Parallel-in-Time Integration of Transient Phenomena in No-Insulation Superconducting Coils Using Parareal

Erik Schnaubelt[1,2] Mariusz Wozniak[1] Julien Dular[1]
Idoia Cortes Garcia[3] Arjan Verweij[1] Sebastian Schöps[2]

[1]CERN, Meyrin, Switzerland
{erik.schnaubelt, mariusz.wozniak, julien.dular, arjan.verweij}@cern.ch

[2]Graduate School CE at TU Darmstadt, Darmstadt, Germany
sebastian.schoeps@tu-darmstadt.de

[3]Eindhoven University of Technology, Eindhoven, Netherlands
i.cortes.garcia@tue.nl

Abstract

High-temperature superconductors (HTS) have the potential to enable magnetic fields beyond the current limits of low-temperature superconductors in applications like accelerator magnets. However, the design of HTS-based magnets requires computationally demanding transient multiphysics simulations with highly non-linear material properties. To reduce the solution time, we propose using Parareal (PR) for parallel-in-time magneto-thermal simulation of magnets based on HTS, particularly, no-insulation coils without turn-to-turn insulation. We propose extending the classical PR method to automatically find a time partitioning using a first coarse adaptive propagator. The proposed PR method is shown to reduce the computing time when fine engineering tolerances are required despite the highly nonlinear character of the problem. The full software stack used is open-source.

1 Introduction

High-temperature superconductors (HTS) for particle accelerator magnets have the potential to exceed the magnetic field limits of low-temperature superconductors. In particular, no-insulation (NI) pancake coils which are coils wound without turn-to-turn (T2T) electrical insulation are of interest due to their increased electrical and thermal stability [7]. Numerical tools like the finite element (FE) method are crucial to understanding the complex transient behavior of such HTS coils to avoid damage and malfunction.

These simulations of transients and specifically quenches, where the coils locally lose their superconducting properties, are multi-physics problems involving coupled electromagnetism, thermodynamics, and, eventually, solid mechanics. A quench is a local, transient phenomenon that generally requires three-dimensional simulation in time with a high number of degrees of freedom. Furthermore, as material properties must be considered over a wide temperature range, the problem is highly non-linear. Therefore, quench simulations are computationally expensive, and parallelization methods in space and time are desirable for reducing the computation time. This paper describes the use of Parareal (PR) [8] for the parallel-in-time integration of coupled magneto-thermal transient phenomena in NI HTS pancake coils.

Section 2 defines the used FE formulation, and Section 3 discusses the PR method, including the automatic time partitioning. Section 4 introduces the model problem of an NI HTS pancake coil and summarizes the results obtained using the FE formulation with PR time integration. The conclusion is presented in Section 5.

2 Weak Magneto-Thermal Formulation

A magneto-thermal problem is considered on a computational domain Ω, consisting of a conducting subdomain Ω_c and an insulating subdomain Ω_i. Denoting by $(X,Y)_\Omega$, resp. $(\vec{X},\vec{Y})_\Omega$, the integral of the product XY, resp. $\vec{X}\cdot\vec{Y}$, over Ω, the weak form reads: From an initial solution at time $t=t_0$, find the magnetic field strength $\vec{H}\in H_{\phi,I}(\mathrm{curl},\Omega)$ and temperature $T\in H^1_g(\Omega_\mathrm{c})$ for $t\in(t_0,t_N]$ such that [10]

$$\begin{aligned}(C_\mathrm{V}\,\partial_t T, T')_{\Omega_\mathrm{c}} + (\kappa\nabla T,\nabla T')_{\Omega_\mathrm{c}} &= \left(\rho\,\vec{J}\cdot\vec{J},T'\right)_{\Omega_\mathrm{c}} \quad \forall T'\in H^1_0(\Omega_\mathrm{c})\,,\\ \left(\partial_t(\mu\vec{H}),\vec{H}'\right)_\Omega + \left(\rho\nabla\times\vec{H},\nabla\times\vec{H}'\right)_{\Omega_\mathrm{c}} &= 0 \quad \forall\vec{H}'\in H_{\phi,0}(\mathrm{curl},\Omega)\,.\end{aligned}$$

