plants. The accuracy of such predictions was critically dependent on the quality of the energy-dependence of basic evaluated nuclear cross section data. One should remember that even if one could solve the full-core transport problem exactly, the solution accuracy depends on the quality of the basic nuclear data employed. As US industry moved from ENDF/B-II in the early 1970's to ENDF/B-VIII in 2018, the accuracy of core

analysis was dramatically improved – regardless of core analysis method (e.g., nodal, full-core heterogeneous transport, or Monte Carlo) employed.

The main point here is that reactor analysis methods are inherently dependent on lattice physics approximations and/or the quality of nuclear data/models employed.

4) Assembly Homogenization/Pin Power Reconstruction

The transverse-integrated nodal methods described in Section 2 were proven to be capable of solving homogenized diffusion problems using assembly-sized (BWRs) or half-assembly-sized meshes (PWRs) with 1-2% accuracy in assembly power relative to reference solutions for numerical diffusion benchmark problems. As nodal mesh sizes are decreased, the nodal solutions are guaranteed to converge to the reference diffusion solutions. However, representing heterogeneous fuel assemblies using assembly-averaged cross section introduces approximations that are not eliminated by refining the nodal mesh. Consequently, methods development in the late 1970's and early 1980's focused on understanding and solving the "homogenization problem".

One can understand the homogenization problem by assuming the existence of a reference fine-mesh whole-core transport solution (e.g., from Monte Carlo or fine-energy transport calculation) and then ask what nodal quantities should be extracted from the reference solution so that a nodal diffusion solution can precisely preserve all node-integrated quantities of the reference solution. Researchers traditionally assumed that the best possible diffusion solution would be obtained by defining node-averaged cross sections from the reference node-integrated reaction rates and fluxes. However, tests indicated that this procedure led to errors in assembly powers as large as 4-5% in PWRs and 10-15% in BWRs – far larger than target accuracies of 1-2%.

Some homogenization methods were developed by finding cross sections that preserved the response matrix of each assembly, but little real progress was made until Klaus Koebke (of KWU in Germany) made a game-changing observation. Koebke proved that the flaw in using the conventional reference node-averaged cross sections was not with the cross sections themselves but rather with the interface continuity conditions imposed on the nodal diffusion solution. Koebke demonstrated that a formally exact nodal solution could be obtained using node-averaged cross sections if the interface scalar fluxes were allowed to be discontinuous and the magnitude of the discontinuity was constructed from the reference solution. Koebke reasoned that the transverse-integrated nodal flux distribution for a given node and direction is uniquely determined by 1) the node-average cross sections, 2) the node surfaced-averaged net current on the two opposite nodal faces, and 3) the intra-nodal transverse leakage shape. Cross sections and interface net currents can easily be edited from reference solutions. For any chosen transverse leakage model (that uses node-averaged transverse leakages as constraints) the intra-nodal leakage shape can be uniquely constructed from the reference solutions. Consequently, Koebke proposed to simply compute the intra-nodal transverse-integrated flux distribution using the two surface-integrated net currents as constraints on the diffusion solution. Koebke used the NEM nodal model to perform such diffusion calculations, and he defined the resulting intra-nodal flux shape as the "homogenized flux" as

$$\phi^i_{Hom}(x).$$

Koebke defined new homogenization parameters (called Heterogeneity Factors) for each direction and interface as the ratio of reference (heterogeneous) flux to the homogenized flux at the + and – interfaces of node *i*

$$HF^i(x^{\pm}) \equiv \frac{\phi^i_{Het}(x^{\pm})}{\phi^i_{Hom}(x^{\pm})}.$$

Since the reference solution must obey continuity of heterogeneous scalar flux between node *i* and node *i+1*,

$$\phi^i_{Het}(x^+) = \phi^{i+1}_{Het}(x^-) \Rightarrow \phi^i_{Hom}(x^+) HF^i_{Hom}(x^+) = \phi^{i+1}_{Hom}(x^-) HF^{i+1}_{Hom}(x^-).$$

This expression implies that

$$\frac{\phi^i_{Hom}(x^+)}{\phi^{i+1}_{Hom}(x^-)} = \frac{HF^{i+1}_{Hom}(x^-)}{HF^i_{Hom}(x^+)}.$$

In other words, the **magnitude of the discontinuity** of homogenized (and transverse-integrated) nodal fluxes at a nodal interface can be directly computed from the reference solution interface-averaged scalar fluxes. Koebke correctly reasoned that **homogenized diffusion scalar fluxes must be discontinuous at nodal interfaces** in order to preserve interface-averaged net currents of the reference solution. This observation is rather stunning, since we know from first principles that reference scalar fluxes are continuous at all points in space. Koebke's observation was key to understanding how to define rigorous homogenization parameters that precisely preserve reference nodal net currents. Moreover, when interface net currents are preserved (and node-integrated cross sections are used) nodal fluxes will identically match the reference node-integrated scalar fluxes, and the reference core eigenvalue is preserved. One consequence of using Koebke's Heterogeneity Factors as derived, was that any energy-group-collapse or transport effects are naturally accounted for by the HFs. Moreover, exact preservation of reference solutions was obtained – independent of the number of energy groups employed, the choice of nodal model order (including CMFD), or the choice of transverse leakage shape representation. Thus, the HFs account for more than just the heterogeneity of the lattices.

