

Towards parallel time-stepping for the numerical simulation of atherosclerotic plaque growth

Frei, Stefan; Heinlein, Alexander

DOI

[10.1016/j.jcp.2023.112347](https://doi.org/10.1016/j.jcp.2023.112347)

Publication date

2023

Document Version

Final published version

Published in

Journal of Computational Physics

Citation (APA)

Frei, S., & Heinlein, A. (2023). Towards parallel time-stepping for the numerical simulation of atherosclerotic plaque growth. *Journal of Computational Physics*, 491, Article 112347. <https://doi.org/10.1016/j.jcp.2023.112347>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy


Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Towards parallel time-stepping for the numerical simulation of atherosclerotic plaque growth



Stefan Frei^{a,*}, Alexander Heinlein^b

^a Department of Mathematics and Statistics, University of Konstanz, 78457 Konstanz, Germany

^b Delft University of Technology, Delft Institute of Applied Mathematics, Mekelweg 4, 2628CD Delft, Netherlands

ARTICLE INFO

Article history:

Received 9 March 2022

Received in revised form 28 June 2023

Accepted 2 July 2023

Available online 10 July 2023

Keywords:

Atherosclerotic plaque growth

Temporal homogenization

Parallel time-stepping

Parareal algorithm

Fluid-structure interactions

ABSTRACT

The numerical simulation of atherosclerotic plaque growth is computationally prohibitive, since it involves a complex cardiovascular fluid-structure interaction (FSI) problem with a characteristic time scale of milliseconds to seconds, as well as a plaque growth process governed by reaction-diffusion equations, which takes place over several months. In this work we combine a temporal homogenization approach, which separates the problem in computationally expensive FSI problems on a micro scale and a reaction-diffusion problem on the macro scale, with parallel time-stepping algorithms. It has been found in the literature that parallel time-stepping algorithms do not perform well when applied directly to the FSI problem. To circumvent this problem, a parareal algorithm is applied on the macro-scale reaction-diffusion problem instead of the micro-scale FSI problem. We investigate modifications in the coarse propagator of the parareal algorithm, in order to further reduce the number of costly micro problems to be solved. The approaches are tested in detailed numerical investigations based on serial simulations.

© 2023 Elsevier Inc. All rights reserved.

1. Introduction

Cardiovascular diseases are by far the most common cause of death in industrialized nations today. One of the most common cardiovascular diseases is the growth of plaque (atherogenesis) in coronary arteries or the pathological accumulation of plaque (atherosclerosis), which can result in heart attacks or strokes that are often fatal. Since the formation of plaque ranges over a long period of time, from months to several years, early diagnosis and treatment to prevent plaque growth can have a good chance of success.

An important driving force for plaque growth is the wall shear stress distribution, which varies significantly within each heartbeat, i.e., every second; see, for instance, [20,57] for a discussion on the dependence of plaque growth on the wall shear stress. Using three-dimensional fluid-structure interaction (FSI) simulations with realistic material models, the heartbeat scale has to be resolved by time steps in the order of milliseconds. We refer to [5,6] for the first studies on using complex wall models that take into account the effect of the reinforcing fibers within three-dimensional FSI simulations. For further references on FSI for cardiovascular applications, see, e.g., [9,10,26,38,40]; for FSI simulations in the context of plaque growth, see also [25,28,30,46,56,60].

* Corresponding author.

E-mail addresses: stefan.frei@uni-konstanz.de (S. Frei), a.heinlein@tudelft.nl (A. Heinlein).

Hence, realistic numerical simulations of plaque growth over several months that resolve each individual heartbeat would easily require $\mathcal{O}(10^9)$ sequential time steps. Consequently, even using today's fastest supercomputers, such a simulation is clearly infeasible.

Therefore, in [29] a temporal homogenization approach for separating the plaque growth time scale (macro scale) in the order of days and the FSI time scale (micro scale) in the order of milliseconds has been introduced. The approach is based on the assumption that the FSI is approximately periodic in time on the micro scale, which makes it possible to upscale the fluid dynamics to the macro scale. Using this approach it is possible to simulate only a moderate number of FSI time steps for each time step of the plaque growth problem. This means that we have to simulate only a few heartbeats (seconds) of the full FSI problem instead of the whole 24 hours of each day; the total number of time steps is reduced by a factor of $\mathcal{O}(10^4)$. Note that, completely neglecting the fine scale may introduce a significant error; see [28,30]. See also the recent PhD thesis of Florian Sonner for further details on temporal homogenization for plaque growth [55].

However, considering realistic three-dimensional simulations the number of time steps corresponding to fine scale FSI problems remains still infeasibly high, even after temporal homogenization. In order to further reduce the computational times, we introduce a new approach based on parallel time-stepping. We focus on a classical parallel time-stepping method, the parareal algorithm, which has been introduced by Lions et al. in [42]; for an overview over the vast literature on parallel-in-time integration methods, we refer to the review papers [31,45] and the references therein.

Due to a phase-shift in a coarse solution for hyperbolic partial differential equations (PDEs), such as the structural problem in FSI, it is generally challenging to apply parallel time-stepping methods to FSI problems; see, e.g., [43,51]. Thus, instead of applying the parareal algorithm to the micro-scale FSI problem, we apply it to the homogenized plaque growth problem on the macro scale. Plaque growth is typically modeled by a system of reaction-diffusion equations; see, e.g., [54, 59]. Instead of being hyperbolic, they have a parabolic character and hence can be solved more efficiently by parallel time-stepping methods.

The focus of this paper is not the computation of accurate plaque-growth predictions in patient-specific geometries. Instead, our objective is to introduce a numerical framework for making the simulation times feasible and to investigate the methods numerically. Therefore, we make several simplifications: We focus on two-dimensional FSI simulations on a simple geometry. Furthermore, we consider two simplified models for the plaque growth, that is, the ordinary differential equation (ODE) model already considered in [28,29] as well as a more complex partial differential equation (PDE) of reaction-diffusion type. We formulate the parareal algorithm for the two-scale problem of plaque growth and also propose variants to increase the efficiency of the coarse-scale propagators. Finally, we investigate the parallel time-stepping methods numerically based on serial simulations. As, even in two dimensions, the micro-scale FSI problems are generally much more expensive than the coarse plaque growth problem, we are able to give some good estimates for the speedup that can be expected in fully parallel simulations.

The main novelties of this publication are as follows:

- To our knowledge this is the first paper that applies a parallel-in-time algorithm to the problem of atherosclerotic plaque growth. Therefore, the parareal algorithm is combined with the temporal homogenization approach presented in [29].
- In contrast to our previous works [28–30,53], we consider also a PDE model of diffusion-reaction type for the modelling of plaque growth.
- The fact that the solution of the micro problems (non-stationary FSI problems) is extremely expensive, for example, compared to the macro problems, allows us to simplify the complexity analysis of the parareal algorithm by only counting the number of micro problems that need to be solved.
- The complexity analysis motivates the development of a new variant of the parareal algorithm, based on the idea of re-using previously computed growth values (instead of solving expensive micro problems) within the coarse propagator. A convergence analysis is given for a prototypical model problem.
- We show numerically that, when using the heuristic averaging approach of [17,30,59] as an approximation for the non-stationary FSI problem in the coarse propagator, parareal convergence is much slower.
- Detailed numerical studies investigate the performance of the algorithms under consideration both for an ODE and a PDE growth model, based on serial simulations.

The paper is structured as follows: In section 2, we introduce our fluid-structure interaction problem as well as the two solid growth models considered for modeling the plaque growth: an ODE-based model in section 2.2.1 and a reaction-diffusion equation-based model in section 2.2.2. Next, in section 3, we briefly introduce the variational formulation of the FSI problem and then describe the temporal homogenization approach including some first numerical results. In section 4, we recap the parareal algorithm and discuss how we estimate the computational costs and the possible speed up in parallel simulations. We also discuss numerical results for the parallel time-stepping approach using the simple ODE plaque growth model. Furthermore, we introduce some ideas for reducing the costs of the coarse-scale propagators. In section 5, we show results for the reaction-diffusion plaque growth model, including a theoretical discussion on the computational costs of the algorithms. We conclude in section 6, where we also give a brief outlook on topics for future work.

2. Model equations

We consider a time-dependent fluid-structure interaction system, where the fluid is modeled by the Navier–Stokes equations and the solid by the Saint Venant–Kirchhoff model. In order to account for the solid growth, we add a multiplicative growth term to the deformation gradient, which is motivated by typical plaque growth models [30,44,54,59].

