TWO-LEVEL MULTIGRID PRECONDITIONING OF A NEUTRON NOISE DIFFUSION SOLVER

Antonios Mylonakis, Paolo Vinai, and Christophe Demazière
Chalmers University of Technology
Department of Physics
Division of Subatomic and Plasma Physics
SE-412 96 Gothenburg, Sweden

antmyl@chalmers.se, vinai@chalmers.se, demaz@chalmers.se

ABSTRACT

This paper presents the utilization of a two-level multigrid preconditioner for the acceleration of a two-energy-group neutron noise diffusion solver for fine grid applications. The highly localized nature of most neutron noise sources requires the fine discretization of the spatial domain leading to large systems of algebraic equations. These systems are solved with iterative methods whose performances are usually determined by the accompanying preconditioners. This work applies a two-level multigrid approach aiming to enhance the convergence behavior of the GMRES iterative linear solver. The results show that the two-level scheme improves significantly the performance of GMRES in the solution of two problems. In particular, it outperforms two general-purpose alternative acceleration methods, i.e. ILU(0) and ILUC.

KEYWORDS: neutron noise, multigrid, preconditioning, GMRES

1. INTRODUCTION

Neutron noise in power nuclear reactors consists of stationary fluctuations of the neutron flux that are induced by perturbations such as vibrations of reactor components, oscillations of coolant temperature or density, etc. Since these perturbations can be problematic for the operation of nuclear reactors, their effects need to be analyzed. In addition, the neutron flux fluctuations can be used for diagnostic purposes, i.e. the early detection of anomalies before they have any inadvertent effect on reactor safety and availability [1]. Consequently, the computation of the dynamic reactor transfer function, that is the neutron flux response to a given neutron noise source, is of prime importance.

CORE SIM [2] is a numerical tool able to compute the dynamic transfer function of a reactor. It contains a solver of the two-energy-group neutron noise diffusion equations in the frequency domain. Although the tool has been successfully utilized for the analysis of various coarse-grid configurations, it has limitations when applied to problems with fine spatial discretization. For this reason, an extension of CORE SIM targeting this kind of applications is under-development and this work is part of that research.

The need for fine discretization is implied by the fact that most neutron noise sources are highly localized and induce strong gradients of the neutron flux, mainly in the vicinity of the perturbation [3]. The implementation of fine grids leads to large systems of algebraic equations. For the solution of such large

systems of equations, iterative methods are often employed whose performances are usually determined by the accompanying preconditioners [4]. The objective of this work is to identify a relatively simple and efficient multigrid preconditioner of Krylov methods when solving the neutron noise diffusion equation. The experience of the computational physics community shows that preconditioners do not need to accurately preserve the physics to be effective [5]. Therefore, all the details that render a multigrid method complex when used as a solver, could be of minor importance when used as an acceleration method. In the current effort, a simple V-cycle algorithm with two levels is applied to accelerate the linear solver. The numerical performance is evaluated in two problems: a 1-D homogeneous reactor and a 2-D heterogeneous system.

The structure of the paper is as follows. Section 2 introduces the governing equations of the problem. Section 3 discusses the developed preconditioner. Section 4 presents the results for the two test cases. In Section 5, conclusions are drawn.

2. TWO-GROUP NEUTRON NOISE DIFFUSION EQUATIONS

The two-energy-group Neutron Noise Diffusion Equations (NNDEs) are obtained starting with the time-dependent two-energy-group neutron diffusion equations. All the time-dependent terms, generically expressed as $X(\boldsymbol{r},t)$, are split into their mean values $X_0(\boldsymbol{r})$ and their fluctuations $\delta X(\boldsymbol{r},t)$ written as: $X(\boldsymbol{r},t) = X_0(\boldsymbol{r}) + \delta X(\boldsymbol{r},t)$. Considering small fluctuations compared to the mean values and stationary processes, and neglecting the second order terms (linear theory), the NNDEs are obtained. Then, a temporal Fourier transform is applied to derive the equations in the frequency domain, which read as:

