International Conference on Mathematics, Computational Methods & Reactor Physics (M&C 2009)
Saratoga Springs, New York, May 3-7, 2009, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2009)

PERFORMANCE OF WHOLE CORE PIN-BY-PIN CALCULATIONS BY DOMAIN DECOMPOSITION THROUGH ALTERNATE DISSECTIONS IN STEADY STATE AND TRANSIENT CALCULATIONS

José-Javier Herrero, Javier Jiménez, José-María Aragonés and Carol Ahnert
Departamento de Ingeniería Nuclear
Universidad Politécnica de Madrid
c/ José Gutiérrez Abascal, 2; 28006 Madrid (Spain)
herrero@din.upm.es; jimenez@din.upm.es; arago@din.upm.es; carol@din.upm.es

ABSTRACT

Domain decomposition through alternate dissections methodology is applied to solve the transport corrected fine mesh finite difference multigroup diffusion problem at the homogenized pin-cell scale level for full cores including a nodal acceleration. Problems are solved in parallel computer architectures and the efficiency of the coding and method is addressed for steady state problems with thermal-hydraulic feedback and transients without thermal-hydraulic feedback on the cross sections for the moment.
The main objective is to be able to compute full core problems using reasonable computing times and resources, but obtaining higher detail in the flux and power distribution, with the ability to lead to detailed values of other variables such as burnup. The diffusion code COBAYA3 is able to use this development together with the cross sections feedback capability already implemented to couple with the parallelized thermal-hydraulics code COBRA-TF using the same domain decomposition approach. Computations are performed for the PWR MOX/UO_2 Core Transient Benchmark in 8 energy groups.

Key Words: fine mesh diffusion, multiscale, multiphysics, domain decomposition

1. INTRODUCTION

Computation of full nuclear reactor cores with a level of detail fixed at the homogenized pin-cell scale has been already investigated by a number of organizations with the objective of avoiding the errors associated with homogenization methods and pin power reconstruction techniques in cores with highly heterogeneous fuel assemblies. Besides, there is a need at the cell scale to efficiently implement the capability of performing feedbacks in the cross sections with a thermal-hydraulics subchannel code able to take into account the detail of the spatial fuel temperature distribution inside each fuel assembly, cross flows and void distributions in two phase flows during a transient simulation from criticality [3].

In this work the efficiency of computing such problems by domain decomposition using the parallelized lattice code COBAYA3 is shown in section 4 for steady states and transients at hot zero power without thermal-hydraulics, and for steady states at hot full power conditions with a thermal-hydraulics solution from COBRA-TF. Also computational efficiency analyses are presented in section 3, showing the influence of taking different numbers of processors for the same case.

The verification cases are based on different configurations using the cross sections provided in the PWR MOX/UO_2 Core Transient Benchmark in 8 energy groups at pin level.

2. DOMAIN DECOMPOSITION

2.1. Convergence of the alternate dissections methodology

All the development of the alternate dissections is based on the Schwarz alternating method, whose convergence is proven for overlapping subdomains [10][11]. And on the heuristic numerical performance in 2D and 2-group diffusion calculations of PWR cores [1].

It is possible to divide a PWR core geometry in quarters of assemblies comprehending the whole axial length, these quarters can be grouped following four alternate dissections formed by single assemblies (one color-sets), for quarters of neighboring assemblies (4 color-sets), or combinations of half assemblies in the horizontal and vertical directions (2 color-sets).

The fundamental paradigm of the alternate dissections resides in the fact that the centerplanes of each subdomain considered in this way coincides with the boundary planes of other subdomains formed in different dissections, so as to compute updated values for the fine-mesh boundary conditions in these centerplanes on each iteration. Moreover, as each of the four partitions considered fully overlap with the rest of partitions; the convergence rate of the recomputed boundary conditions is considerably faster.