Since the intra-nodal flux shape is a function of the nodal diffusion coefficient, the Heterogeneity Factors are naturally dependent on the choice of diffusion coefficients. Koebke went one step further to select arbitrary values of PWR assembly diffusion coefficients for each direction so that

$$HF^i_{Hom}(x^-) = HF^i_{Hom}(x^+).$$

At MIT, Prof. Alan Henry and I were studying homogenization methods at the time (1978) that Koebke devised the Heterogeneity Factors for homogenization, and we appreciated the wisdom of his bold approach. We performed many studies of Heterogeneity Factors for both PWRs and

BWRs, and we demonstrated that a better approach was to use the most physical diffusion coefficient and treat each interface with its own unique homogenization factor. These factors are now referred to as Discontinuity Factors (DFs) to distinguish them from those computed with Koebke's arbitrary direction-dependent diffusion coefficients

$$DF^i(x^{\pm}) \equiv \frac{\phi^i_{Het}(x^{\pm})}{\phi^i_{Hom}(x^{\pm})}.$$

Discontinuity Factors have significant advantages over Heterogeneity Factors in BWR applications because of the large flux asymmetries introduced by 1) wide and narrow water gaps of the assemblies , 2) asymmetric distributions of gadolinium absorber pins, and 3) asymmetric distributions of pin enrichments.

The most significant contribution of my Ph.D. thesis was to demonstrate that reference Discontinuity Factors could be accurately approximated by performing lattice calculations with reflective boundary conditions. Such lattice calculations have spatial flat homogenized flux distributions and Assembly Discontinuity Factors (ADFs) can be uniquely defined for any single assembly (SA) lattice as

$$ADF(x^{\pm}) \equiv \frac{\phi^{SA}_{Het}(x^{\pm})}{\phi^{SA}_{Hom}(x^{\pm})} = \frac{\phi^{SA}_{Het}(x^{\pm})}{\overline{\phi}^{SA}_{Het}}$$

ADFs are nothing more than edited ratios of group-wise surface-averaged fluxes to volume-averaged fluxes. Thus, these new homogenization parameters can be obtained directly from the same lattice solutions used to compute assembly-averaged cross sections – with essentially zero computational cost. It is important to note that if one wants DFs for nodes that have net neutron leakage (like full-core problems or fuel assembly/reflector problems), then the formal step of reconstructing nodal homogenized flux shapes must be performed for each node (and direction) and DFs are defined by the ratio of heterogeneous and homogenized surface-averaged fluxes, as initially defined in this note. This later point is particularly important when defining PWR baffle/reflector homogenized nodal data that preserves core neutron leakage effects.

The power of ADFs lies in the fact that for PWR and BWR applications, the maximum and mean errors in assembly powers are reduced by approximately a factor of 3.0 to 4.0 relative to those observed when global flux distributions are computed assuming continuity of nodal interface scalar fluxes. This makes ADFs a very practical method for eliminating most of the homogenization errors normally encountered in LWR homogenized diffusion calculations. Moreover, ADFs are nearly trivial to introduce into transverse-integrated nodal methods, and there is essentially no computational overhead. One should recognize that ADFs are equally applicable to FEM (or any other) diffusion methods. However, it is exceedingly difficult to constrain FEM multi-dimensional flux expansions to have known scalar flux discontinuities across each nodal (element) interface. ADFs and transverse-integrated nodal methods fit together in a very natural way.

ADFs can also be derived intuitively by simply assuming a factorization of the global heterogeneous flux distribution into a homogeneous (smooth) flux shape and a heterogeneous

(lattice) flux shape within each homogenized node. With this idea in mind, it is easy to understand how one can take a nodal flux solution (with its implicit intra-nodal flux shape) and modulate it with the heterogeneous (lattice) flux shape to approximate the global heterogeneous fluxes. This is often referred to as "pin power reconstruction" or de-homogenization. Note that there are many 3-D intra-nodal flux shapes that can satisfy the 1-D transverse-integrated flux shapes in each of the three directions, and additional constraints are often imposed to obtain non-separable components of the flux shape. These constraints often utilize nodal corner-point continuity constraints and lattice-computed Corner Point Ratios (CPRs) defined in x-y geometry at each lattice corner as

$$CPR(x^{\pm}, y^{\pm}) \equiv \frac{\phi_{Het}^{SA}(x^{\pm}, y^{\pm})}{\phi_{Hom}^{SA}(x^{\pm}, y^{\pm})} = \frac{\phi_{Het}^{SA}(x^{\pm}, y^{\pm})}{\overline{\phi}_{Het}^{SA}}$$

