

UKAEA-CCFE-PR(18)27

D. Samaddar, D.P. Coster, X. Bonnin, L.A. Berry, W.R. Elwasif  
and D.B. Batchelor

# **Application of the parareal algorithm to simulations of ELMs in ITER plasma**

Enquiries about copyright and reproduction should in the first instance be addressed to the UKAEA Publications Officer, Culham Science Centre, Building K1/0/83 Abingdon, Oxfordshire, OX14 3DB, UK. The United Kingdom Atomic Energy Authority is the copyright holder.

# Application of the parareal algorithm to simulations of ELMs in ITER plasma

D. Samaddar,<sup>a</sup> D.P. Coster,<sup>b</sup> X. Bonnin,<sup>c</sup> L.A. Berry,<sup>d</sup> W.R. Elwasif<sup>d</sup>  
and D.B. Batchelor

<sup>a</sup>CCFE, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK

<sup>b</sup>Max-Planck-Institut für Plasmaphysik, Germany

<sup>c</sup>ITER Organization, Route de Vinon sur Verdon, 13115 St Paul Lez Durance, France

<sup>d</sup>Oak Ridge National Laboratory, Oak Ridge, USA



# Application of the parareal algorithm to simulations of ELMs in ITER plasma

D. Samaddar<sup>a</sup>, D.P. Coster<sup>b</sup>, X. Bonnin<sup>c</sup>, L.A. Berry<sup>d</sup>, W.R. Elwasif<sup>td</sup>, D.B. Batchelor<sup>d</sup>

<sup>a</sup>*CCFE, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK*

<sup>b</sup>*Max-Planck-Institut für Plasmaphysik, Germany*

<sup>c</sup>*ITER Organization, Route de Vinon sur Verdon, 13115 St Paul Lez Durance, France*

<sup>d</sup>*Oak Ridge National Laboratory, Oak Ridge, USA*

---

## Abstract

This paper explores the application of the parareal algorithm to simulations of ELMs in ITER plasma. The primary focus of this research is identifying the parameters that lead to optimum performance. Since the plasma dynamics vary extremely fast during an ELM cycle, a straightforward application of the algorithm is not possible and a modification to the standard parareal correction is implemented. The size of the time chunks also have an impact on the performance and needs to be optimized. A computational gain of 7.8 is obtained with 48 processors to illustrate that the parareal algorithm can be successfully applied to ELM plasma.

*Keywords:* time parallelization, parareal algorithm, edge plasma, scrape off layer (SOL), SOLPS, magnetically confined plasmas, Edge Localized Modes or ELMs, ITER, tokamak

*PACS:* 02.70.-c, 52.40.Hf, 52.55.Fa, 52.65.-y, 02.90.+p, 07.05.Tp

---

## 1. Introduction

The ever increasing demand for improving complexities in simulations requires maximizing the efficient use of computational resources. Traditional parallelization techniques (such as space parallelization) although reduce the wallclock time, often reach saturation on modern supercomputing machines.

---

*Email address:* [Debasmita.Samaddar@ukaea.uk](mailto:Debasmita.Samaddar@ukaea.uk) (D. Samaddar)

Parallelizing the time domain ushers in a new possibility of optimizing resource utilization. Parallel in time algorithms are not a replacement for but are complimentary to other schemes of parallelization. This therefore allows utilization of the gain achieved from existing parallelization and adding a large improvement in computational gain.

Various parallel in time algorithms exist such as PITA [1], parareal [2], RIDC [3, 4], PFASST [5]. This work explores the application of the Parareal algorithm to a complex, non-linear simulation in fusion plasma. Of all the algorithms that achieve temporal parallelization, the Parareal algorithm has been the one to have been used the most in non-linear plasma physics simulations over the years [6, 7, 8]. The algorithm was introduced in [2] and its convergence and stability have been analyzed in great detail in [9, 10, 11, 12, 13, 14, 15, 16, 17].

This work seeks to explore temporal parallelization of simulations of Edge Localized Modes (ELMs) that play a crucial role in the successful operation of any fusion device [18, 19, 20, 21, 22]. ELMs are simulated using the multi-fluid codes SOLPS 5.0 [23, 24] and SOLPS-ITER [25, 26]. ELMs have been observed in experiments on present machines[27, 28]. These modes exhibit irregular or quasi-regular periodicity and lead to deposition of huge power fluxes on plasma facing walls . While ELMs are an area of strong active research in the fusion community and there is significant information collected from experimental databases, a complete understanding of ELMs is still elusive to scientists.