$$\left[\nabla \cdot D\left(\boldsymbol{r}\right)\nabla + \Sigma_{dyn}^{crit}(\boldsymbol{r},\omega)\right] \times \begin{bmatrix}\delta\phi_1(\boldsymbol{r},\omega)\\ \delta\phi_2(\boldsymbol{r},\omega)\end{bmatrix} = \\ \phi_r(\boldsymbol{r})\delta\Sigma_r(\boldsymbol{r},\omega) + \phi_a(\boldsymbol{r}) \begin{bmatrix}\delta\Sigma_{a,1}(\boldsymbol{r},\omega)\\ \delta\Sigma_{a,2}(\boldsymbol{r},\omega)\end{bmatrix} + \phi_f^{crit}(\boldsymbol{r},\omega) \begin{bmatrix}\delta\upsilon\Sigma_{f,1}(\boldsymbol{r},\omega)\\ \delta\upsilon\Sigma_{f,2}(\boldsymbol{r},\omega)\end{bmatrix} \quad (1)$$

where Σ_{dyn}^{crit}, ϕ_r, ϕ_a and ϕ_f^{crit} are defined in APPENDIX A. The RHS represents the neutron noise sources which are modelled as fluctuations of the macroscopic nuclear cross-sections around their mean value.

In this work, NNDEs are spatially discretized with the nodal approach described in [2]. According to that method, the nuclear reactor core is considered as a juxtaposition of adjacent nodes and the governing equations are spatially-averaged on each one of them. The diffusion operator is discretized based on finite-differences. The fully discretized NNDEs take the form:

$$\boldsymbol{A\Phi} = \boldsymbol{S} \quad (2)$$

where $\boldsymbol{A}$ is a non-symmetric complex coefficient matrix, $\boldsymbol{\Phi}$ is the vector of the unknowns and $\boldsymbol{S}$ the vector containing the noise sources.

3. TWO-GRID PRECONDITIONING OF THE KRYLOV SOLVER

The applied discretization algorithm generates a sparse system of linear algebraic equations that in case of fine grid problems is also large. Large sparse linear systems are not solved with standard direct methods like matrix inversion or decomposition. The main reasons are that direct methods imply fill-in and they usually scale poorly with the problem size. Instead, iterative solvers like GMRES are better

suited. GMRES [6] is a Krylov-class method oriented to large sparse non-symmetric real and complex linear systems. The convergence rate is determined by the spectral properties of the coefficient matrix. Therefore, GMRES is usually combined with a preconditioner aiming to obtain an eigenvalue distribution that favors convergence. If $\boldsymbol{M}$ is a non-singular matrix that approximates $\boldsymbol{A}$ in some sense, the left-preconditioned form of Eq. 1 has the form:

$$\boldsymbol{M}^{-1}\boldsymbol{A}\boldsymbol{\Phi} = \boldsymbol{M}^{-1}\boldsymbol{S} \tag{3}$$

Incomplete LU decomposition (ILU) is a family of typical methods used to construct $\boldsymbol{M}^{-1}$. However, in case of large systems of equations, they may meet difficulties regarding memory and computational time. Typically, the whole fill-in is discarded during the factorization process generating the simple but still powerful zero fill-in ILU version (denoted as ILU(0)) that is applicable even to large problems.

For fine grid problems, a preconditioner based on the multigrid method (MG) may constitute a good choice because it does not require any inversion or decomposition operation. Besides, the main work is performed on coarser grids than the original one reducing the computational effort. MG [7] is based on the fact that iterative solvers like Jacobi and Gauss-Seidel produce smooth errors in the approximation of the solution. Such iterations nearly remove the high frequency modes of the error in a few iterations. However, low frequency modes are reduced very slowly. The central idea of MG is to transfer the error to a coarser grid, where low frequencies can be removed faster. Many versions of MG exist. In this work, a Two-Level (TL) V-cycle scheme is utilized as a preconditioner of GMRES. Usually, cycles with more than one coarse levels are preferred. However, a TL scheme constitutes a good starting point for investigation.