Note the similarity in definition of the CPRs to the ADFs. It is also easy to understand that since pin power reconstruction is performed after convergence of the nodal solution, these methods produce pin-wise power distributions at very low computational cost. Many other subtle details needed for accurately performing LWR pin power reconstruction are beyond what can be summarized in this short note, but suffice it to say that nodal methods with pin power reconstruction form the backbone of current production LWR analysis. The primary complexity of pin power reconstruction is in treating lattice flux (or power) distributions as additional de-homogenization parameters – because the 100's of pin-wise powers for a lattice significantly increase the size of the "homogenization parameter library" compared to the seven two-group cross sections and eight (2-D) group-wise discontinuity factors. Execution times for computing full-core 3-D pin-wise power distributions (including thermal-hydraulic feedback) with these nodal methods are typically a few seconds and cycle depletion calculations take roughly one minute on today's single-CPU laptops.

5) Observations and Reflections

Being intimately involved in the development of advanced nodal methods for LWR analysis has been a tremendously rewarding personal endeavor. While I initially thought it would only take two or three years to produce production quality nodal codes, I soon discovered countless additional refinements that were needed to produce functional and accurate codes (such as the proper definition of diffusion coefficients, homogenization models for baffles/reflectors, techniques for treating nonlinear feedback, modeling axial control rod cusping effects, and finding accurate/efficient time-integration methods). I also found it necessary to understand and improve fuel pin performance and thermal-hydraulics models that are needed for rapid computation of feedback parameters in transient nodal calculations. It also became my goal to improve lattice transport methods to efficiently handle full-core 2-D lattice calculations. By using these tools for full-core tests, it was possible to understand the sources of error in homogenization/cross section models and to further improve nodal method approximations. I have also learned a tremendous amount from working with utility customers about the engineering features that are needed to make codes easy to use (e.g., automatic coefficient calculations, HFP cycle length searches, power/flow searches to critical, automated stuck-rod minimum shutdown margin calculations, axial-offset constrained boron/control rod bank

insertion sequences, etc.). It would have been overwhelming to realize that it would take more than thirty years to complete what I thought were my M.S. and Ph.D. thesis objectives!

In the last decade of teaching and researching advanced nodal methods at MIT, I have found it very challenging for new researchers to understand and appreciate the level of detailed modeling needed before one can improve the accuracy of production nodal models. It is far easier and more intuitive in academia to teach new heterogeneous transport methods - where feedback is brute-forced onto a very fine spatial neutronic mesh, and accuracy is limited primarily by the available computing resources. Students of nuclear engineering have always been, and will undoubtedly continue to be, attracted to using the latest and most powerful computing resources to solve nuclear core physics problems. We are now at the point of having massively parallel Exa-FLOP computers (10^{12} times the FLOPS of 1970's computers) with many Peta-bytes of aggregate memory at the US national laboratories. Consequently, computer memory seems infinite and the largest of problems can be tackled by throwing more and more cores at nuclear simulations. However, one should remember that modern multi-level cache computers (unlike the totally-fast-memory computers of the 1970's) are limited more by the time required to move data between compute registers and main memory (via multi-level caches) than by available FLOPS – thus one often hears "FLOPS are free". Consequently, **simulation runtimes are often dictated by the amount of node memory addressed**, and using unnecessary memory inevitably makes computations much slower than they need to be. I constantly remind students that near-perfect scalability of methods on massively parallel machines is no substitute for efficiency in reactor simulation models. Students discover over and over that many lessons learned from early nodal method development regarding memory minimization and cache coherency are just as important for modern supercomputer as they were for 1970's vintage computers.

6) Perspectives on Current Reactor Physics Methods Developments

Nodal methods have also been used as a steppingstone to the development of 3-D transport methods. After efficient 2-D transport methods (with NDA) had been developed, researchers realized that pseudo 3-D transport calculations could be performed by running 2-D MOC transport models for each unique plane of a reactor core, and then stitching the axial planes together using nodal models for each radially homogenized pin-cell (using DFs in the radial direction to approximately preserve the MOC solution). Such 3-D calculations can be executed in roughly the time it takes to compute all of the 2-D MOC calculations – which for a 50-group full-core LWR core calculation is typically less than 1000 core-hours on modern multi-core computers.