One major hindrance in studying the behavior of ELMs using computer simulations is their demand for extremely large wall clock times. This is why adding a new dimension to the parallelization is of significance in ELM research. A successful application of the Parareal algorithm makes these simulations much more feasible, allowing the inclusion of more complex physics while still maintaining a reasonable computation time. However, due to the strong non-linear phenomena involved, the application of temporal parallelization to simulations of ELMs is unsurprisingly far from straightforward and this work attempts to identify the computational parameters and regimes that allow the best performance.

The present research is a follow-up of the work detailed in [8]. [8] demonstrated that the parareal algorithm is applicable to edge plasma simulations but relatively 'steady' states were used as test beds. This work extends that application to more strongly coupled non-linear phenomena called ELMs. Due to the complex non-linear behavior of ELMs, the application poses new

challenges. A modification to the application of the parareal correction allows convergence and generates computational gain.

This paper is organized as follows: ELMs are described in sec.2 and sec.3 details how they are implemented in the code to be observed in the simulations. Sec.4 describes the Parareal Algorithm and the results are described in sec.5.

## 2. ELMs

Edge Localized Modes(ELMs) [29, 30] are periodic bursts of instabilities at the plasma edge in a tokamak. A tokamak is a fusion device of toroidal geometry with complex magnetic field lines to confine the burning plasma of a few hundred million Kelvin within the vessel. As has been mentioned in [8], simulations of the plasma edge is particularly challenging due to the interactions with the materials of the walls of the vessel accompanied by both radial (perpendicular to the magnetic field lines) and parallel (parallel to the field lines) transports.

It is desirable to operate a tokamak in the H-mode (High Confinement Mode) [18, 31, 32, 33] since that greatly enhances the energy confinement time of the plasma. The energy confinement time  $\tau$  is the characteristic time scale in which the energy escapes the plasma. In fact, H-mode confinement time is typically twice that of L-mode (Low Confinement mode). The sharp transition from L-mode to H-mode occurs when the input heating power is above a certain threshold. A sharp increase in the profile gradients (such as of pressure, density or temperature) occurs over a very narrow region (typically a few cm) at the plasma edge. This results in the formation of the 'edge transport barrier' or 'H-mode pedestal' which generates the improved confinement. The suppression of turbulent transport at the plasma edge is believed to be the reason behind the improved confinement - although a complete understanding of the mechanisms at play is still an area of active research.

While the improved confinement makes the H-mode an attractive operating regime, the presence of instabilities makes it more challenging. Two MHD instabilities become prominent. The increasing pressure gradient leads to ballooning instabilities and the gradient in the edge current generates the peeling instability. These two instabilities are believed to be responsible for ELMs [18, 19].

ELMs are characterized by sudden bursts of energy and a transient rise of heat loads on the divertor targets of the device. Depending on their amplitudes and frequencies, ELMs are classified into types I (large amplitudes - hence often termed as 'giant' ELMs) and II and III (smaller, hence often known as 'grassy' ELMs). Current predictions estimate [21, 34, 35] the heat flux due to ELMs on the divertor plates at ITER (the world's biggest experimental tokamak) will be 20 times larger than what can be tolerated for a reasonable lifetime of the target materials. This makes research into ELM control and mitigation of primary importance for making fusion a viable source of energy.

Despite the devastating impact of tremendous heat loads on plasma facing components and leading to temporary degradation in confinement, ELMs have strong positive consequences. They help to control plasma densities and remove impurities from the plasma core. As a result, controlled ELMs in H-mode is the desired operational mode for ITER.

### 3. ELMs in the SOLPS code

The SOLPS 5.0 and SOLPS-ITER codes are described in detail in [23, 24, 25, 26]. SOLPS stands for Scrape Off Layer Plasma Simulator which, as the name suggests, simulates the Scrape Off Layer (SOL). The SOL is the region between the wall of the device and the Last Closed Flux Surface (LCFS), characterized by open field lines. Both codes solve the same set of equations and SOLPS-ITER is considered to be an upgraded version of SOLPS5.0. Both codes use the B2.5 package that solve the Braginskii fluid equations for multiple species or all charge states of every individual element present in the plasma [20]. The regimes studied are sufficiently collisional to justify the use of fluid equations. In the present case, the fluid model is used for both the charged species as well as the neutrals in the plasma as in [20]. Ideally, it is desirable to use a kinetic model to simulate the neutral species in the plasma, but using the fluid model makes the computations tractable.