3.1. PRECONDITIONING ALGORITHM

In the context of a Krylov method like GMRES, a preconditioner $\boldsymbol{M}^{-1}$ is acting only in terms of computing matrix-vector products of the form:

$$\boldsymbol{u} = \boldsymbol{M}^{-1}\boldsymbol{b} \tag{4}$$

This is equivalent to obtain the solution of:

$$\boldsymbol{M}\boldsymbol{u} = \boldsymbol{b} \tag{5}$$

The optimal preconditioning, at least as far as convergence is concerned, is the choice $\boldsymbol{M} = \boldsymbol{A}$. This operation would require to solve:

$$\boldsymbol{A}\boldsymbol{u} = \boldsymbol{b} \tag{6}$$

This task is as difficult as the original problem (Eq. 2) and is never performed. Instead, one may solve Eq. 6 loosely to generate an approximate solution. A good preconditioner is one for which $\boldsymbol{M}$ is a good approximation to $\boldsymbol{A}$ and $\boldsymbol{M}^{-1}\boldsymbol{b}$ is inexpective to compute. In this work, a TL V-cycle solver is employed as a preconditioner of GMRES, i.e. it provides GMRES with estimates of $\boldsymbol{u}$ through Eq. 6.

TL V-cycle uses two grids: the fine one denoted by h and a coarser one specified by nh where n stands for an integer larger than 1. The integer n defines the node size ratio between the fine and the coarse grid at each dimension. The method performs m smoothing iterations on h obtaining an approximation of the solution of the linear system $\mathbf{A}_h\boldsymbol{u}_h = \boldsymbol{b}_h$. Then the residual $\boldsymbol{R}_h$ is calculated and transferred to nh. Thereafter, the linear system $\mathbf{A}_{nh}\mathbf{e}_{2h} = \mathbf{R}_{nh}$ is solved to find the error $\mathbf{e}_{nh}$. After a transfer of the error back to h, $\boldsymbol{u}_h$ is corrected as $\mathbf{u}_h \leftarrow \mathbf{u}_h + \mathbf{e}_h$. A predefined number of k cycles are performed. The general overview of the algorithm used in this work is summarized in Alg.1.

Algorithm 1 Two-level V-cycle preconditioner iteration

for $i = 1, k$

1. Smooth: Perform m iterations to approximate $\boldsymbol{u}_h$ from Eq. 6
2. Get residual: $\boldsymbol{R}_h = \boldsymbol{b}_h - \boldsymbol{A}_h \boldsymbol{u}_h$
3. Restrict: $\boldsymbol{R}_{nh} = \boldsymbol{I}_h^{nh} \boldsymbol{R}_h$
4. Solve: $\boldsymbol{A}_{nh} \boldsymbol{e}_{nh} = \boldsymbol{R}_{nh}$
5. Prolong: $\boldsymbol{e}_h = \boldsymbol{I}_{nh}^h \boldsymbol{e}_{nh}$
6. Correct: $\boldsymbol{u}_h \leftarrow \boldsymbol{u}_h + \boldsymbol{e}_h$

end

The design of a MG scheme for a neutron diffusion problem requires the selection of the following components: smoother method, solver method, restriction/prolongation operators and techniques for preserving physical quantities, e.g. reaction rates, during intergrid operations. The selected smoother (step 1 of Alg.1) is the Gauss-Seidel method. The solver of the error equation at step 4, called *inner* solver, is chosen to be GMRES preconditioned with ILU(0). The restriction ($\boldsymbol{I}_h^{nh}$) and the prolongation ($\boldsymbol{I}_{nh}^h$) operators provide the communication of the error between the two grids (steps 3 and 5). The restriction operator is constructed considering that the nodal values on h may be related to the quantities on nh through averaging. Since each node on h corresponds exactly to n^d nodes on nh, where d is the number of dimensions of the spatial domain, the restriction operator performs the following action:

$$R_{nh} = \frac{1}{n^d} \sum_{i=1}^{n^d} R_{h,i} \tag{7}$$

The restriction matrix has dimensions $2N_{nh} \times 2N_h$ where N symbolizes the number of spatial nodes. The prolongation operator is defined as:

$$\boldsymbol{I}_{nh}^h = n^d (\boldsymbol{I}_h^{nh})^\top \tag{8}$$

The operator $\boldsymbol{I}_{nh}^h$ copies a node value of a coarse grid to all the corresponding nodes of the fine one. By using matrix operators instead of loops, the process of transferring quantities between the two grids is significantly accelerated.