Let’s suppose the fission source in each cell is fixed and also the k eigenvalue, then the boundary conditions for the next iteration can be computed from the centerplanes of each set of subdomains at the end of the current iteration; and therefore, a fully converged solution will be achieved after some iterations over the 4 dissections classes. The number of iterations at the level of the cell is very much reduced when a nodal solution using homogenized cross sections and heterogeneity factors from the pin-by-pin solution is introduced in between each set of subdomains solutions; this nodal computation is performed using the ANDES code [12] which provides an accurate solution of the fission source distribution using de ACMFD formulation [2][4]. The nodal values are applied to the detailed solution by interpolating the boundary conditions on each nodal face adding a piecewise constant. After the convergence of the boundary conditions, the fission source for each cell is recomputed together with the core k eigenvalue, and a new fixed source problem is solved till convergence of the full core fission source.

In this development we go one step further and the fission source problem is solved on each subdomain, getting a k eigenvalue for each one, and also solving the eigenvalue problem for the nodal acceleration. In this case, the convergence of the fission source is very much improved by the nodal solution; even more, the nodal acceleration is absolutely necessary when the problem is solved in multigroups and three dimensional geometries. The reason for that comes from the fact that the convergence of the Schwarz alternating method is not ensured for this kind of mixed eigenvalue plus boundary conditions problem.

For the case of the thermal-hydraulics, the same domain decomposition is applied; but in this case the nodal acceleration, or better the channel acceleration, is not needed because the convergence is already ensured as the source terms are not updated for each solved dissection. The subchannel problem is found to converge in a reduced number of iterations for nominal PWR conditions, where the cross flows are not of much importance and there is no void distribution in the coolant.

2.2. Parallelized scheme for coupled neutronics and thermal-hydraulics computations

The core partition in quarters of assemblies is a physically meaningful way of domain decomposition that can be applied both in the neutronics code COBAYA3 [5][6] and the thermal-hydraulics code COBRA-TF [7][8] which are the object of parallelization. Figure 1 is a graphical representation of the level of detail considered for the fine-mesh transport corrected diffusion solver COBAYA3 and the two phase-flow code COBRA-TF.