Herein, $\kappa(\|B\|,T)$ is the thermal conductivity, $C_\mathrm{V}(T)$ the volumetric heat capacity, $\rho(\|B\|,T)$ the electric resistivity, and μ the magnetic permeability. In particular, the resistivity of the HTS is given by the power law

$$\rho_\mathrm{HTS} = \frac{E_\mathrm{c}}{J_\mathrm{c}}\left(\frac{\|\vec{J}_\mathrm{HTS}\|}{J_\mathrm{c}}\right)^{n-1},$$

where E_c is the critical electric field, $J_\mathrm{c}(\vec{B},T)$ the critical current density with $\vec{B}=\mu\vec{H}$, n the power law index and $\vec{J}_\mathrm{HTS}$ the current density in the HTS layer. The function space $H^1_g(\Omega_\mathrm{c})$ is a subspace of $H^1(\Omega_\mathrm{c})$ with Dirichlet condition g enforced on parts of the boundary; $H^1_0(\Omega_\mathrm{c})$ is the special case for $g=0$. The space $H_{\phi,I}(\mathrm{curl},\Omega)$ is a sub-space of $H(\mathrm{curl},\Omega)$ with vanishing curl in Ω_i and strongly imposed source currents via cohomology basis functions [9]. Its subspace with zero current is $H_{\phi,0}(\mathrm{curl},\Omega)$. An $\vec{H}-\phi$ formulation is chosen since a resistivity-based characterization for HTS is an efficient and robust choice for systems without ferromagnetic materials [2]. The T2T electrical [13] and thermal [11] contact resistance is treated using a magneto-thermal thin shell approximation [10]. The HTS coated conductor is homogenized, leading to anisotropic κ and ρ. All further details are found in [10].

3 Parareal Method with Automatic Partitioning of Time

The discretization in space of the strongly coupled weak formulation using lowest order finite elements yields a monolithic system of the form

$$\mathbf{M}\left(\mathbf{u}(t)\right)\mathrm{d}_t\mathbf{u}(t) + \mathbf{K}\left(\mathbf{u}(t)\right)\mathbf{u}(t) = \mathbf{f}\left(\mathbf{u}(t)\right). \tag{1}$$

To use PR, we define a fine propagator $\mathbf{u}_i = \mathcal{F}(t_i, t_0, \mathbf{u}_0)$, that integrates (1) from an initial solution $\mathbf{u}_0 = \mathbf{u}(t_0)$ until time t_i, e.g., an implicit Euler method with fine time step. Furthermore, let $\mathcal{G}$ denote a coarse propagator of lower precision, e.g., an implicit Euler method with a larger time step or loosened tolerances in case of an adaptive scheme.

The total time interval is split into smaller intervals $(t_{j-1}, t_j]$ with $t_0 < t_1 < ... < t_N$. In an iterative manner, we solve (1) concurrently on these intervals $(t_{j-1}, t_j]$ with initial conditions $\mathbf{U}_j$ at t_j using the fine propagator. We impose continuity at the time interval boundaries t_j, that is,

$$\begin{cases} \mathbf{U}_0 - \mathbf{u}_0 & = \mathbf{0}, \\ \mathbf{U}_1 - \mathcal{F}(t_1, t_0, \mathbf{U}_0) & = \mathbf{0}, \\ \vdots & \vdots \\ \mathbf{U}_{N-1} - \mathcal{F}(t_{N-1}, t_{N-2}, \mathbf{U}_{N-2}) & = \mathbf{0}. \end{cases}$$

This non-linear system of equations can be solved with the Newton-Raphson (NR) method [4], which leads to the explicit update formula for the initial conditions at iteration k known from [8]

$$\mathbf{U}_j^{(k)} = \mathcal{F}\left(t_j, t_{j-1}, \mathbf{U}_{j-1}^{(k-1)}\right) + \mathcal{G}\left(t_j, t_{j-1}, \mathbf{U}_{j-1}^{(k)}\right) - \mathcal{G}\left(t_j, t_{j-1}, \mathbf{U}_{j-1}^{(k-1)}\right).$$

The PR algorithm iterates until $\mathrm{err}^{(k)} < \mathrm{tol}_{\mathrm{PR}}$ with $\mathrm{err}^{(k)}$ the error at iteration k and $\mathrm{tol}_{\mathrm{PR}}$ the tolerance of the PR iteration. We denote the iteration index at which PR converged by K. Both $\mathrm{err}^{(k)}$ and $\mathrm{tol}_{\mathrm{PR}}$ will be discussed in more detail for the NI coil in Sect. 4.