It is interesting to note that ELMs are believed to be intrinsically MHD phenomena while SOLPS is not an MHD code. It must be clarified that a code like SOLPS is used to study the impact of ELMs (and provide insight into the effects of individual species which is impossible to perform in an MHD code) rather than the mechanisms for their generation (which is typically studied using MHD codes [36]).



ELMs characteristically enhance the radial transport in a plasma [37]. This is typically simulated by increasing the diffusive or convective coefficient for a brief period (say 1 ms) at regular intervals. The impact of that increase successfully lasts much longer than 1 ms and simulates the behaviour of an ELM-plasma. These simulations have been successfully bench marked against experiments in the JET and TCV tokamaks [21, 22].

[8] studied relatively simpler cases in the tokamaks MAST and DIII-D. The present work simulates a significantly larger device - ITER - with a much higher number of species mix (98) in the plasma. The physics involved in these simulations are very similar to the cases described in [20]. The species include all ions and neutrals of Deuterium, Tritium, Helium, Beryllium, Neon and Tungsten and the complex interactions between the various charge states are included in the simulations. As in [20], prompt redeposition of Tungsten, which is otherwise released into the plasma as a result of sputtering at the divertor targets, is included in the calculations. The density gradient driven diffusivity has been increased by 10 times for 1 ms in every 20 ms to simulate the ELM triggered plasma. This is illustrated in Fig.1 where the diffusivity along the outer midplane is plotted for two different times. The ELM affected area on the outer midplane is estimated to be 40 cm for ITER.

A computational grid of size  $96 \times 36$  as shown in Fig.2 was used as the fine model for the SOLPS computations of ELMs in the ITER plasma. The typical size of the timesteps for the fine computations was  $1E - 5$  sec. A typical ELM cycle is shown in Fig.3 where the electron density (in  $m^{-3}$ ) at the separatrix on the midplane (henceforth will be referred to as *nsep*) and the maximum total power flux (in  $W.m^{-2}$ ) (henceforth referred to as *pwxap*) on the outer target are plotted against time. It is seen that both these quantities are greatly increased during the ELM and then reduce from the peak values during the ELM-crash phase.

#### 4. Parareal algorithm

The parareal algorithm was first introduced in [2]. It is described in detail in [8]. The algorithm uses a predictor-corrector approach with an accurate fine (F) integrator and another coarse integrator (G). Identifying the optimum coarse predictor is generally the biggest challenge in all complex applications. The parareal correction is given by eq.1.

$$\lambda_{i+1}^k = F(\lambda_i^{k-1}) + G(\lambda_i^k) - G(\lambda_i^{k-1}) \quad (1)$$




Figure 1: The particle density gradient driven radial diffusivity is increased by 10 times to simulate ELMs.






Figure 2: A grid of size  $96 \times 36$  was used for the fine or serial computation of the ELM plasma in the ITER device.



(a) The electron density (nsepm) at the separatrix on the outer midplane computed shows a sudden rise followed by a reduction during the ELM cycle.



(b) The maximum total power flux on the <sup>8</sup> outer target for the same case as Fig.3a.

Figure 3: A typical ELM cycle using the SOLPS5.0 code package with the Fine solver is demonstrated.

$\lambda_{i+1}^k$  is the initial state for the  $(i+1)^{th}$  time slice at the  $k^{th}$  iteration. At the  $k^{th}$  iteration  $\lambda_i^k$  is evolved to  $\lambda_{i+1}^k$  using  $F(\lambda_i^k)$  and  $G(\lambda_i^k)$ .

Just as in [8], this paper also utilizes the event-based approach given in detail in [38, 39, 40].

The coarse and fine calculations are repeated across a number of parareal iterations  $k$  until convergence is achieved. Parareal convergence across a time chunk or processor is said to have been achieved when the 'defect' between two successive fine calculations ( $k$  and  $k-1$ ) across that processor is below a tolerance value. The sum of this defect in solutions across a time chunk between  $t_i$  and  $t_{i-1}$  is defined by

$$\zeta_i^k = \int_{t_{i-1}}^{t_i} \frac{|\lambda^k(t) - \lambda^{k-1}(t)|}{|\lambda^{k-1}(t)|} dt. \quad (2)$$

The solution is then converged for time slice  $i$  if,

$$\zeta_i^k < \text{tol}, \quad \forall i \leq I. \quad (3)$$

For the purposes of parareal convergence, eq.2 is applied to two variables `pwmxap` and `pwmxip`. As already mentioned, `pwmxap` is the maximum total power flux (in  $W.m^{-2}$ ) on the outer divertor. `pwmxip` is the same for the inner divertor.