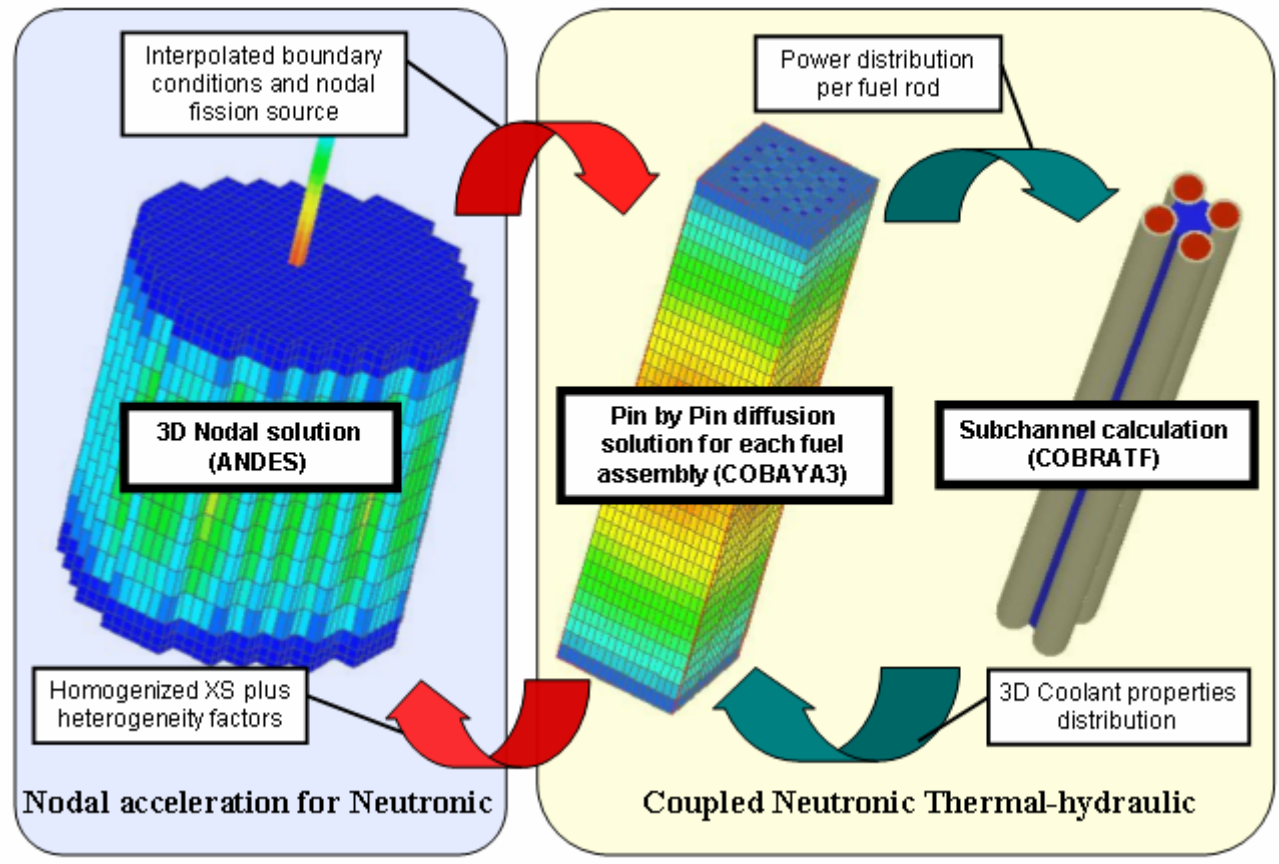

Figure 1. Overview of the multiscale coupling scheme

To reduce computing times, each formed subdomain is solved simultaneously to other subdomains in different processors. The boundary conditions on the subdomain and their returned updated values at the end of each computation for the centerplanes of the subdomain are the information to be transmitted mainly. But also volume information is passed to initialize the

values for the next subdomain computation and keep track of the convergence to the full core solution, and to compute the nodal solution.

The selected boundary conditions to use in the neutronics are the fine-mesh current to flux ratios on each cell interface, because they are independent of the level of flux in the subdomain and leave the diffusion equation in a homogeneous form. Therefore, they are insensible to the fission source levels supposed over the core in order to converge the k-eigenvalue problem; it is a necessary feature of the method because the true fission source levels over the core are not known a priori. On the other hand, the domain decomposition is unable to produce the correct source level for the core, it only can improve the fission source distribution inside the subdomain. The alternate dissections scheme makes it possible to reach the true fission source distribution by "mixing" the quarters of assemblies. In that sense, the nodal solution is in charge of improving the fission source levels for the full core problem which are used to normalize the fluxes of each subdomain in the pin-by-pin solution.

The processors compute the k eigenvalue and fission source distribution for each isolated subdomain, then new current to flux ratios are computed in the centerplanes. When convergence is reached in the full core, k eigenvalue coincides for all the subdomains, the pin powers are stored and then used by the thermal-hydraulics code as a fixed source term.

The thermal-hydraulics solution uses a quasi-stationary method to reach the steady state conditions because, up to now, COBRA-TF does not have a steady state solution implemented (this will be addressed in future developments). The solution initiates at zero time and a void transient is run till convergence is reached for each subdomain, then cross flows are computed in the centerplanes of the geometry which are subchannel boundaries where boundary conditions will be applied for the next dissection. Those boundary conditions are formed by the three fields exchanged between neighboring subchannels in the centerplanes of the subdomains, namely mass, enthalpy, and momentum.

After each dissection, all the subdomains are said to be in the same instant of time, and a restart case is computed for the next dissection using the updated cross flows and the already computed fields as initial conditions. It is remarkable that the amount of information to be transmitted at the moment is 15 times more for thermal-hydraulics than for neutronics, this could be balanced with the inclusion of burnup or equilibrium Xenon in the computation for instance.

A comparison is established between the temperatures and densities fields computed at each iteration and the precedent to decide the moment of convergence. In nominal conditions only two to three different alternate dissections are needed for the thermal-hydraulics to converge the full core temperatures and densities distributions. Afterwards, the cell temperatures and densities are stored and used by the pin-by-pin solver to interpolate in the cross sections tables in a new pin power computation.

Figure 2 represents the flow chart of this coupled calculation which also applies a damping on the thermal-hydraulics solution to accelerate the convergence to the steady state and avoid oscillations. Other coupling schemes at the subdomain level are being investigated.