2.1. Fluid-structure interaction

Consider a partition of an overall domain $\Omega(t) = \mathcal{F}(t) \cup \Gamma(t) \cup \mathcal{S}(t)$ into a fluid part $\mathcal{F}(t)$, an interface $\Gamma(t)$ and a solid part $\mathcal{S}(t)$. The blood flow and its interaction with the surrounding vessel wall is modelled by the following FSI system:

$$\begin{aligned} \rho_f(\partial_t \mathbf{v}_f + \mathbf{v}_f \cdot \nabla \mathbf{v}_f) - \operatorname{div} \boldsymbol{\sigma}_f &= 0, & \operatorname{div} \mathbf{v}_f &= 0 & \text{in } \mathcal{F}(t), \\ \rho_s \partial_t \hat{\mathbf{v}}_s - \operatorname{div}(\hat{\mathbf{F}}_e \hat{\boldsymbol{\Sigma}}_e) &= 0, & \partial_t \hat{\mathbf{u}}_s - \hat{\mathbf{v}}_s &= 0 & \text{in } \hat{\mathcal{S}}, \\ \boldsymbol{\sigma}_f \vec{n}_f + \boldsymbol{\sigma}_s \vec{n}_s &= 0, & \mathbf{v}_f &= \mathbf{v}_s & \text{on } \Gamma(t). \end{aligned} \quad (1)$$

Here, \mathbf{v}_f and $\hat{\mathbf{v}}_s$ denote the fluid and solid velocity, respectively, and $\hat{\mathbf{u}}_s$ the solid displacement. Quantities with a “hat” are defined in Lagrangian coordinates, while quantities without a “hat” are defined in the current Eulerian coordinate framework. Two quantities $\hat{f}(\hat{x})$ and $f(x)$ correspond to each other by a $C^{1,1}$ -diffeomorphism $\hat{\xi} : \hat{\Omega} \rightarrow \Omega(t)$ and the relation $\hat{f} = f \circ \hat{\xi}$. Later, we will also need the solid deformation gradient $\hat{F}_s = I + \hat{\nabla} \hat{\mathbf{u}}_s$, which is the derivative of $\hat{\xi}$ in the solid part. The constants ρ_f and ρ_s are the densities of blood and vessel wall, and \vec{n}_f and \vec{n}_s are outward pointing normal vectors of the fluid and solid domain, respectively.

By $\boldsymbol{\sigma}_f$ and $\boldsymbol{\sigma}_s$ we denote the Cauchy stress tensors of fluid and solid. Using the well-known Piola transformation between the Eulerian and the Lagrangian coordinate systems, we can relate $\boldsymbol{\sigma}_s$ and the second Piola–Kirchhoff stress $\hat{\boldsymbol{\Sigma}}_e$ as follows:

$$\boldsymbol{\sigma}_s(\mathbf{x}) = \hat{\boldsymbol{\sigma}}_s(\hat{x}) = \hat{J}_e^{-1} \hat{\mathbf{F}}_e \hat{\boldsymbol{\Sigma}}_e(\hat{x}) \hat{\mathbf{F}}_e^T,$$

where \hat{F}_e is the elastic part of the deformation gradient \hat{F}_s . For modeling the material behavior of the vessel wall, different approaches have been proposed and investigated in literature; see, e.g., [5,7,8,13,37] for more sophisticated material models, for instance, incorporating an anisotropic behavior due to the reinforcing fibers. For the sake of simplicity, we use in this work the relatively simple Saint Venant–Kirchhoff model with the Lamé material parameters μ_s and λ_s

$$\hat{\boldsymbol{\Sigma}}_e = 2\mu_s \hat{\mathbf{E}}_e + \lambda_s \operatorname{tr}(\hat{\mathbf{E}}_e) I, \quad \hat{\mathbf{E}}_e := \frac{1}{2}(\hat{\mathbf{F}}_e^T \hat{\mathbf{F}}_e - I). \quad (2)$$

The Saint Venant–Kirchhoff model is based on Hooke’s linear material law in a large strain formulation, resulting in a weakly nonlinear material model.

The blood flow is modeled as an incompressible Newtonian fluid, such that the Cauchy stresses are given by

$$\boldsymbol{\sigma}_f = \rho_f \nu_f (\nabla \mathbf{v}_f + \nabla \mathbf{v}_f^T) - p_f I, \quad (3)$$

where ν_f is the kinematic viscosity of blood. A sketch of the computational domain is given in Fig. 1. We split the outer boundary of Ω into a solid part Γ_s with homogeneous Dirichlet conditions, a fluid part Γ_f^{in} with an inflow Dirichlet condition and an outflow part Γ_f^{out} , where a *do-nothing condition* is imposed.

The boundary data is given by

$$\mathbf{v}_f = \mathbf{v}^{\text{in}} \text{ on } \Gamma_f^{\text{in}}, \quad \rho_f \nu_f (\vec{n}_f \cdot \nabla) \mathbf{v}_f - p_f \vec{n}_f = 0 \text{ on } \Gamma_f^{\text{out}}, \quad \hat{\mathbf{u}}_s = 0 \text{ on } \hat{\Gamma}_s, \quad (4)$$

where \mathbf{v}^{in} is the inflow velocity on Γ_f^{in} .

2.2. Modelling of solid growth

Developing a realistic model of plaque growth at the vessel wall is a complex task that involves the interaction of many different molecules and species, see for example [54]. Furthermore, the plaque growth will also strongly depend on the geometry and material model of the arterial wall. In this contribution, our focus does not lie on a realistic modeling of plaque growth, and thus, our model is greatly simplified. However, the two models considered here are chosen to behave similarly (from a numerical viewpoint) to more sophisticated models. We consider a simple ODE-based model in section 2.2.1 and a more complex but still relatively simple PDE model of reaction-diffusion type in section 2.2.2. Both models focus on the influence of the concentration of foam cells c_s on the growth. Numerical results can be found in sections 3.2, 4.1.4 and 4.2 as well as in section 5, respectively.




Fig. 1. Sketch of the computational domain centered at the origin $(0, 0)$; \mathcal{F} and \mathcal{S} are the fluid and solid parts, respectively, and the solid lines correspond to the fluid–solid interface Γ . The plaque growth is initiated at $(0, \pm 1)$.

2.2.1. A simple ODE model for solid growth

In the first model, which is taken from [28,30], the evolution of the foam cell concentration c_s depends only on its current value but not on its spatial distribution. Hence, it can be described by a simple ODE. The rate of formation of these cells depends on the distribution of the wall shear stress σ_f^{WS} at the vessel wall. As a result, we obtain the following simplified ODE model:

$$\partial_t c_s = \gamma(\sigma_f^{WS}, c_s) := \alpha(1 + c_s)^{-1} \left(1 + \frac{\|\sigma_f^{WS}\|_{L^2(\Gamma)}^2}{\sigma_0^2} \right)^{-1}, \tag{5}$$

$$\sigma_f^{WS} := \rho\nu(I_d - \vec{n}_f \vec{n}_f^T)(\nabla \mathbf{v} + \nabla \mathbf{v}^T)\vec{n}_f.$$

The reference wall shear stress σ_0 and the scale separation parameter α are parameters of the growth model. For cardiovascular plaque growth, we have typically $\alpha = \mathcal{O}(10^{-7}) \text{ s}^{-1}$; see also [29].

We model the solid growth by a multiplicative splitting of the deformation gradient $\hat{\mathbf{F}}_s$ into an elastic part $\hat{\mathbf{F}}_e$ and a growth function $\hat{\mathbf{F}}_g$

$$\hat{\mathbf{F}}_s = \hat{\mathbf{F}}_e \hat{\mathbf{F}}_g \Leftrightarrow \hat{\mathbf{F}}_e = \hat{\mathbf{F}}_s \hat{\mathbf{F}}_g^{-1} = [I + \hat{\nabla} \hat{\mathbf{u}}_s] \hat{\mathbf{F}}_g^{-1}; \tag{6}$$

cf. [30,50,59]. In the ODE model, we use the following growth function depending on c_s

$$\hat{g}(\hat{x}, \hat{y}, t) = 1 + c_s \exp(-\hat{x}^2) (2 - |\hat{y}|), \quad \hat{\mathbf{F}}_g(\hat{x}, \hat{y}, t) := \hat{g}(\hat{x}, \hat{y}, t) I. \tag{7}$$