## 5. Results

The cases explored in this work used both the SOLPS5.0 and SOLPS-ITER codes. In some cases, results from SOLPS-ITER were more stable than those from SOLPS5.0. The value of `tol` in eq.3 was 0.005. The choice of this value is explained in [8].

Following the work in [8], a reduced grid model of size  $48 \times 36$  with bigger timesteps was used as a coarse predictor for the Parareal implementation. While this approach was relatively straightforward for the cases explored in [8], the treatment turned out to be much more challenging for a plasma with ELMs. The parareal solutions had spurious rises and falls in the plasma quantities such as temperature and density throughout the ELM cycles. One such 'catastrophic' case is illustrated in Fig.4 where there is a rise in `nsepmp`, the electron density at the separatrix on the midplane (it must be noted that other quantities like temperature and ion densities also behave similarly) when there is expected to be a decrease as is seen in Fig.3a. This was a

simulation using 16 processors, with  $ntim_F = 100$  and  $dt_G = 10 * dt_F$ .  $ntim$  is the number of time-steps solved per processor and  $dt$  represents the size of each time step. The subscripts F and G represent the fine and coarse cases, respectively. It may be noted that  $ntim_F * dt_F = ntim_G * dt_G$  is the size of each 'time chunk' solved on every processor.

A closer look at the parareal computation shows that in this case, the unwanted ELM rise starts at the 14<sup>th</sup> time chunk, which subsequently affects successive time chunks solved by other processors. Plotting  $nesepm$  against time for both  $k= 1$  and  $k= 2$  in Fig.5 showed that the rise occurs at the parareal iteration  $k= 2$ , which somewhat implies that the parareal correction as given by eq.1 might be the cause for this behaviour.

A possible explanation might be obtained if we take into account that the radial profiles and the subsequent plasma dynamics vary extremely rapidly during an ELM cycle. This is illustrated for the electron density and temperature along the midplane in Figs.6 & 7. With the fact that eq.1 is an algebraic correction at all grid points, it is not very surprising that after applying eq.1, the plasma profiles change to an ELM rise phase while it is expected to be in an ELM crash phase or vice versa.

This issue due to the rapidly varying profiles during an ELM is a cause for concern since the success of the parareal algorithm depends on eq.1. However, implementing a restriction to the parareal correction appeared to solve the problem. At the  $k^{th}$  parareal iteration, the value of each primary variable was not allowed to vary more than that of the fine value computed at the  $(k - 1)^{th}$  iteration by  $X\%$ . To ensure that the study was restricted to the correction to the grid points and to rule out any impact of choosing a  $dt_G$  greater than  $dt_F$ ,  $dt_G = dt_F = 1E - 5s$  was chosen with a grid of  $48 \times 36$  for the coarse predictor. A series of simulations were performed with different values of  $X = 1, 5, 20, 40$ . While the undesirable behaviour in the solutions was eradicated with a decreasing value of  $X$ , the number of iterations required for parareal convergence was minimum at  $X = 1$ .

With this promise, the next numerical parameter that was explored was the size of the time chunk solved per processor. This involved varying  $ntim_F$  and subsequently varying  $ntim_G$ , keeping  $dt_F$  and  $dt_G$  constant. The size of the time chunk has been a factor influencing parareal performances in a large number of previous applications [6, 8] and was found to be a strong one in the present case. Once again with restricting the time step sizes such that  $dt_G = dt_F = 1E - 5s$ , a series of simulations were performed with 8 processors or time chunks varying the values of  $ntim$ . The results are