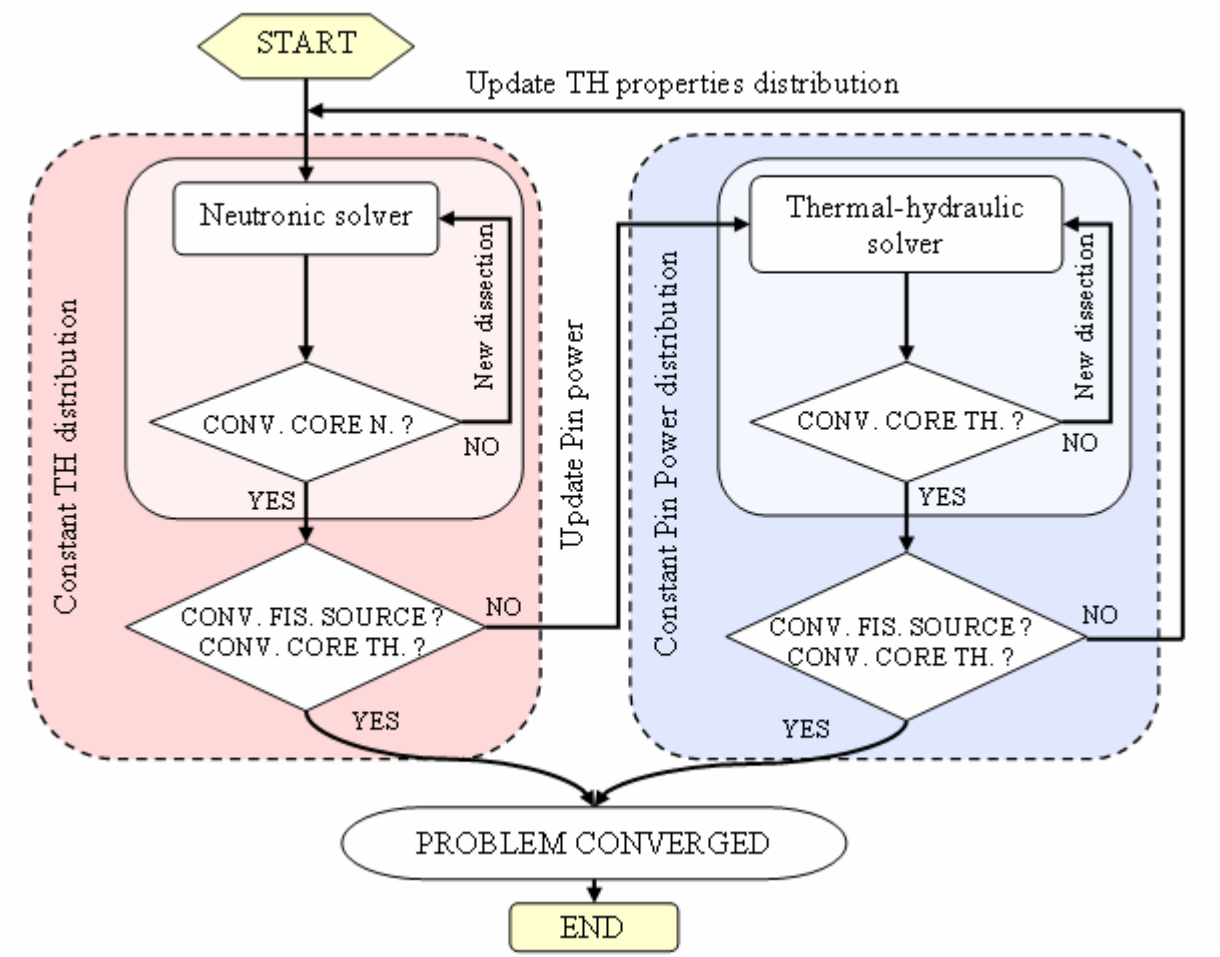

Figure 2. Implementation of the convergence logic in Alternate Dissection

After each neutronic or thermal-hydraulic full core solution, the new values are compared against the ones stored at the beginning of the alternate dissections. When the change is small enough for the pin powers and for the temperatures and densities fields used in the cross-section feedback, the problem is considered solved.