In the general case, the selected intergrid operators do not preserve the physics. An accurate restriction operator would consider a noise-weighted cross-section homogenization process. In addition, the prolongation operation could be based on shape functions etc. Utilization of accurate operators would imply not only high developmental work, but also increase of the computational effort. For example, the restriction operator should be constructed on-the-fly within GMRES based on the current estimate of the solution vector. However, MG may not need to preserve the physics accurately to be effective since it operates as an acceleration method of the Krylov solver and not as a solver itself. This attractive feature has also be pointed out in [5]. Another reason of avoiding cross-section homogenization comes from the nature of the problems of interest. Since neutron noise problems are strongly localized, the applied computational grids must be much finer than the size of the homogenized parts of the configuration. This implies that even when transferring to the coarsest mesh, no spatial cross-section homogenization may be needed. Consequently, the intergrid operators are constructed only once, at the initialization of the solver, making the preconditioning algorithm simple and fast.

4. NUMERICAL RESULTS

The developed preconditioner is evaluated with two test cases: a 1-D one-region homogeneous reactor core and a 2-D heterogeneous system. In both cases, a localized, fluctuating perturbation is introduced in the system. The performance of the solver is compared with the ones of free GMRES, and GMRES preconditioned with ILU(0) and ILUC(*droptol*). For ILUC a drop tolerance equal to 10^{-5} has been selected. For the first test case, the solver is benchmarked with a semi-analytical solution.

4.1. 1-D one-region homogeneous reactor

The first problem consists of a 1-D one-region homogeneous reactor near criticality with a size of 301 cm. The macroscopic cross-sections and the diffusion coefficients are selected to be representative of a typical Light Water Reactor. Their values along with the dynamic parameters are reported in Tables (1 & 2). A uniform grid with a node size of Δx = 0.125 cm is used (2408 nodes) leading to a linear system of 4816 equations. The specified configuration generates a coefficient matrix **A** with an estimated condition number equal to 5.0684×10^6. The large condition number indicates that GMRES will not perform well without efficient preconditioning.

Table 1: Macroscopic cross-sections and diffusion coefficients of the 1-D test case.

$\Sigma_{a,1}\ (cm^{-1})$	$\Sigma_{a,2}\ (cm^{-1})$	$\upsilon\Sigma_{f,1}\ (cm^{-1})$	$\upsilon\Sigma_{f,2}\ (cm^{-1})$	$\Sigma_r\ (cm^{-1})$	$D_1\ (cm)$	$D_2\ (cm)$
0.0115	0.1019	0.0057	0.1425	0.0151	1.4376	0.3723

Table 2: Dynamic parameters of the 1-D test case.

$\beta_{eff}\ (-)$	$l\ (s^{-1})$	$\nu_1\ (cm/s)$	$\nu_2\ (cm/s)$
0.00535	0.0851	18230400	413067

First, the convergence property of the TL V-cycle iteration is studied. The preconditioner can be considered as an iterative process of the form [8]:

$$\boldsymbol{u}_h^{(new)} = \boldsymbol{M}_h \boldsymbol{u}_h^{(old)} + \boldsymbol{g}_{M_h} \tag{9}$$

where $\boldsymbol{M}_h$ is the iteration matrix of the process. Such a process converges for all initial guesses if $\rho(\boldsymbol{M}_h) < 1$, where ρ symbolizes the spectral radius. To obtain an estimation of the spectral radius of TL V-cycle, the iteration matrix is built following the steps of Alg. 1. The first step includes the Gauss-Seidel smoother that is based on the following coefficient matrix splitting: $\boldsymbol{A} = \boldsymbol{L} + \boldsymbol{D} + \boldsymbol{U}$ where $\boldsymbol{D}$ is a diagonal matrix and $\boldsymbol{L}$ and $\boldsymbol{U}$ are strictly lower and upper triangular matrices, respectively. Then the Gauss Seidel iteration can be defined as:

$$\boldsymbol{u}_{h,GS}^{(new)} = \boldsymbol{G}_h \boldsymbol{u}_h^{(old)} + f(\boldsymbol{b}) \tag{10}$$

where $\boldsymbol{G}_h = \boldsymbol{I} - (\boldsymbol{D} + \boldsymbol{L})^{-1}\boldsymbol{A}$, with $\boldsymbol{I}$ being the identity matrix, and $f(\boldsymbol{b}) = (\boldsymbol{D} + \boldsymbol{L})^{-1}\boldsymbol{b}$. For the purpose of convergence analysis it is enough to consider the homogeneous linear system [7], i.e taking $\boldsymbol{b} = \boldsymbol{0}$. Then, one can notice that the result of the smoother is a repetitive application of the Gauss-Seidel