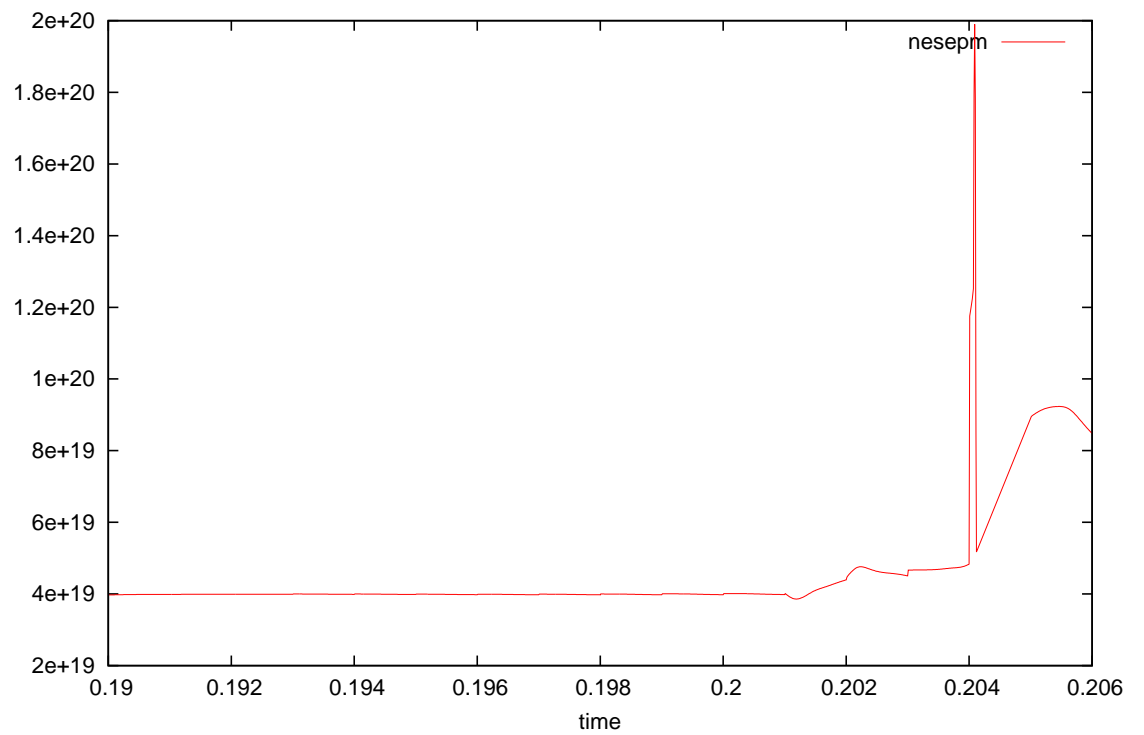




Figure 4: A spurious rise in the electron density at the outer midplane obtained from the parareal computation using a grid of  $48 \times 36$  and  $dt_G = 10 * dt_F$  as the coarse solver.



(a) nesepm is expected to decrease during the time 0.203 and 0.204 sec. This behaviour is evident in the parareal computation at  $k=1$ .



(b) At the parareal iteration  $k=2$ , <sup>12</sup>nesepm starts to increase affecting subsequent time chunks solved using the parareal algorithm eventually leading to the (wrong) solution shown in Fig.4.

Figure 5: The parareal solutions are very sensitive to initial values in case of ELMs.






Figure 6: The radial profile of the electron density,  $n_e$  very rapidly varies during an ELM cycle. The distance is in metres.




Figure 7: The radial profile of the electron temperature,  $t_e$  very rapidly varies during an ELM cycle. The distance is in metres.

listed in Table 1. It was observed that  $ntim = 200$  allowed the best parareal performance. This dependence on  $ntim$  has been observed many times but a full mathematical understanding of the parareal algorithm with respect to it is still unclear. It may be argued that with a large  $ntim$ , the coarse solution deviates sufficiently such that the parareal correction can no longer recover the 'correct' solution. Interestingly, with  $dtim = 1E - 5$ ,  $ntim = 2000$  actually is the time that includes an entire ELM, that is an ELM rise and crash phases. One might therefore expect, that parareal should be easier to implement when  $ntim = 2000$ , as most of the radical changes in the plasma dynamics are confined within a time chunk, and the breaks in the time line occur at 'quiet phases' in an ELM cycle. However, it was found that parareal convergence was not at all achieved with chunk size,  $ntim = 2000$ . This could indicate that the plasma dynamics were still varying rapidly during the 'quiet phase' and with such a long time chunk, the coarse predictor was deviating enough to hinder parareal convergence.

Table 1: The number of iterations were minimum at  $ntim_F = ntim_G = 200$  with parareal correction restricted to  $X = 1$ .

<b>ntim</b>	<b>Parareal iteration k</b>
20	4
200	2
300	3
400	No convergence
2000	No convergence

It is believed that SOLPS-like simulations with  $dt = 1E - 5s$  marginally resolve ELM physics with very rapidly varying plasma dynamics. A step size of  $1E - 6s$  is desired. But a serial computation of a single ELM cycle would then take about 16.25 days! However, once the optimum parameters ( $ntim_F = 200$  &  $X = 1$ ) were identified, a set of parareal simulations were performed with  $dt_F = 1E - 6s$  &  $dt_G = 1E - 5s$ . With this coarsening of the predictor with respect to the fine integrator resulted in an increase in parareal iteration ( $k$ ) as in Table 1. The case of  $ntim_F = 200$  &  $X = 1$  now converged at  $k = 3$ . With the serial and parareal wallclock times defined as  $T_{ser}$  and  $T_{PR}$  respectively, the computational gain is defined as  $gain = \frac{T_{ser}}{T_{PR}}$ . A computational gain of 7.8 was achieved with 48 processors, which is expected to rise with increasing processor counts.