3. IMPLEMENTATION IN THE COBAYA3 CODE

The methodology presented has been implemented in the COBAYA3 code using the OpenMPI library, the code is developed in FORTRAN 90 and uses LAPACK and SPARSKIT numerical libraries to solve the linear systems.

The master-slave methodology chosen uses one single reserved process to manage the information treatment; this process prepares the values to be sent to each subdomain and sends that information to one subdomain at a time. All the information is packed in a single buffer and then sent in a single hit to the MPI library, so as to avoid generating constantly all the wrap-up that accompanies message sending.

The master process sends sequentially each subdomain to free processors, meaning that they are not computing a subdomain at that time. When there are no more processors available, it waits until one finishes his computations and sends results, and then it continues sending new subdomains until the full core has been computed for the corresponding type of dissection.

Thereafter, when the fission source is being computed, a nodal solution is performed in the master process and the interpolation values are generated. Again a new dissection, different from the previous one, begins where the updated nodal values are used to interpolate the boundary conditions in each slave process and the nodal fission source level is used to renormalize the subdomains.

The developed code has been tested in a single 8 processors machine, a cluster of 48 processors, and a supercomputer with more than 1000 processors available. Actually, the optimum number of processors needed clearly depends on the size of the problem, thus, having more processors than the maximum number of subdomains formed, does not reduce the overall CPU time. And, on the other side, having less available processors than the maximum number of subdomains can lead to a situation in which all the subdomains are computed except one, and the execution in the master process have to wait till the last one is finished before going for a new dissection.

One important weakness for the master-slave method is that the amount of memory required to store all the detailed values for the neutronics and the thermal-hydraulics for the case of a full core is more than the available memory in typical computing nodes. The memory required for a single fuel assembly including the whole active height is about 300 MB when using 40 axial levels in the thermal-hydraulics and 160 axial levels in the neutronics. Therefore, for a typical PWR core with 157 assemblies, the amount of memory needed is about 46 GB which has to be distributed among all the available processors.

Thus, a new scheme will be developed for the alternate dissections where all the computing nodes act, at the same time, as master nodes storing part of the full core problem, and sending and receiving information to and from the rest of processors; and act as slaves nodes computing the pin-by-pin coupled solutions for the given subdomains. In this way, the detailed solution of a full core would be affordable and the message passing load would be also better distributed.

To illustrate the speedup of the parallelized neutronics solver, Figure 3 shows the convergence behavior with the number of processors used to compute the ARI-1 steady-state case at HZP for PWR MOX/UO_2 benchmark in 3D with 8 energy groups. Computational times go from 4.5 hours in one processor to 0.4 hours using 128 processors.

As showed in the figure, the speedup saturates over 128 processors because the time fraction in the parallelized part of the code lowers its importance with increasing number of processors: the unitary benefit of adding one more processor gets lower and lower. The cause for this behavior is of course the non-parallelized part of the code which corresponds with the nodal solver and the message passing from the slaves to the subdomains.

Thus two ways of improving this performance have to be considered, one improving the nodal solver speed, may be using a lower order nodal solution and/or an also parallelized solver, and

other sending less amount of memory from the master to the slaves. However, the amount of time used in the message passing time is very small when compared with the computing time, and the true gain should be in the improvement of the nodal solver performance.