iteration matrix on the initial guess. Following the subsequent steps of Alg. 1, the following relation is built:

$$\boldsymbol{u}_h^{(new)} = (\boldsymbol{G}_h)^m \boldsymbol{u}_h^{(old)} - \boldsymbol{I}_{nh}^h \boldsymbol{A}_{nh}^{-1} \boldsymbol{I}_h^{nh} \boldsymbol{A}_h (\boldsymbol{G}_h)^m \boldsymbol{u}_h^{(old)} \tag{11}$$

where m corresponds to the number of smoothing iterations. After rearrangement of the terms one obtains:

$$\boldsymbol{u}_h^{(new)} = [(\boldsymbol{I} - \boldsymbol{I}_{nh}^h \boldsymbol{A}_{nh}^{-1} \boldsymbol{I}_h^{nh} \boldsymbol{A}_h)(\boldsymbol{G}_h)^m] \boldsymbol{u}_h^{(old)} \tag{12}$$

Therefore, the iteration matrix of the TL scheme is:

$$\boldsymbol{M}_h = (\boldsymbol{I} - \boldsymbol{I}_{nh}^h \boldsymbol{A}_{nh}^{-1} \boldsymbol{I}_h^{nh} \boldsymbol{A}_h)(\boldsymbol{G}_h)^m \tag{13}$$

Since this test case generates relatively small matrices, the spectral radius can be computed as the maximum modulus of the eigenvalues of matrix $\boldsymbol{M}_h$. Additionally, the spectral radius is estimated based on the asymptotic error reduction of the preconditioner iteration [8]. Fig. 1 illustrates the spectral radii for $n = 4$ and $m = 1, 2, 3, 4$ and 5, estimated with the two methods. The actual behavior of the preconditioning iteration is in good agreement with the predicted behavior. It is also illustrated that both 4 and 5 smoothing iterations lead to low spectral radius indicating satisfactory numerical performance. The spectral radius is an asymptotic measure of convergence because it predicts the worst-case error reduction over many iterations. Thus, it does not generally predict the behavior for a small number of iterations that a preconditioner can perform. Nevertheless, it provides an indication of the quality of the developed iterative process.

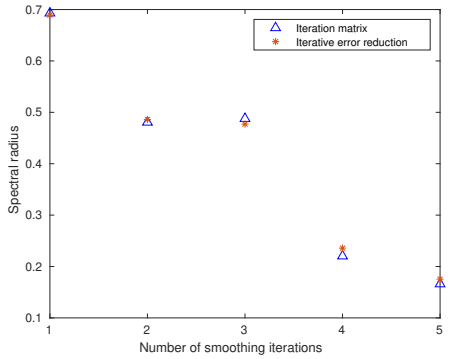

Figure 1: Spectral radius with smoothing iterations.

After the convergence analysis, a point-like 1 Hz perturbation of the macroscopic removal cross-section (Σ_r) located at $x = -50$ cm is simulated. GMRES is accelerated with TL V-cycle using $n = 4$, $m = 5$ and $k = 2$. The convergence tolerance of the inner solver (step 4, Alg. 1) was chosen to be 10^{-10}. Fig. 2a confirms the difficulty of the free GMRES in which the relative residual error is reduced very slowly with the number of iterations. TL V-cycle is proven very efficient leading to a rapid convergence of GMRES, i.e. after only 7 iterations. ILU(0) provides the lowest gain in terms of iteration reduction as expected. ILUC(10^{-5}) reduces the number of iterations compared to ILU(0). However, the selected drop tolerance leads to a matrix that is $\sim$55 times denser than $\mathbf{A}$. Such a dense matrix is characterized by high construction cost and implies slower matrix-vector multiplications. These features may counterbalance the profit from the reduction in the number of iterations. In addition, such a dense matrix could be

proven non-applicable to 3-D full-core problems regarding computer memory and construction time. Fig. 2b shows that the selected grid reproduces accurately the analytical solution. The steep gradients around the point-like noise source are properly captured.