This means that the shape and position of the plaque growth is prescribed, but the growth rate depends on the variable c_s . As the simulation domain is centered around the origin, see Fig. 1, growth is concentrated at $(0, \pm 1)$, in the center of the domain. It follows that

$$\hat{\mathbf{F}}_g = \hat{g}I \Rightarrow \hat{\mathbf{F}}_e := \hat{g}^{-1} \hat{\mathbf{F}}_s, \tag{8}$$

and the elastic Green–Lagrange strain is given by

$$\hat{\mathbf{E}}_e = \frac{1}{2}(\hat{\mathbf{F}}_e^T \hat{\mathbf{F}}_e - I) = \frac{1}{2}(\hat{g}^{-2} \hat{\mathbf{F}}_s^T \hat{\mathbf{F}}_s - I) \tag{9}$$

resulting in the Piola–Kirchhoff stresses

$$\hat{\mathbf{F}}_e \hat{\Sigma}_e = 2\mu_s \hat{\mathbf{F}}_e \hat{\mathbf{E}}_e + \lambda_s \text{tr}(\hat{\mathbf{E}}_e) \hat{\mathbf{F}}_e = 2\mu_s \hat{g}^{-1} \hat{\mathbf{F}}_s \hat{\mathbf{E}}_e + \lambda_s \hat{g}^{-1} \text{tr}(\hat{\mathbf{E}}_e) \hat{\mathbf{F}}_s. \tag{10}$$

2.2.2. A PDE reaction-diffusion model

Secondly, we consider a slightly more complex growth model, where the concentration of foam cells $\hat{c}_s = \hat{c}_s(\hat{x}, t)$ ($\hat{x} \in \hat{\mathcal{S}}$) is now governed by a PDE model, namely a non-stationary reaction-diffusion equation

$$\begin{aligned} \partial_t \hat{c}_s - D_s \hat{\Delta} \hat{c}_s + R_s \hat{c}_s (1 - \hat{c}_s) &= 0 \quad \text{on } \hat{\mathcal{S}}, & D_s \hat{\partial}_n \hat{c}_s &= \gamma(\sigma_f^{WS}) = \alpha \delta(x) \left(1 + \frac{\|\sigma_f^{WS}\|_{L^2(\Gamma)}^2}{\sigma_0^2} \right)^{-1} \quad \text{on } \hat{\Gamma}, \\ \hat{c}_s &= 0 \quad \text{on } \hat{\Gamma}_s, & \hat{c}_s(\cdot, 0) &= 0 \quad \text{on } \hat{\mathcal{S}}, \end{aligned} \tag{11}$$

where $\delta(x) = \min\{0, (x-1)(x+1)\}^2$ and $\|\cdot\|$ is the Euclidean norm. Note that the norm $\|\sigma_f^{WS}\|$ on the right-hand side of eq. (11) is a function in space and time, whereas the right-hand side of eq. (5), including $\|\sigma_f^{WS}\|_{L^2(\Gamma)}$, depends only on time and is a constant in space. Furthermore, we have σ_0 and α as in the ODE model eq. (6), as well as positive diffusion and reaction coefficients D_s and R_s , respectively. The underlying idea is that the vessel wall is initially damaged in the central part of the interface around $\hat{x}_1 = 0$ and monocytes can penetrate into \hat{S} in the part of the interface corresponding to $\hat{x}_1 \in (-1, 1)$, where δ is actually positive.

Moreover, in contrast to the first numerical example, we do not prescribe the shape of the plaque growth by a growth function \hat{g} . Instead, the growth part F_g depends on the spatial distribution of the foam cell concentration via the relation

$$\hat{g}(\hat{x}, t) = 1 + \hat{c}_s(\hat{x}, t), \quad \hat{\mathbf{F}}_g(\hat{x}, t) := \hat{g}(\hat{x}, t) I. \quad (12)$$

The Green Lagrange strain $\hat{\mathbf{E}}_e$ and the Piola–Kirchhoff stresses $\hat{\Sigma}_e$ are then defined as above in eqs. (9) and (10), with \hat{g} given in eq. (12).

Note that, solving the PDE model has higher computational cost compared to the ODE model described in section 2.2.1. On the other hand, a sophisticated reaction-diffusion problem may yield more realistic results for the plaque growth. Nonetheless, such a model remains significantly cheaper compared to a fully coupled FSI problem; this discrepancy will only increase when moving to three-dimensional simulations with complex material models for the arterial wall. This observation is essential for the efficiency of our parallel time-stepping approach, as we will discuss in more detail in the following sections.

3. Numerical framework

In this work, we use an Arbitrary Lagrangian Eulerian (ALE) approach to solve the FSI problem in eq. (1). The ALE approach is the standard approach for FSI with small to moderate structural deformations; see, e.g., [10,16,18,21,26,48,58]. This assumption holds generally true for the simulation of plaque growth, unless the interest is to simulate a complete closure of the artery. To simulate a full closure, a Fully Eulerian formalism [22,27,47] has been used in [27,30]; see also [2, 14,15] for further works on FSI-contact problems in a Fully Eulerian or Lagrange-Eulerian formalism.

Given a suitable ALE map $\hat{\xi}_f = \hat{x} + \hat{\mathbf{u}}_f$, its gradient $\hat{\mathbf{F}}_f = I + \hat{\nabla} \hat{\mathbf{u}}_f$ and determinant $\hat{J}_f = \det \hat{\mathbf{F}}_f$, the ALE formulation of the FSI problem is given by:

Variational formulation 1 (Non-stationary FSI problem). Find velocity $\hat{\mathbf{v}} \in \mathbf{v}^{\text{in}} + \mathcal{V}$, deformation $\hat{\mathbf{u}} \in \mathcal{W}$ and pressure $\hat{p}_f \in \mathcal{L}_f$, such that

$$\begin{aligned} (\rho_f \hat{J}_f \hat{\partial}_t \hat{\mathbf{v}}_f, \hat{\phi})_{\hat{\mathcal{F}}} + (\rho_f \hat{J}_f \hat{\nabla} \hat{\mathbf{v}}_f \hat{\mathbf{F}}_f^{-1} (\hat{\mathbf{v}}_f - \hat{\partial}_t \hat{\mathbf{u}}_f), \hat{\phi})_{\hat{\mathcal{F}}} + (\rho_s^0 \hat{\partial}_t \hat{\mathbf{u}}_s, \hat{\phi})_{\hat{S}} \\ + (\hat{J}_f \hat{\sigma}_f \hat{\mathbf{F}}_f^{-T}, \hat{\nabla} \hat{\phi})_{\hat{\mathcal{F}}} - (\hat{J}_f \nu_f \hat{\mathbf{F}}_f^{-1} \hat{\nabla} \hat{\mathbf{v}}_f \hat{\mathbf{F}}_f^{-1} \hat{n}_f, \hat{\phi})_{\hat{\Gamma}^{\text{out}}} + (\hat{\mathbf{E}}_e(\hat{g}) \hat{\Sigma}_e(\hat{g}), \hat{\phi})_{\hat{S}} = 0 \quad \forall \hat{\phi} \in \mathcal{V}, \\ (\widehat{\text{div}}(\hat{J}_f \hat{\mathbf{F}}_f^{-1} \hat{\mathbf{v}}_f), \hat{\xi}_f)_{\hat{\mathcal{F}}} = 0 \quad \forall \hat{\xi}_f \in \mathcal{L}_f, \\ (\hat{d}_t \hat{\mathbf{u}}_s - \hat{\mathbf{v}}_s, \hat{\psi}_s)_{\hat{S}} = 0 \quad \forall \hat{\psi}_s \in \mathcal{L}_s, \end{aligned}$$

where $\hat{\sigma}_f := \rho_f \nu_f (\hat{\nabla} \hat{\mathbf{v}}_f \hat{\mathbf{F}}_f^{-1} + \hat{\mathbf{F}}_f^{-T} \nabla \hat{\mathbf{v}}_f^T) - \hat{p}_f I$.

There are different ways to compute the ALE map depending on the displacement $\hat{\mathbf{u}}_s$ of the structural domain \hat{S} . In our simulations, we extend $\hat{\mathbf{u}}_s$ into the fluid domain using harmonic extensions (called $\hat{\mathbf{u}}_f$); that is, we solve a Laplacian problem with right hand side zero and $\hat{\mathbf{u}}_s$ as Dirichlet boundary data on the fluid domain.

Note that $\hat{\Sigma}_e$ and $\hat{\mathbf{E}}_e$ depend on \hat{g} (and hence on c_s), as specified in eq. (10). The function spaces are given by

$$\mathcal{V} = [H_0^1(\hat{\Omega}; \hat{\Gamma}_f^{\text{in}} \cup \hat{\Gamma}_s)]^d, \quad \mathcal{W} = [H_0^1(\hat{S}; \hat{\Gamma}_s)]^d, \quad \mathcal{L}_s := L^2(\hat{S}), \quad \mathcal{L}_f = L^2(\hat{\mathcal{F}}).$$

Furthermore, we assume that the solution $\hat{\mathbf{v}}_f$ has higher regularity, such that the trace of $\hat{\nabla} \hat{\mathbf{v}}_f$ is well-defined on $\hat{\Gamma}^{\text{out}}$, as needed in Variational formulation 1.

