Annals of Nuclear Energy 63 (2014) 22–37

Contents lists available at ScienceDirect

Annals of Nuclear Energy

journal homepage: www.elsevier.com/locate/anucene

A multi-level surface rebalancing approach for efficient convergence acceleration of 3D full core multi-group fine grid nodal diffusion iterations

René van Geemert *

AREVA GmbH, Paul-Gossen-Strasse 100, 91052 Erlangen, Germany

ARTICLE INFO

Article history:
Received 20 March 2013
Received in revised form 10 July 2013
Accepted 11 July 2013
Available online 9 August 2013

Keywords:
Nodal diffusion
Iterative solution method
Multi-grid
Rebalancing
Eigensystem

ABSTRACT

A new multi-level surface rebalancing (MLSR) approach has been developed, aimed at enabling an improved non-linear acceleration of nodal flux iteration convergence in 3D steady-state and transient reactor simulation. This development is meant specifically for anticipating computational needs for solving envisaged multi-group diffusion-like SP_N calculations with enhanced mesh resolution (i.e. 3D multi-box up to 3D pin-by-pin grid). For the latter grid refinement regime, the previously available multi-level coarse mesh rebalancing (MLCMR) strategy has been observed to become increasingly inefficient with increasing 3D mesh resolution. Furthermore, for very fine 3D grids that feature a very fine axial mesh as well, non-convergence phenomena have been observed to emerge. In the verifications pursued up to now, these problems have been resolved by the new approach. The novelty arises from taking the interface current balance equations defined over all Cartesian box edges, instead of the nodal volume-integrated process-rate balance equation, as an appropriate restriction basis for setting up multi-level acceleration of fine grid interface current iterations. The new restriction strategy calls for the use of a newly derived set of adjoint spectral equations that are needed for computing a limited set of spectral response vectors per node. This enables a straightforward determination of group-condensed interface current spectral coupling operators that are of crucial relevance in the new rebalancing setup. Another novelty in the approach is a new variational method for computing the neutronic eigenvalue. Within this context, the latter is treated as a control parameter for driving another, newly defined and numerically more fundamental eigenvalue to unity. This paper presents a rigorous derivation of the new approach, followed by a comparison on convergence efficiencies, for a number of 3D full core nodal grid resolution regimes, between the previously available multi-level rebalancing setup and the new multi-level surface rebalancing concept. The surface rebalancing methodology and a number of related concepts are covered in the patents EP2091049, EP2287855, EP2287854 and EP2287853 that were granted in 2012.

1. Introduction

In a number of AREVA's 3D reactor simulation tools (such as PANBOX (Böer et al., 1992)), PRISM (Grummer, 2000) and ARTEMIS™ (Hobson, 2008), a default use of multi-level rebalancing (van Geemert, 2006) enables a substantial run time reduction for industrial whole-core fuel cycle and transient simulations. Once implemented, the efficiency of a multi-level coarse mesh rebalancing (MLCMR) setup can be quantified in terms of the enabled final run time reduction compared to unaccelerated scenarios. This success is determined by the effected flux error decay rate and the cost of a single nodal multi-group SP_N iteration sweep on the one hand, versus the computational cost of a single multi-level rebalancing cycle. Anticipating the trend to pursue multi-group (i.e. 8- up to 16-group) SP_N simplified transport computations for industrial applications on very fine spatial grids, a new multi-level surface rebalancing (MLSR) concept has been developed. Recent investigations have shown that, on top of the benefit of doing rebalancing at the level of the nodal interfaces instead of merely the nodal volumes, it pays off to extend the rebalancing grid hierarchy with a number of additional space-energy levels corresponding to coarsened (i.e. numerically lumped) neutron energy grids. Previously, a restriction from the fine multi-group space-energy mesh to the finest rebalancing mesh meant lumping *all* neutron energy group spectral information per node into a *single, lumped group*. This is a simple and effective concept for accelerating few-group coarse nodal diffusion computations. However, for accelerating substantially more expensive multi-group fine grid computations (i.e. with 8 or more groups) this approach can be improved. Specifically, one

* Tel.: +49 9131 90095537.
E-mail address: rene.vangeemert@areva.com

0306-4549/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.anucene.2013.07.012

can do this by not only doing restrictions over nodal interfaces but by also defining intermediate coarsened space-energy rebalancing levels featuring a specific energy group hierarchy defined by certain lumping strategies. From this, a generically enhanced numerical proximity results between different rebalancing space-energy grids and the 3D multi-group nodal diffusion iterations. Combined with carefully pursued algebraic manipulations for making sure that the rebalancing operators are optimally sparse, and the use of efficient and lean iteration approaches at rebalancing level, a significant upgrade of error decay rate proved achievable.