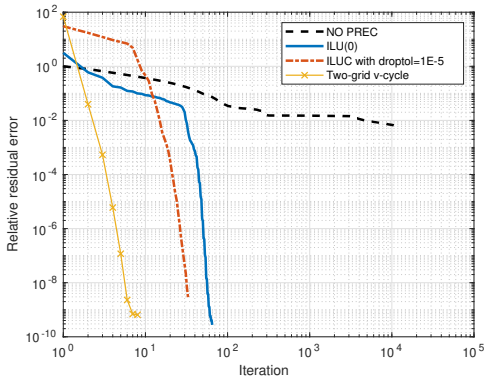

(a) Residual error of the linear solver.

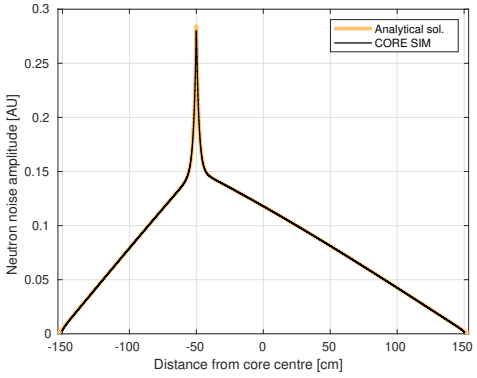

(b) Amplitude of the thermal neutron noise.

Figure 2: GMRES convergence (a) & numerical vs analytical solution (b).

4.2. 2-D heterogeneous system

The second test case is related to the simulation of a neutron noise source in a 2-D heterogeneous configuration based on the C3 benchmark on power distribution within assemblies, available from the OECD/NEA [9]. The configuration consists of four fuel assemblies, two containing UO2 fuel and two MOX fuel. The system has a size of 43.84 cm $\times$ 43.84 cm and is radially surrounded by reflective boundaries. The noise source is a 5 $\%$ perturbation of the fast and thermal absorption cross-sections of the fuel cell (19,16) located inside one of the MOX fuel assemblies (Fig. B.1). The frequency of the pertubation is 1 Hz. A uniform grid with $\Delta x = \Delta y = 0.1575$ cm is used (73,984 nodes) leading to a linear system with 147,968 algebraic equations. The configuration described above generates a coefficient matrix with an estimated condition number equal to 3.99×10^6. The large condition number indicates that free GMRES will not converge, as in the previous test case.

GMRES is accelerated with TL V-cycle using $n = 4$, $m = 4$ and $k = 2$. For the inner solver, a tight convergence tolerance (10^{-10}) was again selected. Fig. 3 demonstrates that the relative residual error of free GMRES reduces very slowly with the number of iterations, showing the inadequacy of the method for the current class of problems. TL V-cycle is however able to notably accelerate GMRES in terms of both number of iterations and computational time, as Fig. 3 and Table 3 show. More specifically, it allowed GMRES to converge very fast, i.e. after only 7 iterations. This is translated into a reduction in the number of iterations by factors of $\sim$1785 and $\sim$828 compared with ILU(0) and ILU(10^{-5}), respectively. In terms of computational time, it implied a reduction by factors of $\sim$5 and $\sim$22 compared with the two alternative preconditioners, respectively. The neutron noise computed with GMRES-TL V-cycle is illustrated in Fig. B.2. The maximum relative differences in the calculated neutron noise amplitude and phase between GMRES-TL V-cycle and GMRES-ILUC(10^{-5}), are 6.5×10^{-6} $\%$ and 9.22×10^{-7} $\%$. Thus, the two solutions are in good agreement.

ILUC(10^{-5}), as also concluded in Section 4.1, would be proven non-applicable to realistic cases because of the prohibitive computer memory requirements caused by the substantial fill-in ($O(10^8)$). Large fill-in also makes matrix-vector multiplications very slow counterbalancing the achieved reduction in the number of iterations. ILU(0), although significantly faster than ILUC, fails to reach a tight convergence tolerance even after 12500 iterations. These findings highlight that general-purpose algebraic methods may be inefficient in accelerating GMRES when applied to fine-grid neutron noise problems. This outcome reinforces the significance of using MG methods to solve the target problems efficiently.