The choice of time interval boundaries t_j is important to ensure load-balancing of the parallel loop and thus speed-up of the method. For NI coils with non-trivial transient phenomena, a simple choice of equidistant time steps is not suitable. Instead, a black box approach for finding t_j automatically is desirable for use by non-experts. To this end, we define an initial coarse propagator $\left(\hat{\mathbf{U}}, \hat{\mathbf{t}}\right) =$

$\hat{\mathcal{G}}(t_N, t_0, \mathbf{u}_0)$ that is using an adaptive time integration scheme with large tolerances. Its M time steps are denoted as $\hat{\mathbf{t}} = \left[\hat{t}_0 = t_0, \hat{t}_1, ..., \hat{t}_M = t_N\right]^{\mathrm{T}}$ and its solution at these time steps is denoted as $\hat{\mathbf{U}} = \left[\hat{\mathbf{u}}_0 = \mathbf{u}_0, \hat{\mathbf{u}}_1, ..., \hat{\mathbf{u}}_M\right]^{\mathrm{T}}$. We split these time steps into time windows of size floor (M/N) or floor $(M/N)+1$ using the floor function to transform its rational argument into an integer index, i.e.,

$$t_j = \hat{t}_{\mathrm{floor}\left(\frac{M}{N}j\right)} \quad \text{for } j \in [0, ..., N] \text{ and assuming } M \geq N. \tag{2}$$

For iterations $k > 1$, we use the coarse propagator $\mathcal{G}$ with the fixed time steps $\hat{\mathbf{t}}$ from the adaptive propagator $\hat{\mathcal{G}}$. The resulting algorithm is summarized in Algorithm 1. One advantage lies in PR being a non-intrusive method that does not require direct access to the FE system matrices, thus the FE solver can be used as a black box [14].

Algorithm 1 Parareal based on [8] extended with a first adaptive coarse integration to find the time intervals t_j.

```
init: $\mathbf{U}_0^{(k)} \leftarrow \mathbf{u}_0$ for all $k \in [0, N]$; set counter: $k \leftarrow 0$;
while $k < 1$ or $\mathrm{err}^{(k)} > \mathrm{tol}_{\mathrm{PR}}$ do
    increment counter: $k \leftarrow k + 1$;
    if k == 1 then
        solve coarse adaptively: $\hat{\mathbf{U}}, \hat{\mathbf{t}} \leftarrow \hat{\mathcal{G}}(t_N, t_0, \mathbf{u}_0)$;
        define time windows: $t_j \leftarrow \hat{t}_{\mathrm{floor}\left(\frac{M}{N}j\right)}$ for all $j \in [0, N]$;
        extract first coarse solution at $t_j$: $\mathbf{U}_j^{(1)} \leftarrow \bar{\mathbf{u}}_j^{(1)} \leftarrow \hat{\mathbf{U}}|_{t=t_j}$ for all $j \in [0, N]$;
    else
        for $j \leftarrow 1$ to $N$ do
            solve coarse: $\bar{\mathbf{u}}_j^{(k)} \leftarrow \mathcal{G}(t_j, t_{j-1}, \mathbf{U}_{j-1}^{(k)})$;
            update: $\mathbf{U}_j^{(k)} \leftarrow \tilde{\mathbf{u}}_j^{(k-1)} + \bar{\mathbf{u}}_j^{(k)} - \bar{\mathbf{u}}_j^{(k-1)}$;
        end for
    end if
    parallel for $j \leftarrow 1$ to $N$
        solve fine: $\tilde{\mathbf{u}}_j^{(k)} \leftarrow \mathcal{F}(t_j, t_{j-1}, \mathbf{U}_{j-1}^{(k)})$;
    end parallel for
end while
```

4 Numerical Example: No-Insulation HTS Coil

A small single three-dimensional NI pancake coil with 10 turns is considered as shown in Fig. 1. The geometrical and material parameters of the coil are identical to those in [10] with two exceptions; the number of turns has been reduced from 24 to 10, and an additional contact resistance between winding and terminal has been included as described in [1].