All rebalancing concepts presented here do boil down to forming space-energy-coarsened blocks of different terms. These may pertain to the process-rate balance equations in MLCMR or the interface current balance equations that are the restriction basis of the new MLSR approach that is the main topic of the current paper. Subsequently, dimensionless multiplicative correction factors (''driving factors'') are solved repeatedly that are associated with nodal volume and/or interface and which are applied for non-linear acceleration of the iterative solution process. Upon overall convergence, all driving factors converge to 1. This concept is somewhat different from the also well-known coarse mesh finite difference (CMFD) methodology (Lewis, 2005; Yoon and Joo, 2008; Tatsumi and Yamamoto, 2003). The latter does have the same non-linear acceleration intention and is characterized by a systematic recomputation of coarse mesh fluxes at the projected coarser mesh level. From there, correction factors are derived as ratios between CMFD-iterated coarse mesh fluxes vs. previously projected coarse mesh fluxes that followed from the integration of previous fine-mesh fluxes. Within the CMFD framework, the projection of the nodal equations to a finite-difference-like coarse mesh system for merely (coarse mesh) nodal fluxes features the systematic accompanying recomputation of coupling coefficients that enable the formulations of the successive (and converging) CMFD systems. These coupling coefficients follow from solutions of 2-node problems based on Fick's law. Both the MLCMR and the MLSR approaches presented in this paper are direct multi-grid projection methods that do not require the computation and subsequent use of additional coupling coefficients in the overall multi-level iteration setup.

This document describes the developed MLSR strategy and presents reports on numerical verification and observed impacts on overall computational efficiency. It was observed that, in case of using classical coarse mesh rebalancing (van Geemert, 2006), the average flux error decay gets substantially poorer when going to higher grid resolution. Furthermore, for very fine 3D grids that feature a very fine axial mesh as well, non-convergence phenomena have been observed to emerge. In all verifications pursued thus far, these problems have been resolved completely by the MLSR approach.

2. A brief history of rebalancing applications for accelerated solution of 3D nodal diffusion equations

Coarse mesh rebalancing, which was proposed originally as a ''variational acceleration technique'' for basic diffusion iterations by Wachspress (1966) in the 1960s, provides a suppression of otherwise slowly decaying low-frequency non-fundamental mode components in the not-yet-converged diffusion solution. This is realized through a systematic multiplicative correction of nodal fluxes and interface outcurrents, prior to each full core nodal diffusion sweep, with iteratively obtained ratios between coarse mesh rebalanced and prebalanced fluxes (Nakamura, 1977). These ratios, conventionally referred to as *driving factors*, emerge from restarted preconditioned Krylov subspace procedures defined at the computationally substantially cheaper coarse mesh equation level. Upon their application, they push down the long wavelength non-fundamental mode components of the nodal diffusion residual to a very significant extent. The acceleration effect realized in this way depends on the numerical proximity of the highest coarse mesh level in the multi-level hierarchy to the full-core diffusion level. Because of the multiplicative rather than additive nature of the correction, coarse mesh rebalancing clearly belongs to the class of *non-linear acceleration approaches* (Trottenberg et al., 2001). Speed-up factors higher than 10, compared to unaccelerated computations, are common for typical steady-state and transient calculations with reactor codes that are accelerated using multi-level rebalancing.

The currently available rebalancing standard (van Geemert, 2006) in PANBOX, PRISM and ARTEMIS™ is aimed at suppressing node-to-node volume-averaged neutronic balance residuals and yielding nodal driving factors (i.e. N driving factors for N volume nodes in the system) with which nodal fluxes and nodal interface currents are corrected multiplicatively before pursuing subsequent computationally expensive nodal diffusion iterations. The principal effect of rebalancing on coarsened space-energy grids is the suppression of slowly converging non-fundamental modes in the eigenspectrum which, in an unaccelerated iteration procedure, would lead to dramatically slow convergence or even convergence stagnation.

With an error reduction ratio σ, a commonly required five-orders-of-magnitude decrease in error would require 5 times $\log\left(\frac{1}{10}\right)/\log(\sigma)$ iterative steps. For example, with $\sigma = 0.98$, these are about 570 steps, and with $\sigma = 0.998$, these amount to about 10,750 steps. Pursuing this many such operations would take clearly far too much time within the framework of large-scale practical applications. This applies even within the context of present-day computational power and in spite of the relatively low computational cost of a single matrix–vector multiplication pertaining to the iterative solution of nodal diffusion equations based on the nodal expansion method (NEM) (Finnemann et al., 1977). Now, for well-converging full core steady-state and transient cases, the currently available rebalancing methodology can generate asymptotic error-reduction ratios of between 0.8 and 0.9 between successive NEM-iterations. For more demanding cases, ratios of between 0.9 and 0.95 still enable acceptable run times. Tedious cases with ratios of higher than 0.95, in spite of applying rebalancing, are known (for example, stuck rod cases under cold conditions). In case no acceleration would be applied, dramatically poor error reduction ratios between 0.98 and 0.997 are common in any case (indicating a typically very small difference between the eigenvalues of the fundamental mode and the first higher mode). Generally, each additional 9 behind the comma makes the convergence about 10 times as slow. This means that the currently available rebalancing methodology is imperative for the ability to pursue full core three-dimensional flux computations within practical time frames. And, in case one would like to be able to pursue full core computations on increasingly finer 3D grids, one should always be interested in noticeable robustness and efficiency improvements for a rebalancing methodology as principal convergence accelerator. Now, one needs to keep in mind that the effectiveness of rebalancing depends heavily on its ability to capture, at the level of the restricted equations, the higher 3D modes of the SP_N equations on the full core grid. An aspect of utmost importance here is the eigenspectrum density of the lower modal part, and in particular the numerical distance of the first higher eigenvalues from the fundamental eigenvalue. Recent numerical investigations have confirmed that, in terms of numerical proximity of neighboring eigenvalues, the eigenspectrum is substantially denser on 3D full core grids than on 2D full core grids. The eigenspectrum can be expected to become even denser when making the 3D grid finer (i.e. from 1box/fa to multi-box/fa up to 3D pin-by-pin resolution).