3.1. Temporal two-scale approach

Even for the simplified two-dimensional configuration considered in this work, a resolution of the micro-scale dynamics with a scale of milliseconds to seconds is unfeasible over the complete time interval of interest $[0, T_{\text{end}}]$, with T_{end} being several months up to a year. For instance, when considering a relatively coarse micro-scale time step of $\delta\tau = 0.02$ s, the number of time steps required to simulate a time frame of a whole year would be $365 \cdot 86400 \cdot \frac{1\text{s}}{\delta\tau} \approx 1.58 \cdot 10^9$, each step corresponding to the solution of a mechano-chemical FSI problem.

This dilemma is frequently solved by considering a heuristic averaging: as the micro-scale is much smaller than the macro scale one considers a stationary limit of the FSI and solves for the stationary FSI problem on the macro scale (e.g., $\delta t \approx 1$ day); see, e.g., [17,30,59]. The wall-shear stress $\hat{\sigma}_f^{WS}$ of the solution of this stationary FSI problem is then used to advance the foam cell concentration in eq. (5) or eq. (11).




Fig. 2. Schematic representation of the two-scale algorithm as described in section 3.1 and summarized in Algorithm 1: instead of solving the FSI problem and updating the foam cell concentration c_s on the micro time scale ($\delta\tau$, black), we compute an averaged growth function $\bar{\gamma}(\sigma_f^{WS})$ over a few periods (heartbeats) of the micro-scale problem (right); this average is used to update c_s on the macro scale (δt , green in the left part). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Variational formulation 2 (Stationary FSI problem). Find velocity $\bar{\mathbf{v}} \in \mathbf{V}^{in} + \mathcal{V}$, deformation $\bar{\mathbf{u}} \in \mathcal{W}$ and pressure $\bar{p}_f \in \mathcal{L}_f$, such that

$$\begin{aligned} &(\rho_f \hat{J}_f \hat{\nabla} \bar{\mathbf{v}}_f \bar{\mathbf{F}}_f^{-1} \bar{\mathbf{v}}_f, \hat{\phi})_{\hat{\mathcal{F}}} + (\bar{J}_f \bar{\boldsymbol{\sigma}}_f \bar{\mathbf{F}}_f^{-T}, \hat{\nabla} \hat{\phi})_{\hat{\mathcal{F}}} \\ &- (\hat{J}_f \nu_f \hat{\mathbf{F}}_f^{-1} \hat{\nabla} \hat{\mathbf{v}}_f \hat{\mathbf{F}}_f^{-1} \hat{\mathbf{n}}_f, \hat{\phi})_{\hat{\mathcal{F}}^{out}} + (\bar{\mathbf{F}}_e(\hat{\mathbf{g}}) \bar{\boldsymbol{\Sigma}}_e(\hat{\mathbf{g}}), \hat{\phi})_{\hat{\mathcal{S}}} = 0 \quad \forall \hat{\phi} \in \mathcal{W}, \\ &(\widehat{\text{div}}(\bar{J}_f \bar{\mathbf{F}}_f^{-1} \bar{\mathbf{v}}_f), \hat{\xi}_f)_{\hat{\mathcal{F}}} = 0 \quad \forall \hat{\xi}_f \in \mathcal{L}_f \end{aligned}$$

The foam cell concentration is then advanced by the ODE in eq. (5) resp. the PDE in eq. (11), with σ_f^{WS} replaced by $\bar{\sigma}_f^{WS}$. It has, however, been shown that $\gamma(\sigma_f^{WS})$ is not necessarily a good approximation of $\gamma(\sigma_f^{WS})$, which depends on the pulsating blood flow; see the numerical results in [28,30] and the analysis in [29].

A more accurate two-scale approach has been presented by Frei and Richter in [29]. The numerical approach can be cast in the framework of the Heterogeneous Multiscale Method (HMM); see, e.g., [1,23,24]. In [29], a periodic-in-time micro-scale problem is solved in each time step of the macro scale, for instance each day. The growth function $\gamma(\sigma_f^{WS})$ is then averaged by integrating over one period of the heart beat, its average will be denoted by $\bar{\gamma}(\sigma_f^{WS})$. This averaged growth function is applied to advance the foam cell concentration by eq. (5) or eq. (11), respectively. A schematic illustration of the two-scale algorithm is given in Fig. 2.

To be precise, we divide the macro-scale time interval $[0, T_{end}]$ into N_f time steps of size δt

$$0 = t_0 < t_1 < \dots < t_{N_f} = T_{end}, \quad N_f = \frac{T_{end}}{\delta t}. \tag{13}$$

As c_s varies significantly on the macro scale only, using $c_s(t_m)$ as a fixed value for the growth variable on the micro scale results in a sufficiently good approximation in the time interval $[t_m, t_{m+1}]$. Then, one cycle of the pulsating blood flow problem (around 1 s) is to be resolved on the micro scale $\delta\tau$

$$0 = \tau_0 < \tau_1 < \dots < \tau_{N_s} = 1s, \quad N_s = \frac{1s}{\delta\tau}. \tag{14}$$

It has been shown in [29] (for a simplified flow configuration) that this approach leads to a model error $\mathcal{O}(\epsilon)$ compared to a full resolution of the micro scale, where $\epsilon = \frac{1s}{T_{end}}$ denotes the ratio between macro- and micro time scale and is in the range of $\mathcal{O}(10^{-7})$ for a typical cardiovascular plaque growth problem. The relations in eqs. (13) and (14) imply $\delta\tau = \epsilon\delta t$. In the model problems formulated above, the scale separation is induced by the parameter $\alpha = \mathcal{O}(\epsilon)$.

A difficulty lies in solving the periodic micro-scale problem. If accurate initial conditions $\mathbf{w}^0 := (\mathbf{v}^0, \mathbf{u}^0)$ are available on the micro-scale, a periodic solution can be computed using a time-stepping procedure for one cycle. If the initial conditions are known approximately, convergence to the periodic solution may still be obtained after simulating a few cycles of the micro-scale problem due to the dissipation of the flow problem; see [28,29]. Numerically, it can be checked after each cycle if the solution is sufficiently close to a periodic state. In this work, we apply a stopping criterion based on the computed averaged growth value:

$$|\bar{\gamma}(\sigma_f^{WS,r}) - \bar{\gamma}(\sigma_f^{WS,r-1})| < \epsilon_p,$$

where $r = 1, 2, \dots$ denotes the iteration index with respect to the number of cycles of the micro-scale problem. The algorithm is summarized as Algorithm 1, where we use the abbreviation $\mathbf{w}^{r,s} := (\mathbf{v}^{r,s}, \mathbf{u}^{r,s})$.

Algorithm 1: Two-Scale Algorithm

Set suitable starting values $\mathbf{w}^{0,0} = (\mathbf{v}^{0,0}, \mathbf{u}^{0,0})$ and time-step sizes $\delta t, \delta T$.

for $n = 1, 2, \dots, N_f$ **do**

1.) **Micro problem:** Set $r \leftarrow 0$

while $|\bar{\gamma}(\sigma_f^{WS,r}) - \bar{\gamma}(\sigma_f^{WS,r-1})| > \epsilon_p$ **do**

1.a) Solve micro-scale problem in eq. (1) in $I_n = (t_n, t_n + 1s)$

$$\{\mathbf{w}^{r,0}, c_s^{n-1}\} \mapsto \{\mathbf{w}^{r,m}\}_{m=1}^{N_s}$$

1.b) Compute the averaged growth function

$$\bar{\gamma}(\sigma_f^{WS,r+1}) = \frac{1}{N_s} \sum_{m=1}^{N_s} \gamma(\sigma_f^{WS,m}(\mathbf{v}^{r+1,m}, c_s^{n-1}))$$

and set $\mathbf{w}^{r+1,0} = \mathbf{w}^{r,N_s}$, $r \leftarrow r + 1$.

2.) **Macro problem:** Update the foam cell concentration c_s^n by eq. (5) or eq. (11).

As starting values $\mathbf{w}^{r,0} = (\mathbf{v}^{r,0}, \mathbf{u}^{r,0})$ in step 1. of the algorithm, we use the variables $\mathbf{w}^{r-1,N_s} = (\mathbf{v}^{r-1,N_s}, \mathbf{u}^{r-1,N_s})$ from the quasi-periodic state of the previous macro step. It has been observed in [28] that these are usually closer to the starting values of the periodic state than the solution of an averaged stationary problem on the macro scale.