Before analysing the parareal solutions, it may be noted that very slight changes to the integrator or the code itself can result in slight variations in the solutions. This is evident in the results obtained from SOLPS 5.0 and SOLPS-ITER. Although both versions of the code solve the same sets of equations, the solutions (pwmxip and pwmxap) over a complete ELM cycle (2ms) are illustrated in Fig.8. The Root Mean Squared Errors (RMSE) for these two datasets are computed. RMSE is defined by the formula given in eq.4 and is normalized by eq.5, where  $N$  is the total number of datapoints,  $v_n$  and  $u_n$  are the  $n^{th}$  variables of SOLPS-ITER and SOLPS 5.0 respectively. Comparing solutions from SOLPS 5.0 and SOLPS-ITER, the normalized RMSE for pwmxip is 1.557E-02 and for pwmxap it is 2.0203E-02.

$$RMSE = \sqrt{\frac{\sum_{n=1}^N (v_n - u_n)^2}{N}} \quad (4)$$


$$NormalizedRMSE = \frac{RMSE}{u_{max} - u_{min}} \quad (5)$$

A similar comparison is performed using the parareal solutions with respect to the corresponding SOLPS 5.0 solutions in 9. The normalized RMSE for the parareal computation was determined to be 1.95E-02 for pwmxip and 6.43E-03 for pwmxap.


## 6. Conclusion

The parareal algorithm is shown to work for a complex case of ELM simulations in ITER plasma with 98 species. The rapidly changing radial profiles during an ELM cycle pose unique challenges for this application. A modification to the application of the parareal algorithm alleviates the problem. The parareal performance is also found to be sensitive to the size of the time chunk solved per processor. An optimum value for it is identified, and a simulation with 48 processors yielded a computational gain of 7.8.

This application illustrates that ELM simulations can become more tractable using the parareal algorithm. As a result more complex physics can be incorporated into the model, such as kinetic neutrals, and long simulations of multiple ELM cycles may be performed within much shorter wall clock time.




(a) The maximum total power flux on the inner target for an ELM cycle, obtained from SOLPS 5.0 and SOLPS-ITER. The normalized RMSE defined by eq.5 for these data sets is  $1.557 \times 10^{-2}$ .




(b) The maximum total power flux on the outer target for the same case as Fig.8a. The normalized RMSE for these data sets is  $2.0203 \times 10^{-2}$ .

Figure 8: The solutions obtained from serial versions of SOLPS 5.0 and SOLPS-ITER are compared to one another.



(a) The maximum total power flux on the inner target for an ELM cycle, obtained from SOLPS 5.0 and parareal. The normalized RMSE defined by eq.5 for these data sets is  $1.95E-02$ .



(b) The maximum total power flux on the outer target for the same case as Fig.9a. The normalized RMSE for these data sets is  $6.43E-03$ .

Figure 9: The solutions obtained from serial versions of SOLPS 5.0 and the parareal computation are compared. The error bars represent the maximum absolute error for each variable (pwmxi and pwmxi) between SOLPS 5.0 and SOLPS-ITER calculations using a serial processor.

## 7. Acknowledgements

This work has been funded within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. This work has been part-funded by the RCUK Energy Programme [grant number EP/I501045]. To obtain further information on the data and models underlying this paper please contact PublicationsManager@ccfe.ac.uk. Also, the views and opinions expressed herein do not necessarily reflect those of the ITER Organization. The computations have been performed on the EUROfusion Gateway Cluster.

## References

- [1] C. Farhat, M. Chandesris, Time-decomposed parallel time-integrators: theory and feasibility studies for fluid, structure, and fluid-structure applications., *Int. J. Numer. Meth. Engng.* 58 (9) (2003) 1397.
- [2] J. Lions, Y. Maday, G. Turinici, A "parareal" in time discretization of PDE's, *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics* 332 (7) (2001) 661–668.
- [3] A. J. Christlieb, B. W. Ong, J. M. Qiu, Integral deferred correction methods constructed with high order Runge-Kutta Methods, *Math. Comp.* 79 (2009) 761–783.
- [4] A. Christlieb, C. Macdonald, B. Ong, Parallel high-order integrators, *SIAM J. Sci. Comput.* 32 (2) (2010) 818–835.
- [5] M. Emmett, M. L. Minion, Toward an efficient parallel in time method for partial differential equations, *Communications in Applied Mathematics and Computational Science* 7 (1) (2012) 105.
- [6] D. Samaddar, D. E. Newman, R. Sánchez, Parallelization in time of numerical simulations of fully-developed plasma turbulence using the parareal algorithm, *Journal of Computational Physics* 228 (2010) 6558.