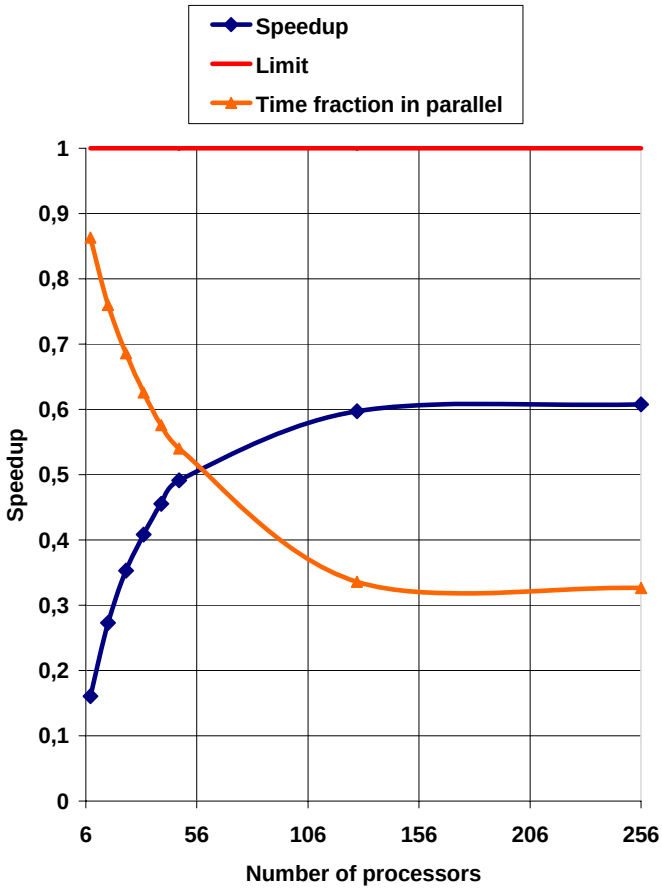

Figure 3. Speedup of the parallelized neutronics solver

4. NUMERICAL VERIFICATION

The OECD/NEA/USNRC PWR MOX/UO_2 Core Transient Benchmark [9] provides homogenized pin level cross sections but without cell discontinuity factors to account for cell heterogeneities, transport and mesh size effects. It is used to compute full core results in comparison to the ones submitted to the final report.

First we present a pin-by-pin solution for the full core at HZP as described in Part 3 of the benchmark specification, where each pin, cladding and surrounding coolant is homogenized in one cell. Table I shows a comparison of some parameters with DeCART transport solution which was taken as reference and the PARCS nodal solution.

Table I. HZP 3D steady state comparison

	DeCART	COBAYA3	PARCS 8G
Critical boron	1265	1313	1334
PWE (%)	ref	1.17	1.20
EWE (%)	ref	2.79	2.85

The values are found to be in between both approaches, and they would be very much improved if the interface discontinuity factors [1] would have been included in the cell library of the benchmark because of the mesh, transport and heterogeneity effects to be corrected. Table II shows the assembly power relative error compared with the DeCART solution and Table III compared with the PARCS nodal solution. The highest error is concentrated in the center of the core around the inserted control rods mainly due to transport effects as the comparison with the nodal solution produces smaller errors.

Table II. HZP 3D steady state assembly power relative error (%) with DeCART

-6,81	-8,42	-2,52	-1,01	1,56	1,15	1,31	-1,18
-8,42	-7,49	-3,65	-0,24	0,80	1,64	-0,04	-2,00
-2,52	-3,65	-0,43	-0,44	0,91	1,34	2,13	-0,32
-1,01	-0,24	-0,44	0,85	0,30	0,62	0,53	-0,91
1,56	0,80	0,91	0,30	2,86	0,22	0,42	
1,15	1,64	1,34	0,62	0,22	0,38	0,63	
1,31	-0,04	2,13	0,53	0,42	0,63		
-1,18	-2,00	-0,32	-0,91				

Table III. HZP 3D steady assembly power relative error (%) with PARCS 8G

1,81	-2,01	1,62	-1,39	-0,42	0,48	0,62	-1,62
-2,01	-0,98	-1,47	-0,54	-0,52	0,13	0,79	-1,17
1,62	-1,47	1,76	-1,33	-0,27	0,75	0,65	-1,72
-1,39	-0,54	-1,33	-0,77	-0,59	0,58	2,12	-0,03
-0,42	-0,52	-0,27	-0,59	3,42	0,25	0,98	
0,48	0,13	0,75	0,58	0,25	2,57	0,87	
0,62	0,79	0,65	2,12	0,98	0,87		
-1,62	-1,17	-1,72	-0,03				

Regarding the coupled calculations, several mini-core configurations have been computed to test the ability of the scheme to converge the pin powers and the temperatures and densities distribution. From all those mini-core cases, we have chosen a full axial length symmetric one with 101 fuel assemblies surrounded with reflector of assembly size as the most representative. Figure 4 shows the rod cladding temperature and the subchannel water density at two thirds of

the axial domain for half of the core. All these mini-core calculations used the cell cross section libraries from the MOX benchmark. This simulation took 36 hours to converge using 40 processors.