3.2. Numerical example

Before we present the different approaches for parallel time-stepping, let us illustrate the two-scale approach by a first numerical example. Therefore, we use the simple ODE growth model introduced in section 2.2.1, eq. (5). The test configuration introduced here will be used in section 4 as well.

Concerning the geometry we use a two-dimensional channel of length 10 cm and an initially constant width $\omega(0)$ of 2 cm as illustrated in Fig. 1. The solid parts on the top and bottom corresponding to the arterial wall have an initial thickness of 1 cm each. Fluid density and viscosity are given by $\rho_f = 1 \text{ g/cm}^3$ and $\nu_f = 0.04 \text{ cm}^2/\text{s}$, respectively. The growth parameters are set to $\sigma_0 = 30 \frac{\text{g cm}}{\text{s}^2}$ and $\alpha = 5 \cdot 10^{-7} \frac{1}{\text{s}}$, which yields a realistic time-scale for the arterial plaque growth, see [29]. The solid parameters are $\rho_s = 1 \text{ g/cm}^3$, $\mu_s = 10^4 \text{ dyne/cm}^2$, and $\lambda_s = 4 \cdot 10^4 \text{ dyne/cm}^2$. As an inflow boundary condition, we prescribe a pulsating velocity inflow profile on Γ_f^{in} given by

$$\mathbf{v}_f^{\text{in}}(t, x) = 30 \begin{pmatrix} \sin(\pi \frac{t}{\mathcal{P}})^2 (1 - x_2^2) \\ 0 \end{pmatrix} \text{ cm/s.} \quad (15)$$

Here, $\mathcal{P} = 1 \text{ s}$ is the period of a single heartbeat. The symmetry of the configuration can be exploited in order to reduce computational cost by simulating only on the lower half of the computational domain and imposing the symmetry condition $\mathbf{v}_f \cdot \vec{n} = 0$, $\tau^T \sigma_f \vec{n} = 0$ on the symmetry boundary Γ_{sym} , where τ denotes a tangential vector.

We discretize the FSI problem (variational formulation 1) in time using the backward Euler method. For discretizing the ODE growth model, we use the forward Euler method, which results in

$$c_s^n = c_s^{n-1} + \delta t \bar{\gamma}(\sigma_f^{WS}, c_s^{n-1}), \quad n = 1, \dots, N_f. \quad (16)$$

For spatial discretization, we use biquadratic (Q_2) equal-order finite elements for all variables and LPS stabilization [11] for the fluid problem. Our mesh, containing both fluid and solid, consists of 160 rectangular grid cells; this corresponds to a total of 3 157 degrees of freedom. The time-step sizes are chosen as $\delta \tau = 0.02 \text{ s}$ and $\delta t = 0.3 \text{ days}$; the tolerance for periodicity of the micro-scale problem as $\epsilon_p = 10^{-3}$. All the computational results have been obtained with the finite element library Gascoigne3d [12]. We use a fully monolithic approach for the FSI problem following Frei, Richter & Wick [30, 48].

In Fig. 3, we compare results obtained with the two-scale approach for different macro time-step sizes δt with the heuristic averaging approach outlined above; see variational formulation 2. We see that the pure averaging approach underestimates the growth significantly. Even a very coarse discretization of the macro-scale time interval $\delta t = 15 \text{ days}$ in the two-scale approach gives a much better approximation. We note, however, that too coarse time steps might introduce different issues, as the starting values $\mathbf{v}^{0,0}$ and $\mathbf{u}^{0,0}$ might not be good approximations for the periodic state on the micro scale anymore. In combination with a significant mesh deformation for $t \approx 300 \text{ days}$, this led to divergence of Algorithm 1 for an even coarser macro time-step size $\delta t = 30 \text{ days}$. For $\delta t \leq 15 \text{ days}$, a near-periodic state was reached in 2-3 iterations in each macro-step. A detailed convergence study for a very similar problem has been given in [28].

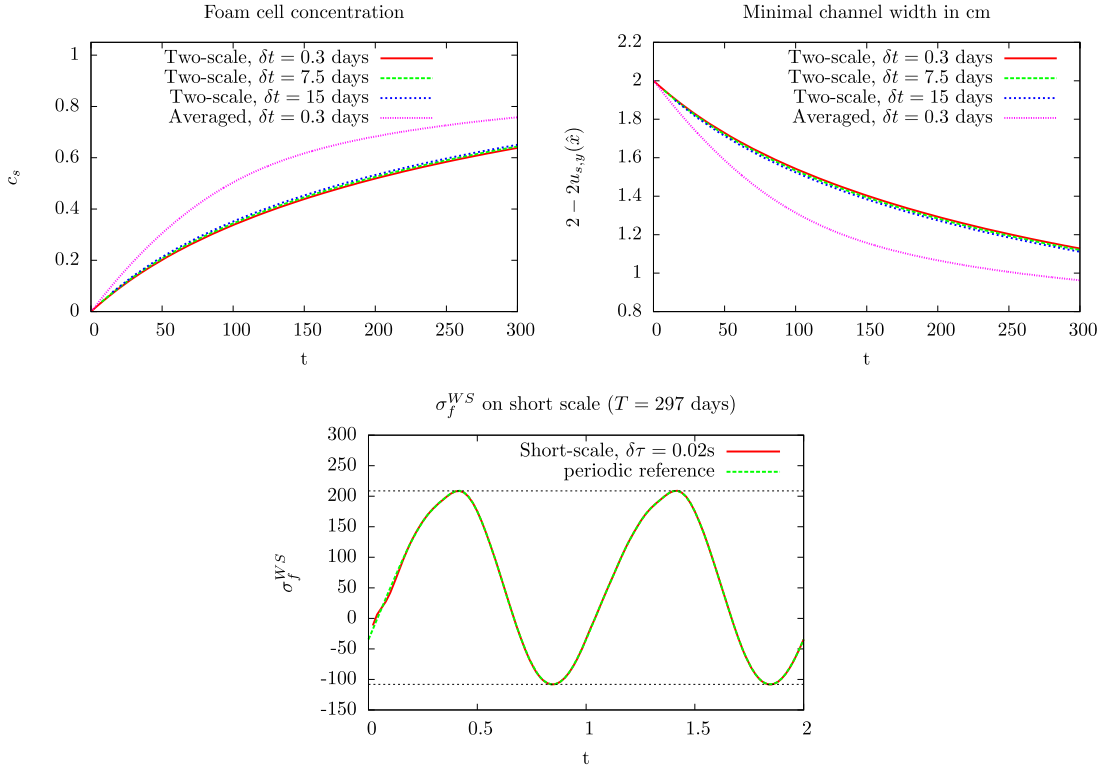


Fig. 3. Top: Comparison of a pure averaging approach with the two-scale approaches with different macro-scale time steps δT . **Top left:** Concentration c_s over time. **Top right:** Channel width over time. **Bottom:** Wall shear stress σ_f^{WS} over two periods of the micro scale at $t = 297$ days for $\delta t = 9$ days. The initial values are taken from the previous micro problem at $t = 288$ days.

At the bottom of Fig. 3, we show the $L^2(\Gamma)$ -norm of the wall shear stress over two periods, i.e., 2 seconds, of the micro problem for $\delta t = 0.02$ s and $\delta T = 9$ days. The initial values are taken from the periodic state of the micro problem at the previous macro time-step, i.e., 9 days before. We see that the wall shear stresses converge very quickly to the periodic state. An initial deviation of $6.28 \frac{\text{g cm}}{\text{s}^2}$ to the periodic reference solution at time $\tau = 0.02$ is reduced to $6 \cdot 10^{-4} \frac{\text{g cm}}{\text{s}^2}$ at time $\tau = 1$ within only one period.

4. Parallel time-stepping

The main cost in Algorithm 1 lies in the solution of the non-stationary micro problem in step 1.a, which needs to be solved in each time step of the macro problem. Considering a relatively coarse micro-scale discretization of $\delta t = 0.02$ s, as used in the previous section, 50 time steps are necessary to compute a single period of the heart beat. The simulation of two or more cycles might be necessary to obtain a near-periodic state in step 1.a of Algorithm 1; cf. the discussions in [28,29]. In a realistic scenario, each time step of the micro problem corresponds to the solution of a complex three-dimensional FSI problem, which makes already the solution of one micro problem very costly.