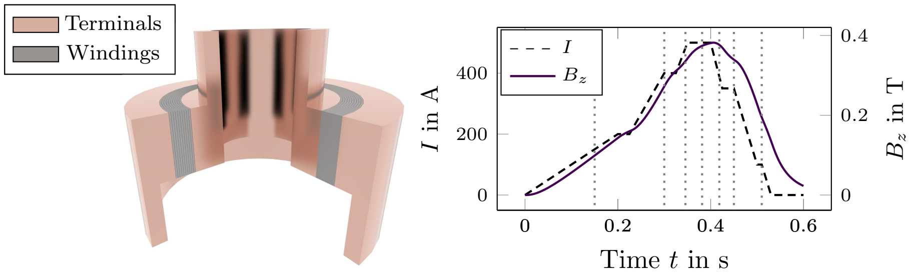

Figure 1: (left) The simulated NI coil consists of 10 turns of an HTS coated conductor surrounded by copper terminals that are used to connect to the source current. (right) Time evolution of the source current that results in a delayed central axial magnetic flux density B_z. The vertical, dashed lines show the division of the time intervals t_j for $N = 8$ calculated with Algorithm 1.

All software used is free and open-source. The model has been created using the open-source Pancake3D module [1] of the Finite Element Quench Simulator FiQuS [15] developed at CERN. It uses Gmsh [6] for creating FE meshes and GetDP [3] as the FE solver with the automatically created cohomology basis for the $\vec{H} - \phi$ formulation provided by a Gmsh plugin [9]. The PR algorithm has been implemented in Python, including MPI-parallel calls to GetDP, can be found in [12] and is based on the implementation in [5].

The pancake coil is excited by imposing a current ramping scheme shown in Fig. 1. The central axial magnetic flux density is also depicted and shows the typical characteristic delay for NI coils between field and source current [13].

The non-linear equations are linearized using an NR scheme as described in [10]. Convergence of the iterative NR scheme is assumed when the absolute change of maximum temperature between two subsequent iterations is below a certain tolerance $\mathrm{tol}_{\mathrm{NR}}$. This convergence criterion based on the maximum temperature is suitable for quench simulations as it is a key quantity of interest and allows for easy physical interpretation [11]. Therefore, the PR convergence is also based on the discontinuity of the maximum temperature at the time window boundaries t_j, i.e.,

$$\mathrm{err}^{(k)} \;=\; \max_j \left| \max_T \mathbf{U}_j^{(k)} - \max_T \mathcal{F}\left(t_j, t_{j-1}, \mathbf{U}_j^{(k-1)}\right) \right|$$

for all $j \in [0, N]$, where $\max_T \mathbf{u}$ extracts the maximum temperature $T_{\max}$ from the solution vector $\mathbf{u}$.

An implicit Euler scheme with an adaptive time step as implemented in GetDP is used for the fine propagator $\mathcal{F}$ as well as the adaptive first coarse propagator $\hat{\mathcal{G}}$. GetDP uses a prediction by polynomial extrapolation based on past results. The method accepts a time step if the approximated local truncation error of the maximum temperature is below tol_t. Both adaptive propagators are coded to enforce a time step at a time when the ramping rate changes (e.g., from linear ramp-up to plateau). All subsequent coarse solves using $\mathcal{G}$ use a fixed time step implicit Euler scheme provided by GetDP with the same time steps $\hat{\mathbf{t}}$ defined by $\hat{\mathcal{G}}$.