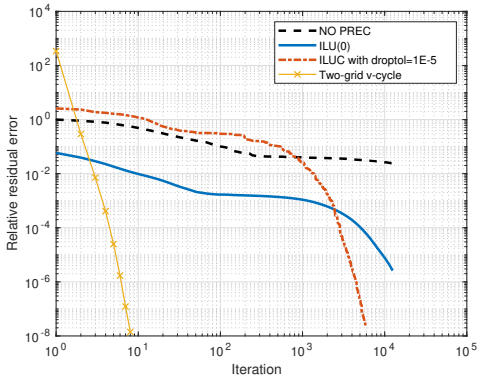

Figure 3: Residual error of the linear solver.

Table 3: Iteration and CPU time of GMRES solver accelerated by different preconditioners.

Preconditioner	Preconditioner building time (s)	Solver iterations	Solver time (s)
-	-	12500	545
ILU(0)	0.02	12500	634
ILUC(10^{-5})	297	5800	2760
TG V-cycle	2.3	7	121

5. CONCLUSIONS

A simple and fast TL V-cycle preconditioner was utilized to accelerate GMRES when solving the two-energy-group neutron noise diffusion equations. The numerical analysis of the technique was carried out for neutron noise examples in a 1-D homogeneous reactor core and in a 2-D heterogeneous system. A fine spatial discretization of the systems was applied for an accurate evaluation of the sharp gradients of the neutron noise that arise from the localized character of the perturbation. TL V-cycle was proven very efficient in accelerating GMRES in terms of both number of iterations and computational time. On the other hand, GMRES coupled with the general purpose ILU(0) was not able to reach a tight convergence tolerance. ILUC, although it allowed GMRES to converge, is much slower and possibly impractical because it induces very high fill-in. These findings highlight the strength of the developed preconditioner for 1-D and 2-D problems and suggest that the investigation is extended to 3-D configurations.

ACKNOWLEDGEMENTS

The research leading to these results has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 754316.

REFERENCES

[1] I. Pazsit and C. Demaziere. Noise Techniques in Nuclear Systems. In: *Handbook of Nuclear Engineering*, **3**, Springer, Boston, MA (2010).

[2] C. Demaziere. "CORE SIM: A multi-purpose neutronic tool for research and education." *Annals of Nuclear Energy*, **38**, pp. 2698–2718 (2011).

[3] A. Mylonakis, H. Yi, P. Vinai, and C. Demaziere. "Neutron noise simulations: a comparison between a diffusion-based and a discrete ordinates solver." In *International Conference on Mathematics and Computational Methods applied to Nuclear Science and Engineering - M&C* (2019).

[4] M. Benzi. "Preconditioning Techniques for Large Linear Systems: A Survey." *Journal of Computational Physics*, **182**, pp. 418–477 (2002).

[5] R. Slaybaugh, T. Evans, G. Davidson, and P. Wilson. "Multigrid in energy preconditioner for Krylov solvers." *Journal of Computational Physics*, **242**, pp. 405 – 419 (2013).

[6] Y. Saad and M. Schultz. "GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems." *SIAM J Sci Stat Comput*, **7**, pp. 856–869 (1986).

[7] W. L. Briggs. *A Multigrid Tutorial*. SIAM (1987).

[8] Y. Saad. *Iterative methods for sparse linear systems*. SIAM (Society for Industrial and Applied Mathematics) (2003).

[9] C. Cavarec, J. Perron, D. Verwaerde, and J. West. "Benchmark calculations of power distribution within assemblies." Technical Report NEA/NSC/DOC (94) 28 (1994).

APPENDIX A. Definition of Σ_{dyn}^{crit}, ϕ_f^{crit}, ϕ_r and ϕ_a

The matrix Σ_{dyn}^{crit} is:

$$\Sigma_{dyn}^{crit}(\boldsymbol{r},\omega) \equiv \begin{bmatrix} -\Sigma_1^{crit}(\boldsymbol{r},\omega) & \frac{\upsilon\Sigma_{f,2,0}(\boldsymbol{r})}{k_{eff}}\left(1-\frac{\iota\omega\beta}{\iota\omega+\lambda}\right) \\ \Sigma_{\boldsymbol{r},0}(\boldsymbol{r}) & -\left(\Sigma_{a,2,0}(\boldsymbol{r})+\frac{\iota\omega}{v_2}\right) \end{bmatrix} \tag{A.1}$$

where

$$\Sigma_1^{crit}(\boldsymbol{r},\omega) \equiv \Sigma_{a,1,0}(\boldsymbol{r}) + \frac{\iota\omega}{v_1} + \Sigma_{\boldsymbol{r},0}(\boldsymbol{r}) - \frac{\upsilon\Sigma_{f,1,0}(\boldsymbol{r})}{k_{eff}}\left(1-\frac{\iota\omega\beta}{\iota\omega+\lambda}\right) \tag{A.2}$$