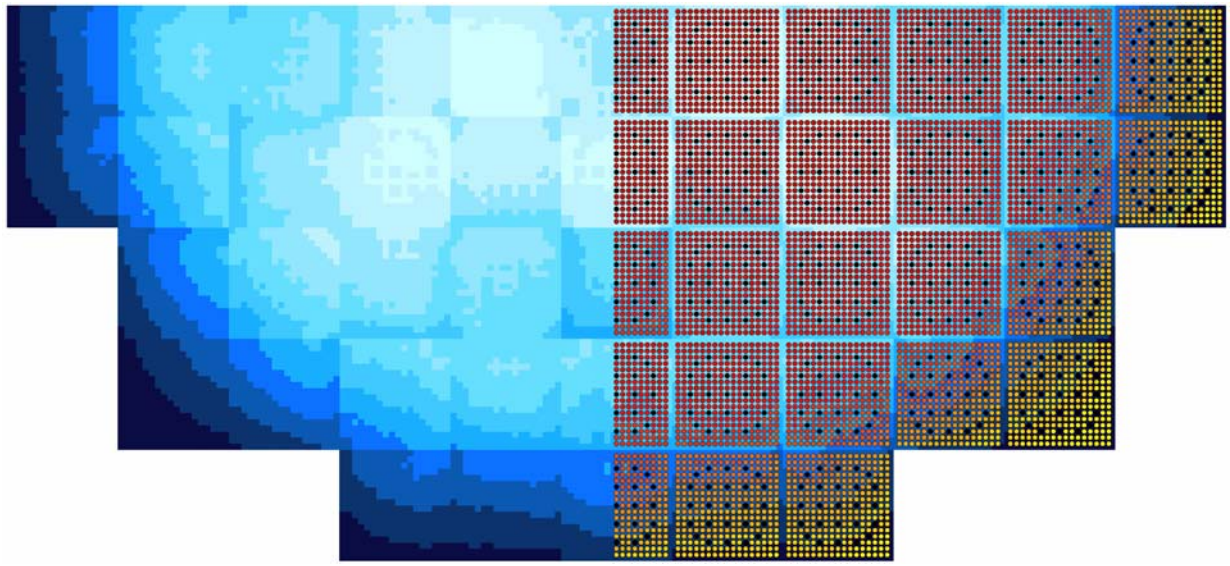

Figure 4. Subchannel water density (left) and cladding temperature (right) for a 101 assemblies mini-core

5. CONCLUSIONS

The efficiency achievable in full core problems by using a parallelized code for the alternate dissections domain decomposition scheme was evaluated. Computing times were notably reduced when using various CPUs simultaneously to the scale of minutes for steady states without thermal-hydraulic feedback, hours for steady states with a thermal-hydraulic feedback, and days for full core transients.

The nodal acceleration introduced in the neutronics solution is a necessary computation to achieve full core convergence and to accelerate the method. However it weakens the scalability of the method for the highest numbers of processors. Improvements in the nodal solver or even its parallelization should be advantageous. A nodal/channel acceleration is also envisioned for transient computations in order to lower computing times.

The introduction of a thermal-hydraulical solution based on quasi-stationary computations expands the computing time of the first steady state but this part of the code holds better scalability properties as a channel acceleration is not used.

A distributed memory architecture where there is not a defined master can be implemented except for the nodal acceleration where it is necessary to store all the detailed information needed to compute a full reactor.

A pin-by-pin transport corrected diffusion computation coupled with an eight equations thermal-hydraulical model yields a level of detail useful to take as a reference calculation for comparison

with coupled nodal-channel codes or to generate nodal cross sections libraries in situations where the thermal-hydraulical distribution within the assemblies could be of importance in off-nominal conditions. Besides, it provides directly more accurate results on the margins at the pin-by-pin scales.