- [7] D. Samaddar, T. Casper, S. Kim, L. A. Berry, W. R. Elwasif, D. B. Batchelor, W. Houlberg, Application of the parareal algorithm to advanced operation scenario simulations of ITER plasma using the CORSICA code, *Applied Mathematics and Computation* (submitted).
- [8] D. Samaddar, D. Coster, X. Bonnin, C. Bergmeister, E. Havlíčková, L. A. Berry, W. R. Elwasif, D. B. Batchelor, Temporal parallelization of edge plasma simulations using the parareal algorithm and the solps code, *Computer Physics Communications* 221 (2017) 19–27.
- [9] L. Baffico, S. Bernard, Y. Maday, G. Turinici, G. Zérah, Parallel in time molecular dynamics simulations, *Phys. Rev. E* 66 (5) (2002) 057706.
- [10] G. Bal, Y. Maday, A parareal time discretization for non-linear PDEs with application to the pricing of an American put, Vol. 23 of *Lect. Notes Comput. Sci. Eng.*, Springer, Berlin, 2002, p. 189202.
- [11] G. Bal, Parallelization in time of (stochastic) ordinary differential equations (2003).  
URL <http://www.columbia.edu/~gb2030/PAPERS/parallelttime.pdf>
- [12] P. F. Fischer, F. Hecht, Y. Maday, A parareal in time semi-implicit approximation of the Navier-Stokes equations, in: *Proceedings of Fifteen International Conference on Domain Decomposition Methods*, Vol. 40, Springer Verlag, 2004, pp. 433–440.
- [13] M. J. Gander, E. Hairer, Nonlinear convergence analysis for the parareal algorithm, in: *Domain Decomposition Methods in Science and Engineering XVII*, Vol. 60, Springer, 2008, p. 45.
- [14] I. Garrido, M. S. Espedal, G. E. Fladmark, A convergent algorithm for time parallelization applied to reservoir simulation, in: *Proceedings of Fifteen International Conference on Domain Decomposition Methods*, Vol. 40, Springer, Berlin, 2004, p. 469.
- [15] G. Staff, G. Rønquist, Stability of the parareal algorithm, in: *Proceedings of Fifteen International Conference on Domain Decomposition Methods*, Springer Verlag, 2003, pp. 449–456.



- [16] J. Reynolds-Barredo, D. E. Newman, R. Sánchez, D. Samaddar, L. A. Berry, W. R. Elwasif, Mechanisms for the convergence of time-parallelized, parareal turbulent plasma simulations, *Journal of Computational Physics* 231 (2012) 7851.
- [17] A.-M. Baudron, J.-J. Lautard, Y. Maday, M. K. Riahi, J. Salomon, Parareal in time 3D numerical solver for the LWR Benchmark neutron diffusion transient model, *Journal of Computational Physics* (submitted).
- [18] J. Connor, Edge-localized modes-physics and theory, *Plasma physics and controlled fusion* 40 (5) (1998) 531.
- [19] J. Connor, R. Hastie, H. Wilson, R. Miller, Magnetohydrodynamic stability of tokamak edge plasmas, *Physics of Plasmas* 5 (7) (1998) 2687–2700.
- [20] D. Coster, A. Chankin, H.-J. Klingshirn, R. Dux, E. Fable, X. Bonnin, A. Kukushkin, A. Loarte, Solps modelling of w arising from repetitive mitigated elms in iter, *Journal of Nuclear Materials* 463 (2015) 620–623.
- [21] B. Gulejová, R. Pitts, M. Wischmeier, R. Behn, D. Coster, J. Horacek, J. Marki, Solps5 modelling of the type iii elming h-mode on tcv, *Journal of nuclear materials* 363 (2007) 1037–1043.
- [22] B. Gulejová, R. Pitts, D. Coster, X. Bonnin, M. Beurskens, S. Jachmich, A. Kallenbach, et al., Solps5 simulations of type i elming h-mode at jet, *Journal of Nuclear Materials* 390 (2009) 412–416.
- [23] R. Schneider, X. Bonnin, K. Borrass, D. Coster, H. Kastelewicz, D. Reiter, V. Rozhansky, B. Braams, Plasma edge physics with b2-eirene, *Contributions to Plasma Physics* 46 (1-2) (2006) 3–191.
- [24] X. Bonnin, A. Kukushkin, D. Coster, Code development for iter edge modelling–solps5. 1, *Journal of Nuclear Materials* 390 (2009) 274–277.
- [25] X. Bonnin, W. Dekeyser, R. Pitts, D. Coster, S. Voskoboynikov, S. Wiesen, Presentation of the new solps-iter code package for tokamak plasma edge modelling, *Plasma and Fusion Research* 11 (2016) 1403102–1403102.