The matrix ϕ_f^{crit} is:

$$\phi_f^{crit}(\boldsymbol{r},\omega) \equiv \begin{bmatrix} -\frac{\phi_{1,0}(\boldsymbol{r})}{k_{eff}}(1-\frac{\iota\omega\beta}{\iota\omega+\lambda}) & -\frac{\phi_{2,0}(\boldsymbol{r})}{k_{eff}}(1-\frac{\iota\omega\beta}{\iota\omega+\lambda}) \\ 0 & 0 \end{bmatrix} \tag{A.3}$$

The matrix ϕ_a is:

$$\boldsymbol{\phi}_a(\boldsymbol{r}) \equiv \begin{bmatrix} \phi_{1,0}(\boldsymbol{r}) & 0 \\ 0 & \phi_{2,0}(\boldsymbol{r}) \end{bmatrix} \tag{A.4}$$

The column vector ϕ_r is:

$$\phi_r(r) \equiv \begin{bmatrix} \phi_{1,0}(\boldsymbol{r}) \\ -\phi_{1,0}(\boldsymbol{r}) \end{bmatrix} \tag{A.5}$$

APPENDIX B. Neutron noise source location & numerical solution of Section 4.2

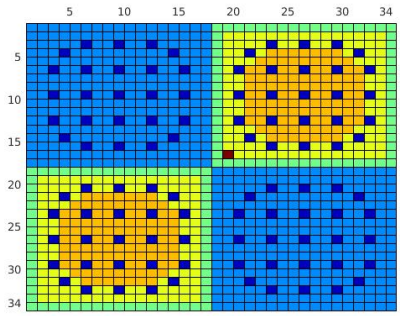

Figure B.1: C3 benchmark configuration and location of the noise source (indicated in red).

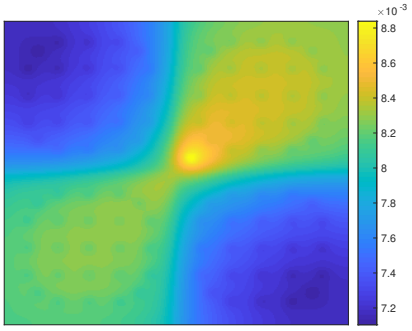

(a) Amplitude of fast noise [AU].

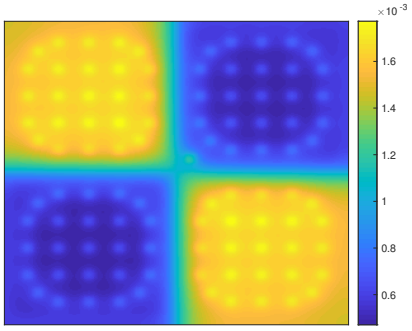

(b) Amplitude of thermal noise [AU].

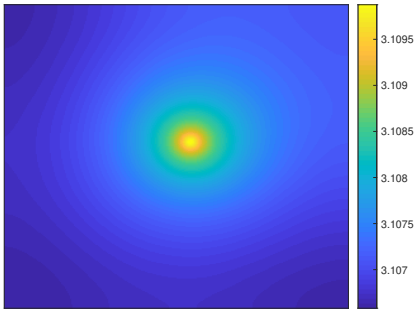

(c) Phase of fast noise [rad].

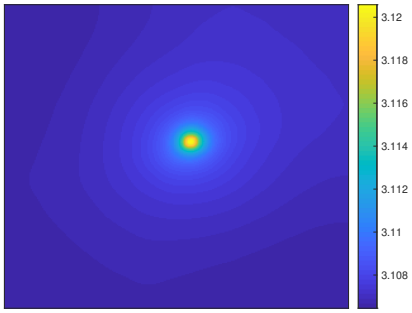

(d) Phase of thermal noise [rad].

Figure B.2: Computed neutron noise for the 2-D heterogeneous system.