ACKNOWLEDGMENTS

This work is partially funded by the EC Commission under the 6th and 7th EURATOM Framework Programs, within the Integrated and Collaborative Projects NURESIM "European Platform for Nuclear Reactor Simulations" and NURISP "Nuclear Reactor Integrated Simulation Project", under contracts 516560 (FI6O) and 232124 (FI7O). It is also funded by the Spanish Science and Innovation Ministry with the FPU Program for teaching and researching formation under grant AP2005-0667 for the first author, and by the Spanish Consejo de Seguridad Nuclear for the second author.

The authors thankfully acknowledge the computer resources, technical expertise and assistance provided by the Centro de Supercomputación y Visualización de Madrid (CeSVIMA) and the Spanish Supercomputing Network.

REFERENCES

[1]. ARAGONÉS, J.M. and AHNERT, C., 1986. A Linear Discontinuous Finite-Difference Formulation for Synthetic Coarse-Mesh Few-Group Diffusion Calculations. *Nuclear Science and Engineering*, **94**(4), pp. 309-322.

[2]. ARAGONÉS, J.M., AHNERT, C. and GARCÍA-HERRANZ, N., 2005. The Analytic Coarse-Mesh Finite-Difference Method for Multigroup and Multidimensional Diffusion Calculations, *Proceedings of Mathematics and Computations, Supercomputing, Reactor Physics and Nuclear and Biological Applications,* 2005, pp. 194.

[3]. CUERVO, D., AHNERT, C. and ARAGONÉS, J.M., 2005. Analysis of the influence of Subchannel Effects on the PWR Average Channel Calculation, *International Topical Meeting on Nuclear Thermal-Hydraulics NURETH-11*, 3-6 October 2005, pp. 396.

[4]. GARCÍA-HERRANZ, N., CABELLOS, O., ARAGONÉS, J.M. and AHNERT, C., 2003. Analytic coarse-mesh finite-difference method generalized for heterogeneous multidimensional two-group diffusion calculation. *Nuclear Science and Engineering*, **144**(1), pp. 23-35.

[5]. HERRERO, J.J., AHNERT, C. and ARAGONÉS, J.M., 2007. 3D whole core fine mesh multigroup diffusion calculations by domain decomposition through alternate dissections, *Proceedings of Mathematics and Computations and Supercomputing in Nuclear Applications,* 15-19 April 2007.

[6]. HERRERO, J.J., AHNERT, C. and ARAGONÉS, J.M., 2007. Spacial domain decomposition for LWR cores at the pin scale, *Transactions of the American Nuclear Society,* 11-15 November 2007.

[7]. JIMENEZ, J., AVRAMOVA, M., CUERVO, D. and IVANOV, K., 2008. Comparative analysis of neutronics/thermal-hydraulics multi-scale coupling for LWR analysis, *Proceedings of the International Conference on the Physics of Reactors,* 14-19 September 2008.

[8]. JIMENEZ, J. et al., A domain decomposition methodology for pin by pin coupled neutronic and thermal–hydraulic analyses in COBAYA3. *Nucl. Eng. Des.* (2009), doi:10.1016/j.nucengdes.2009.01.012
[9]. KOZLOWSKI, T. and DOWNAR, T.J., 2006. Pressurized Water Reactor MOX/UO2 core transient benchmark. Final Report. NEA/NSC/DOC 20.
[10]. LIONS, P.L., 1989. On the Schwarz Alternating Method II: Stochastic Interpretation and Order Properties, *Domain Decomposition Methods*, 1989, SIAM.
[11]. LIONS, P.L., 1988. On the Schwarz Alternating Method I, *First International Symposium on Domain Decomposition Methods for Partial Differential Equations*, 1988, SIAM pp. 1-42.
[12]. LOZANO, J.A., ARAGONÉS, J.M. and GARCÍA-HERRANZ, N., 2007. Development and performance of the analytic nodal diffusion solver ANDES in multigroups for 3D rectangular geometry, *Proceedings of Mathematics and Computations and Supercomputing in Nuclear Applications,* 15-19 April 2007.