In the very recent years, NDA-accelerated 3-D MOC transport and billion-neutron Monte Carlo solutions have now been obtained for 3-D LWR cores. Both approaches require 10,000 to 100,000 core-hours to solve single reactor states, and application of such models remains out of reach for today's commercial reactor analysts. Successful development of applications for GPU-like architectures may ultimately bring heterogeneous 3-D reactor transport to the point at which it becomes accessible to the average reactor analyst. However, it is important for developers of high-fidelity reactor analysis method to come to grips with **storing the nuclide compositions for a full-core 3-D LWR core models – which requires nearly a Terabyte of memory** (i.e., 200 assemblies x 264 pins/assembly x 10 rings/pin x 200 axial mesh per pin x 400 nuclides x 8

bytes/variable = 0.34 x 10^{12} bytes). One should not make the mistake of developing methods that work well for problems having a small number of unique fuel compositions – because every fuel mesh in a reactor model ultimately develops a unique nuclide composition.

One should also not forget that 3-D core analysis methods must be extended to the transient case for reactor safety analysis, and this adds another factor of 100 or 1000 to the computational requirements. While transient extensions are mathematically straightforward, many computational challenges remain for nuclear methods developers of both fine-mesh deterministic and Monte Carlo transport. Multi-physics feedback effects drive reactor transient analysis, and methods developers must address the non-linear nature of feedback between physics fields. Nevertheless, it is my expectation that new and powerful techniques for improving the computational efficiency of high-fidelity transient methods will be successfully developed, and such methods/tools will someday replace existing transient nodal methods.

Reactor vendors and utilities have always been a decade behind national laboratories in terms of available computing capabilities – because the costs of computing hardware and human resources remain a major consideration. Thus, it is very likely that the 45-year-old advanced nodal methods will continue to be used for another decade or more – simply because high-fidelity transport methods are tens of thousands to a million times less computationally efficient at achieving the desired solution accuracy across the full spectrum of existing nodal method applications (e.g., core design and cycle depletions, core loading optimization, operational support, transient safely analysis, real-time core monitoring, and real-time operator training).

A Very Few Selected References

Nodal methods

- K. Smith, "An Analytic Nodal Method For Solving the Two-Group Multidimensional Static and Transient Neutron Diffusion Equation," N.E. and M.S. Thesis, MIT, Cambridge, MA (1979).
- M. Wagner, K. Koebke, and H-J Winter, "A nonlinear extension of the nodal expansion method," *International Topical Meeting on Advances in Mathematical Methods for the Solution of Nuclear Engineering Problems*, Vol 2, p. 43, Munich, 27-29 April (1981).
- K. Smith, et al., "QPANDA: An Advanced Nodal Method for LWR Analysis," *Trans. Am. Nucl. Soc.*, Vol. 50, p. 532, San Francisco, CA (1985).
- R. Lawrence, "Progress in Nodal Methods for the Solution of the Neutron Diffusion and Transport Equations," *Progress in Nuclear Energy*, **17**, pp. 271-303 (1986).

Nonlinear diffusion acceleration (NDA)

- K. Smith, "Nodal Method Storage Reduction by non-linear Iteration," *Trans. Am. Nuc. Soc.,* Vol. 44, p. 265 (1983).

Homogenization and pin power reconstruction

- K. Koebke. "A New Approach to Homogenization and Group Condensation," Lugano, IAEA-TECDOC-231 (1978).
- K. Smith, "Spatial Homogenization Methods for Light Water Reactors," Ph.D. Thesis, MIT, Cambridge, MA (1980).
- K. Koebke. "Advances in Homogenization and Dehomogenization," *Int. Topical Meeting on Advances in Math. Methods for the Solution of Nucl. Eng. Problems,* Munich (1981).
- K. S. Smith. "Assembly Homogenization Techniques for Light Water Reactor Analysis," *Progress in Nuclear Energy*, Vol. 17, No. 3, pp. 303-335 (1986).

NDA for acceleration of LWR transport calculations

- K. S. Smith and J. D. Rhodes III, "Full-Core, 2-D, LWR Core Calculations with CASMO-4E," *Proc. Int. Conf. New Frontiers of Nuclear Technology: Reactor Physics, Safety and High-Performance Computing* (PHYSOR 2002), Seoul, Korea, October 7–10 (2002).
- G. Gunow, B.Forget, and K. Smith. "Full core 3-D simulation of the BEAVRS benchmark with OpenMOC", *Annals of Nuclear Energy* 134 (2019).
- G. Guidicelli, "A Novel Equivalence Method for High Fidelity Hybrid Stochastic-Deterministic Neutron Transport Simulations," Ph.D. Thesis, MIT, Cambridge, MA (2020).