To ensure a fair comparison of the speed-up achievable by PR, i.e. the ratio between computational time taken by PR and a sequential reference run using the fine propagator $\hat{\mathcal{F}}$ from t_0 to t_N, it is important to consider the required engineering tolerance for $\hat{\mathcal{F}}$. Fig. 4 shows the time evolutions of the maximum temperature and the absolute error compared to a fine reference for different tolerances $\mathrm{tol}_{\mathrm{NR}} = \mathrm{tol}_t$. A numerical quench study typically includes a sensitivity analysis of these tolerances, i.e., solving with finer tolerances to ensure that all dynamic effects are captured accurately.

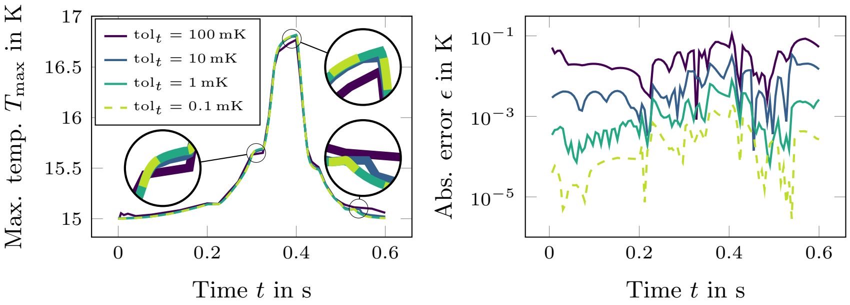

Figure 2: The maximum temperature $T_{\max}$ is shown as a result of a sequential solve with different tolerances $\mathrm{tol}_{\mathrm{NR}} = \mathrm{tol}_t$. The absolute error ϵ of $T_{\max}$ has been computed using a fine reference run with $\mathrm{tol}_{\mathrm{NR}} = \mathrm{tol}_t = 0.01\,\mathrm{mK}$. Both graphs share the same legend.

Table 1 shows the results for PR runs with $N \in \{8, 16, 24\}$ time windows for different tolerances $\mathrm{tol}_{\mathcal{F}} = \mathrm{tol}_{\mathrm{NR},\mathcal{F}} = \mathrm{tol}_{t,\mathcal{F}}$ for the fine propagator. For all runs, the PR tolerance $\mathrm{tol}_{\mathrm{PR}} = 10\,\mathrm{mK}$ is used. For $\mathcal{G}$ and $\hat{\mathcal{G}}$, the NR scheme tolerance is set to $\mathrm{tol}_{\mathrm{NR},\mathcal{G}} = \mathrm{tol}_{\mathrm{NR},\hat{\mathcal{G}}} = 10\,\mathrm{mK}$ and the time stepping tolerance for $\hat{\mathcal{G}}$ is $\mathrm{tol}_{t,\hat{\mathcal{G}}} = 50\,\mathrm{mK}$. These parameters have led to the highest speed-up among the parameter sets tried in this study. No speed-up has been achieved for the highest $\mathrm{tol}_{\mathcal{F}}$. For $\mathrm{tol}_{\mathcal{F}} = 1\,\mathrm{mK}$ and $\mathrm{tol}_{\mathcal{F}} = 0.1\,\mathrm{mK}$, a maximum speed-up of 1.59 and 2.74, respectively, has been achieved. In both cases, the best speed-up has been achieved for 16 time windows.

While considerable speed-up is achieved, it is clearly below the optimal speed-up of N/K. Fig. 3 shows the cumulative time taken by $\hat{\mathcal{G}}$ and $\mathcal{G}$ for all PR iterations k and the minimum, average, and maximum over time windows j for the cumulative time taken by $\mathcal{F}$ for all PR iterations k. Furthermore, the load balancing l, i.e. the ratio between the aforementioned minimum and maximum, is shown as well. Firstly, we observe that the computational time for the coarse propagators is not negligible. Quite fine tolerances for the coarse propagator need to be chosen for PR to converge in a few iterations and to ensure $l \gg 0$ on the fine level. For example, an initial value far away from the actual solution will cause many costly corrective steps of the adaptive prop-

Table 1: PR results for different number of time windows and fine tolerances, for which $\mathrm{tol}_{\mathrm{PR}} = 10\,\mathrm{mK}$, $\mathrm{tol}_{\mathrm{NR},\mathcal{G}} = \mathrm{tol}_{\mathrm{NR},\hat{\mathcal{G}}} = 10\,\mathrm{mK}$ and $\hat{\mathcal{G}}$ $\mathrm{tol}_{t,\hat{\mathcal{G}}} = 50\,\mathrm{mK}$ have been used.