For this reason, parallelization needs to be exploited in different ways: First, one can make use of spatial parallelization using a scalable solver, for instance, based on domain decomposition or multigrid methods; see, e.g., [9,19,33,35,36,39,58]. Since the focus of this contribution is on parallelization in time, we will not discuss this aspect here. In particular, we expect that even when the speedup due to spatial parallelization saturates, the computing times for the whole plaque growth simulations will remain unfeasible. Therefore, we have to make use of an additional level of parallelization, that is, temporal parallelization of the macro-scale problem. In particular, we will use the parareal algorithm, which we will recap in the next subsection.

In order to motivate the algorithmical developments in this section, we will already present some first numerical results for the ODE growth model in eq. (5) within the section.

4.1. The parareal algorithm

First, the time interval of interest $[0, T_{\text{end}}]$ on the macro scale is divided into P sub-intervals $I_p = [T_{p-1}, T_p]$ of equal size, where




Fig. 4. Schematic representation of the parareal algorithm applied to the homogenized (two-scale) problem as shown in Fig. 2: micro-scale solution (black), two-scale solution (green), parareal iterates in iteration k (light grey) and $k + 1$ (dark grey), fine scale solution (red), and coarse scale solution (blue).

$$0 = T_0 < T_1 < \dots < T_p = T_{\text{end}}. \tag{17}$$

In order to define the parareal algorithm, suitable fine and coarse problems need to be introduced. Note again that we will apply the parareal algorithm only on the macro scale as introduced in section 3.1; hence, both the fine and the coarse scale of parareal correspond to the macro scale of the homogenization approach. The fine problem advances the growth variable c_s from time T_p to T_{p+1} by solving Algorithm 1 with a smaller time-step size δt (e.g., 0.3 days) on the corresponding fine time discretization of $[T_p, T_{p+1}]$:

$$T_p = t_{p,0} < t_{p,1} < \dots < t_{p,n_p} = T_{p+1}, \quad n_p = \frac{T_{p+1} - T_p}{\delta t}, \quad t_{p,q} := t_{p,n_p+q} = t_{p,q-1} + \delta t.$$

The fine propagator on a process p consists of a time-stepping procedure to advance $c_s(T_p)$ to $c_s(T_{p+1})$ with the fine time step δt . We write

$$c_s^{\text{fine}}(T_{p+1}) = \mathcal{F}(c_s^{\text{fine}}(T_p)).$$

The efficiency of the parareal algorithm depends strongly on the computational cost of the coarse propagator. In particular, it needs to be much cheaper than the fine problems since it is defined globally on $[0, T_{\text{end}}]$ and, hence, introduces synchronization. Thus, we use a large time-step δT and

$$T_p = \bar{T}_{p,0} < \dots < \bar{T}_{p,N_p} = T_{p+1}, \quad N_p = \frac{T_{p+1} - T_p}{\delta T}, \quad \bar{T}_{p,q} := \bar{T}_{p,N_p+q} = \bar{T}_{p,q-1} + \delta T.$$

For simplicity, we assume that both time-step sizes δt and δT are uniform throughout $[0, T_{\text{end}}]$. In order to keep the cost of the coarse propagator as low as possible, we will mainly focus on the case that the coarse time steps coincide with the P sub-intervals I_p in the numerical results, i.e., $N_p = 1$, such that the total number of coarse time steps

$$N_c := P \cdot N_p$$

is equal to P . We denote the coarse propagation from T_p to T_{p+1} by

$$c_s^{\text{coarse}}(T_{p+1}) = \mathcal{C}(c_s^{\text{coarse}}(T_p)).$$

We use capital letters T_p to denote the coarse discretization of $[0, T_{\text{end}}]$ into P parts and for the time-steps $T_{p,q}$ on the coarse level. By small letters $t_{p,q}$ we denote the finer discretization on each sub-interval; the two discretizations yield the first level (fine problem) and second level (coarse problem) of the parareal algorithm. Of course, it is also possible to employ coarse time steps which differ from the sub-intervals on the fine level, but for the sake of simplicity, we will not discuss this case in this work. In the example given above with $T_{\text{end}} = 300$ days, δT would be 30 days for $P = N_c = 10$, while δt is 0.3 days. On the micro scale, the times t_i and time-step size $\delta \tau$ are defined locally in $[t_i, t_i + 1s]$; in the example above, we had $\delta \tau = 0.02$ s. Note that the micro scale influences the parareal algorithm only indirectly due to the temporal homogenization approach.

Then, given an iterate $\bar{c}_s^{(k)}$ for some $k \leq 0$, the parareal algorithm computes \bar{c}_s^{k+1} by setting

$$\bar{c}_s^{(k+1)}(T_{p+1}) = \mathcal{C}(\bar{c}_s^{(k+1)}(T_p)) + \mathcal{F}(\bar{c}_s^{(k)}(T_p)) - \mathcal{C}(\bar{c}_s^{(k)}(T_p)) \quad \text{for } p = 0, \dots, P - 1. \tag{18}$$

Algorithm 2: Parareal algorithm

(I) **Initialization:** Compute $\{\bar{c}_s^{(0)}(T_p), \mathbf{w}^{(0)}(T_p)\}_{p=1}^P$ by means of Algorithm 1 with a coarse macro time-step size $\delta T := (T_p - T_{p-1})/N_p$. Set $k \leftarrow 0$

(II) **while** $|c_s^{(k+1),\text{fine}}(T_{\text{end}}) - c_s^{(k),\text{fine}}(T_{\text{end}})| > \epsilon_{\text{par}}$ **do**

(II.a) **Fine problem:**

for $p = 1, \dots, P$ **do**

(i) Initialize $c_s^{(k+1),\text{fine}}(T_p) = \bar{c}_s^{(k)}(T_p)$, $\mathbf{w}^{(k+1)}(T_p) = \mathbf{w}^{(k)}(T_p)$

(ii) Compute $\{(c_s^{(k+1),\text{fine}}(t_{p,q}), \mathbf{w}^{(k+1)}(t_{p,q}))_{q=1}^{n_p}\}$ by Algorithm 1 with fine time-step size δt and set $\mathcal{F}(\bar{c}_s^{(k)}(T_p)) = c_s^{(k+1),\text{fine}}(t_{p,n_p})$

(II.b) **Coarse problem**

for $p = 1, \dots, P$ **do**

(i) Compute $\mathcal{C}(\bar{c}_s^{(k+1)}(T_p))$ by solving one time step of Algorithm 1 with time-step size $\delta T = (T_p - T_{p-1})/N_p$

(ii) Parareal update

$\bar{c}_s^{(k+1)}(T_{p+1}) = \mathcal{C}(\bar{c}_s^{(k+1)}(T_p)) + \mathcal{F}(\bar{c}_s^{(k)}(T_p)) - \mathcal{C}(\bar{c}_s^{(k)}(T_p))$.

$k \leftarrow k + 1$

This can be seen as a predictor-corrector scheme, where the coarse predictor $\mathcal{C}(\bar{c}_s^{(k+1)}(T_p))$ is corrected by fine-scale contributions that depend only on the previous iterate $\bar{c}_s^{(k)}$; thus, the fine problems can be computed fully in parallel. A schematic illustration of the parareal algorithm is given in Fig. 4.

Let us analyze the application of eq. (18) to the two-scale problem (Algorithm 1) in more detail. The first term in eq. (18) requires the solution of one micro problem and an update of the foam cell concentration (by eq. (5) or eq. (11)) in each coarse time step $T_p \rightarrow T_{p+1}$. Within the fine-scale propagator (second term in eq. (18)) $n_p = \lceil N_f/P \rceil$ time-steps need to be computed per process, where $\lceil g \rceil$ denotes the next-biggest natural number to g ($\hat{=} \text{ceil}(g)$) and N_f is the number of macro time steps; see eq. (13). Each time step requires the solution of one micro problem and an update of the foam cell concentration. The last term in eq. (18) has already been computed in the previous iteration (compare the first term on the right-hand side of the same equation) and thus introduces no additional computational cost. The algorithm is summarized in Algorithm 2.

We use the FSI variables $w^{(k)}(T_p)$ from the previous parareal iteration as initial values in step (II.a)(i). The initial values $\bar{c}_s^{(k)}(T_p)$ in step (II.a)(i) are taken from step (II.b) of the coarse problem. Moreover, to initialize the variables $\{\bar{c}_s^{(0)}(T_p), \mathbf{w}_s^{(0)}(T_p)\}_{p=1}^P$ before the first parareal iteration, we apply the coarse propagator once with a large time-step size $\delta T = (T_p - T_{p-1})/N_p$.