- [26] S. Wiesen, D. Reiter, V. Kotov, M. Baelmans, W. Dekeyser, A. Kukushkin, S. Lisgo, R. Pitts, V. Rozhansky, G. Saibene, et al., The new solps-iter code package, *Journal of nuclear materials* 463 (2015) 480–484.
- [27] Resonant magnetic perturbation experiments on mast using external and internal coils for elm control.
- [28] O. Gruber, A. Kallenbach, M. Kaufmann, K. Lackner, V. Mertens, J. Neuhauser, F. Ryter, H. Zohm, M. Bessenrodt-Weberpals, K. Büchl, et al., Observation of continuous divertor detachment in h-mode discharges in asdex upgrade, *Physical review letters* 74 (21) (1995) 4217.
- [29] J. Wesson, Vol. 149, Oxford University Press, 2011.
- [30] P. C. Stangeby, *The plasma boundary of magnetic fusion devices*, Vol. 7, Taylor & Francis, 2000.
- [31] F. Wagner, G. Fussmann, T. Grave, M. Keilhacker, M. Kornherr, K. Lackner, K. McCormick, E. Müller, A. Stäbler, G. Becker, et al., Development of an edge transport barrier at the h-mode transition of asdex, *Physical Review Letters* 53 (15) (1984) 1453.
- [32] M. Keilhacker, H-mode confinement in tokamaks, *Plasma Physics and Controlled Fusion* 29 (10A) (1987) 1401.
- [33] F. Wagner, A quarter-century of h-mode studies, *Plasma Physics and Controlled Fusion* 49 (12B) (2007) B1.
- [34] J. Gunn, S. Carpentier-Chouchana, F. Escourbiac, T. Hirai, S. Panayotis, R. Pitts, Y. Corre, R. Dejarnac, M. Firdaouss, M. Kočan, et al., Surface heat loads on the iter divertor vertical targets, *Nuclear Fusion* 57 (4) (2017) 046025.
- [35] T. Eich, B. Sieglin, A. Thornton, M. Faitsch, A. Kirk, A. Herrmann, W. Suttrop, et al., Elm divertor peak energy fluence scaling to iter with data from jet, mast and asdex upgrade, *Nuclear Materials and Energy* 12 (2017) 84–90.
- [36] G. Huysmans, O. Czarny, Mhd stability in x-point geometry: simulation of elms, *Nuclear fusion* 47 (7) (2007) 659.

- [37] A. Loarte, G. Saibene, R. Sartori, D. Campbell, M. Becoulet, L. Horton, T. Eich, A. Herrmann, G. Matthews, N. Asakura, et al., Characteristics of type i elm energy and particle losses in existing devices and their extrapolation to iter, *Plasma Physics and Controlled Fusion* 45 (9) (2003) 1549.
- [38] W. Elwasif, D. Bernholdt, A. Shet, S. Foley, R. Bramley, D. B. Batchelor, L. A. Berry, The Design and Implementation of the SWIM Integrated Plasma Simulator, in: *Distributed and Network-based Processing (PDP)*, Pisa, Italy, 2010.
- [39] W. R. Elwasif, S. Foley, D. Bernholdt, L. A. Berry, D. Samaddar, D. E. Newman, R. Sanchez, A dependency-driven formulation of parareal: parallel-in-time solution of PDEs as a many-task application, in: *MTAGS '11 Proceedings of the 2011 ACM international workshop on Many task computing on grids and supercomputers*, ACM New York, NY, USA, 2011, p. 15.
- [40] L. A. Berry, W. R. Elwasif, J. Reynolds-Barredo, D. Samaddar, R. Sanchez, D. E. Newman, Event-based parareal: A data-flow based implementation of parareal, *Journal of Comp. Phys.* 231 (17) (2012) 5945.