# time windows N	8			16			24		
Fine tol. $\mathrm{tol}_{\mathrm{NR},\mathcal{F}} = \mathrm{tol}_{t,\mathcal{F}}$ in mK	10	1	0.1	10	1	0.1	10	1	0.1
Convergence after K iterations	2	2	2	4	2	2	3	3	3
Absolute error $\mathrm{err}^{(K)}$ in mK	5.2	2.4	2.5	5.8	5.1	5.5	7.4	1.5	1.5
Max. possible speed-up N/K	4	4	4	4	8	8	8	8	8
Actual speed-up	0.63	1.35	2.13	0.36	1.59	2.74	0.47	1.03	1.88

agator $\mathcal{F}$. Furthermore, there is overhead for the adaptivity of $\hat{\mathcal{G}}$. Secondly, the loads between the time windows of the fine propagator $\mathcal{F}$ are not balanced perfectly, i.e., $l < 1$. Future work should therefore focus on improving the load balancing and finding cheaper coarse propagators while maintaining a black box usage possible.

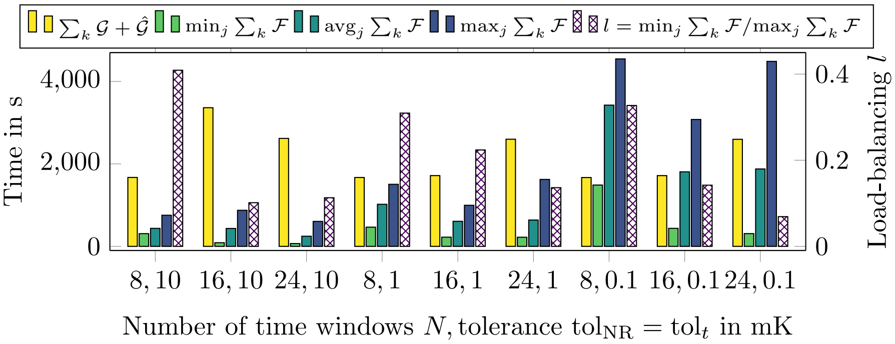

Figure 3: Cumulative time taken by $\hat{\mathcal{G}}$ and $\mathcal{G}$ for all PR iterations k and the minimum, average, and maximum over time windows j for the cumulative time taken by $\mathcal{F}$ for all k. The load balancing l, the ratio between the minimum and maximum, is also shown.

5 Conclusion

The PR method has been used to simulate the magneto-thermal transient behavior of a highly nonlinear superconducting pancake coil during a powering cycle. A first adaptive coarse propagator provided a suitable non-equidistant time window partitioning without requiring user-defined time window boundaries. It has been shown that a speed-up of up to 2.75 for 16 time windows can be achieved if an accurate solution is required. The whole software stack used is open-source. Quench simulations are an interesting challenge for parallel-in-time integration, and future work should address the imperfect load balancing on the fine level and develop methods to reduce the time taken by the coarse propagators.

Acknowledgments. The work of E. Schnaubelt is supported by the Graduate School CE within the Centre for Computational Engineering at TU Darmstadt and by the Wolfgang Gentner Programme of the German Federal Ministry of Education and Research (grant no. 13E18CHA). This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955701. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Belgium, France, Germany, and Switzerland.

References

[1] Sina Atalay, Erik Schnaubelt, Mariusz Wozniak, Julien Dular, Georgia Zachou, Sebastian Schöps, and Arjan Verweij. An open-source 3D FE quench simulation tool for no-insulation HTS pancake coils. *Superconductor Science and Technology*, 2024, in press. `doi:10.1088/1361-6668/ad3f83`.