For the ODE growth model, we use the value of c_s^{fine} at the end time T_{end}

$$|c_s^{(k+1),\text{fine}}(T_{\text{end}}) - c_s^{(k),\text{fine}}(T_{\text{end}})| \leq \epsilon_{\text{par}}, \quad (19)$$

as the stopping criterion for the parareal algorithm.

4.1.1. Parallelization approaches

As mentioned earlier, in this work, we focus on studying the effectivity and efficiency of our approach by investigating the convergence of the parareal algorithm based on a serial implementation. Since the computational cost of the micro FSI problems is very high, we expect the communication cost to be negligible; cf. section 4.1.2 as well as sections 5.2 and 5.3 for a more detailed discussion of the computational cost for the ODE and the PDE growth model, respectively. Even though our implementation used in the numerical examples below is only serial, we want to briefly discuss two potential parallelization approaches for a parallel implementation of Algorithm 2: *master-slave parallelization* and *distributed parallelization*; cf. Fig. 5 for a sketch of both approaches for our parareal algorithm. Note that, as also indicated by Fig. 5, we always assume a one-to-one correspondence of fine problems and processes in this paper.

In the master-slave parallelization approach, the coarse problem is assigned to a single process and the local problems on the fine level are distributed among the remaining processes. Each slave process p has to communicate $\mathcal{F}(\bar{c}_s^{(k)}(T_p))$ to the master after computation of the fine problem for the parareal update (step II.a in Algorithm 2), and the master process has to communicate $\bar{c}_s^{(k+1)}(T_p)$ back after the parareal update (step II.b in Algorithm 2); this corresponds to all-to-one and one-to-all communication patterns.

In the distributed parallelization approach, the parallelization of the coarse problem is different. In particular, time intervals are assigned to the different processes, and each process computes both the fine problem and the part of the coarse problem on those time intervals; for simplicity, we do not discuss the case that the partition of the coarse problem onto the processes is different from the local (fine) problems, which would also be possible. As can be seen in Fig. 5, a major benefit of this approach is the different communication pattern: instead of all-to-one and one-to-all communication, each process sends the parareal update $\bar{c}_s^{(k+1)}(T_p)$ to the next process in line. Note that, in terms of the coarse problem, the processes can only perform their computations in serial, whereas the local problems can, again, be computed concurrently.

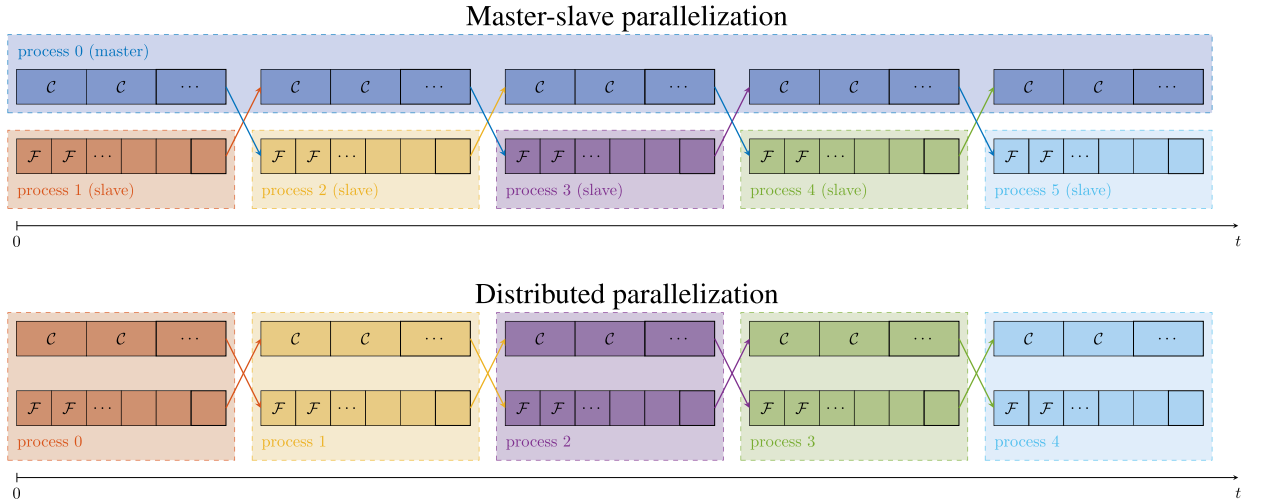


Fig. 5. Master-slave and distributed parallelization schemes for the parareal approach. In the master-slave case, the coarse propagation is performed on a separate process (blue) and information is communicated between master and slave. In the distributed case, the coarse propagation and the fine propagation is performed on the same processes. Communication is indicated by an arrow in color of the sender.

In some implementations of two-level methods, the coarse problem is assigned to one of the processes dealing with the local problems instead of an additional process; if the memory and computational cost of the coarse problem is negligible, this can be beneficial since all available processes can be employed for the local problems. It can be seen as a mixture of the master-slave and distributed parallelization approaches. In our case, this might not be feasible since the coarse problem is defined on the same spatial mesh as the local problems, leading to significant memory cost. Furthermore, more sophisticated approaches such as task-based scheduling [4] can be employed for the parallelization.

An investigation of the differences in performance of the different approaches can only be performed using an actual parallel implementation. This is out of the scope of this paper, and therefore, we will leave it to future work.

4.1.2. Computational costs

Even in the simplified two-dimensional configuration considered in this work, it suffices to count the number of micro problems to be solved to estimate the computational cost of corresponding parallel computations and to compare different algorithms. Remember that, in each micro problem, $r_p \cdot \frac{1}{\delta T}$ FSI systems need to be solved, where r_p is the number of cycles required to reach a near-periodic state. In the example considered above, these were at least $2 \cdot \frac{1}{0.02} = 100$ costly FSI problems. Hence, our discussion in the following sections will be based on the assumption that the computational cost for all other steps of the parareal algorithm as well as the communication between the processes can be neglected. This is particularly obvious for the ODE growth model in eq. (5). For the PDE growth model in eq. (11), we will discuss the computational and communication cost in more detail in sections 5.2 and 5.3.

To analyze the computational cost of Algorithm 2 under this assumption, let us denote the total number of iterations of the parareal algorithm by k_{par} . As discussed in sections 4.1 and 4.1.1, we assume that the number of coarse time steps is the same as the number of fine problems and processes, respectively. Therefore, we need to solve N_c micro problems in step I, $k_{\text{par}} \cdot \lceil N_f/P \rceil$ micro problems on each of the P processes in step II.a (ii), and $k_{\text{par}} \cdot N_c$ micro problems within the coarse propagation in step II.b (i). This corresponds to the solution of

$$\underbrace{k_{\text{par}} \cdot \lceil N_f/P \rceil}_{\text{fine level (} P \text{ parallel processes)}} + \underbrace{(k_{\text{par}} + 1) \cdot N_c}_{\text{coarse level (1 serial process)}} \quad (20)$$

serial micro problems; this is the case for both parallelization approaches discussed in section 4.1.1. If the coarse time step δT is chosen independently of T , N_c is fixed and the computational cost tends to saturate for large P (at least if we assume that the number of required parareal iterations k_{par} is independent of P). The cheapest possible coarse propagator, on the other hand, is to use one coarse time step per coarse propagation, i.e., $N_p = 1$ and $N_c = P$. In this case the computational time increases for large P , if we assume that the number of parareal iterations k_{par} is independent of P . The minimum computational time is attained for $P \approx \sqrt{N_f}$.

As a serial implementation requires the solution of N_f micro problems, the speed-up of the proposed parareal algorithm is given by

$$\text{speedup}(P) = \frac{N_f}{k_{\text{par}} \cdot \lceil N_f/P \rceil + (k_{\text{par}} + 1) \cdot N_c} \approx \frac{1}{\frac{k_{\text{par}}}{P} + (k_{\text{par}} + 1) \frac{N_c}{N_f}}, \quad (21)$$

if we assume a perfect load balancing on the fine processes. This is a special case of the standard analysis for the speed-up of the parareal algorithm; see, e.g., [52].

The aim of this work is to show potential for the parareal algorithm as a parallel time-stepping method for plaque growth simulations. A final assessment of the parallelization capabilities can, of course, only be made based on computing times of an actual parallel implementation. We will leave this to future work.