[2] Julien Dular, Christophe Geuzaine, and Benoît Vanderheyden. Finite-element formulations for systems with high-temperature superconductors. *IEEE Trans. Appl. Super.*, 30(3):1–13, 2020. `doi:10.1109/TASC.2019.2935429`.

[3] P. Dular, C. Geuzaine, F. Henrotte, and W. Legros. A general environment for the treatment of discrete problems and its application to the finite element method. *IEEE Trans. Magn.*, 34(5):3395–3398, 1998. `doi:10.1109/20.717799`.

[4] Martin J. Gander and Ernst Hairer. Nonlinear convergence analysis for the parareal algorithm. In Ulrich Langer, Marco Discacciati, David E. Keyes, Olof B. Widlund, and Walter Zulehner, editors, *Domain Decomposition Methods in Science and Engineering XVII*, pages 45–56. Springer, 2008. `doi:10.1007/978-3-540-75199-1_4`.

[5] Martin J. Gander, Iryna Kulchytska-Ruchka, Innocent Niyonzima, and Sebastian Schöps. A new parareal algorithm for problems with discontinuous sources. *SISC*, 41(2):B375–B395, 2019. `doi:10.1137/18M1175653`.

[6] Christrophe Geuzaine and Jean-François Remacle. Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *Int. J. Numer. Meth. Eng.*, 79:1309–1331, 2009. `doi:10.1002/nme.2579`.

[7] Seungyong Hahn, Dong Keun Park, Juan Bascunan, and Yukikazu Iwasa. HTS pancake coils without turn-to-turn insulation. *IEEE Trans. Appl. Super.*, 21(3):1592–1595, 2011. doi:10.1109/TASC.2010.2093492.

[8] Jacques-Louis Lions, Yvon Maday, and Gabriel Turinici. A parareal in time discretization of PDEs. *Comptes Rendus de l'Académie des Sciences – Series I – Mathematics*, 332(7):661–668, 2001. doi:10.1016/S0764-4442(00)01793-6.

[9] Matti Pellikka, Saku Suuriniemi, Lauri Kettunen, and Christophe Geuzaine. Homology and cohomology computation in finite element modeling. *SISC*, 35(5):B1195–B1214, 2013. doi:10.1137/130906556.

[10] Erik Schnaubelt, Sina Atalay, Mariusz Wozniak, Julien Dular, Christophe Geuzaine, Nicolas Marsic, Benoît Vanderheyden, Arjan Verweij, and Sebastian Schöps. Magneto-thermal thin shell approximation for 3D finite element analysis of no-insulation coils. *IEEE Trans. Appl. Super.*, 34(3):1–6, 2024. doi:10.1109/TASC.2023.3340648.

[11] Erik Schnaubelt, Mariusz Wozniak, and Sebastian Schöps. Thermal thin shell approximation towards finite element quench simulation. *Supercond. Sci. Technol.*, 36(4):044004, 2023. doi:10.1088/1361-6668/acbeea.

[12] Erik Schnaubelt, Mariusz Wozniak, and Sebastian Schöps. Parareal implementation in Python to study transient phenomena in no-insulation superconducting coils, 2024. doi:10.5281/zenodo.10886175.

[13] Erik Schnaubelt, Mariusz Wozniak, Sebastian Schöps, and Arjan Verweij. Electromagnetic simulation of no-insulation coils using $h - \phi$ thin shell approximation. *IEEE Trans. Appl. Super.*, 33(5):1–6, 2023. doi:10.1109/TASC.2023.3258905.

[14] Sebastian Schöps, Innocent Niyonzima, and Markus Clemens. Parallel-in-time simulation of eddy current problems using Parareal. *IEEE Trans. Magn.*, 54(3):1–4, 2018. doi:10.1109/TMAG.2017.2763090.

[15] Andrea Vitrano, Mariusz Wozniak, Erik Schnaubelt, Tim Mulder, Emmanuele Ravaioli, and Arjan Verweij. An open-source finite element quench simulation tool for superconducting magnets. *IEEE Trans. Appl. Super.*, 33(5):1–6, 2023. doi:10.1109/TASC.2023.3259332.