4.1.3. Convergence theory

Let us briefly recap the standard convergence theory for the parareal algorithm; see, e.g., [32]. For this purpose, let us consider the simpler ODE model

$$\partial_t c_s = \gamma(c_s), \quad c_s(0) = 0, \tag{22}$$

and its discretization by the explicit Euler method

$$c_s^n = c_s^{n-1} + \delta T \gamma(c_s^{n-1}), \quad c_s^0 = 0. \tag{23}$$

We assume that the function γ and its derivative are Lipschitz continuous in c_s , i.e., we assume

$$|\gamma(c_s^1) - \gamma(c_s^2)| \leq L|c_s^1 - c_s^2|, \quad \left| \frac{d}{dt} \gamma(c_s^1) - \frac{d}{dt} \gamma(c_s^2) \right| \leq L|c_s^1 - c_s^2| \quad \text{for all } c_s^1, c_s^2 \in \mathbb{R} \tag{24}$$

for some $L > 0$. As in [32], we also assume for simplicity that the fine-scale propagator advances the ODE exactly, i.e.,

$$c_s(T_n) = \mathcal{F}(c_s(T_{n-1})) = c_s(T_{n-1}) + \int_{T_{n-1}}^{T_n} \gamma(c_s(s)) ds. \tag{25}$$

This is motivated by the fact that the time discretization error in the fine propagator is typically small compared to the coarse propagator. The following convergence result is shown by Gander & Hairer in [32]:

Lemma 1. *Let $c_s^* \in C^1(0, T_{\text{end}})$ be the exact solution of eq. (22), and let $\{\bar{c}_s^{(k)}(T_n)\}_{n=1}^P$ be the k -th iterate of the parareal algorithm in eq. (18) using the forward Euler method in eq. (23) as the time discretization. Under the assumptions made above, it holds for the error $e_n^{(k)} = |\bar{c}_s^{(k)}(T_n) - c_s^*(T_n)|$ and $k \in \mathbb{N}_0, n \in \{1, \dots, P\}$ that*

$$e_n^{(k)} \leq (cT_n)^{k+1} \beta^{n-k-1} \delta T^{k+1} \max_{t \in [0, T_{\text{end}}]} |\partial_t c_s^*(t)|, \tag{26}$$

with a constant $c > 0$ and $\beta = 1 + L\delta T$.

Remark 1. The assumptions in eq. (24) cannot be easily verified for the ODE model in eq. (5), as the fluid forces and $\bar{\gamma}(\sigma_f^{WS})$ depend in a highly nonlinear way on c_s . In the numerical examples given below, we will, however, observe that the convergence behavior is similar to the one shown in Lemma 1.

4.1.4. Numerical results

We use, again, the simple example described in section 3.2 with the ODE growth model in eq. (5) and set $\epsilon_{\text{par}} = 10^{-3}$. Results for $P = N_c = 10$ and $\delta t = 0.3$ days are shown in Fig. 6, where the first 3 iterates are compared against a reference solution $(c_s^*, \mathbf{u}_s^*, \mathbf{v}_f^*)$, which was computed by a standard serial time-stepping scheme, as in section 3.2. We observe fast convergence towards the reference curve in all three quantities of interest. The stopping criterion in eq. (19) was satisfied after 4 iterations of the parareal algorithm.

In Table 1, we show the deviations in the foam cell concentration $|c_s^{(k), \text{fine}}(T_{\text{end}}) - c_s^*(T_{\text{end}})|$ at final time T_{end} after each iteration of the parareal algorithm for $P = N_c = 10, 20, 30, 40$ and 50 processes. We observe that the number of iterations k_{par} decreases from four to three for $P \geq 20$. This is due to the fact that the coarse problem is solved with the smaller time-step size $\delta T = \lceil \frac{T_{\text{end}}}{P} \rceil$, which makes the coarse problem more expensive but also more accurate. In particular, we obtain a better approximation for the coarse values $\bar{c}_s^{(k)}(T_p)$ that are used as initial values in the fine problems in the next iterate. We can also observe that, for fixed P , the error decreases at least by a factor $\frac{1}{P}$ in each parareal iteration. This is in agreement with Lemma 1, which predicts (for a simpler model problem) a reduction factor $c \cdot \delta T = c \cdot \lceil \frac{T_{\text{end}}}{P} \rceil$ for some constant $c > 0$, compared to the previous iterate.

Since the number of parareal iterations k_{par} is constant for all cases with $P \geq 20$, we get the lowest computational cost for $P = 30$, which is close to $\sqrt{N_f} \approx 31.6$; cf. the discussion in section 4.1.2. For $P = 30$, $3 \cdot 34 = 102$ micro problems need to be solved on each process (step II.a) and $4 \cdot 30 = 120$ micro problems within the coarse propagators (steps I and II.b), i.e., 222 micro problems in total. Compared to a serial time-stepping scheme with the same fine-scale time-step size, which requires the solution of $N_f = 1000$ serial micro problems, this corresponds to a speed-up of $\frac{1000}{222} \approx 4.5$; see also eq. (21).

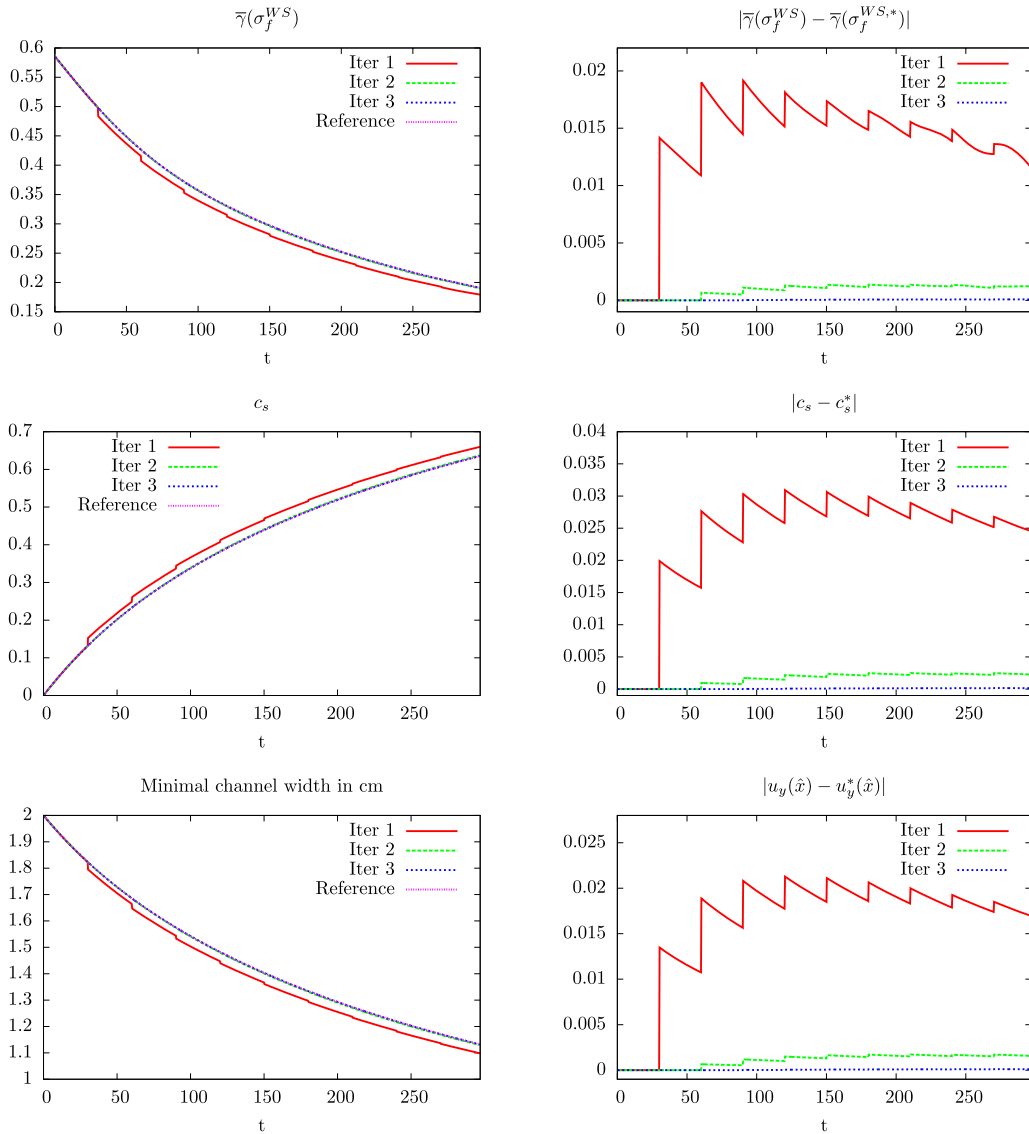


Fig. 6. Behavior of the first three iterates of the parareal algorithm in the first numerical example (ODE growth) for $P = 10$. A description of the problem is given in section 4.1.4, numerical values in Table 1. The iterates converge very quickly towards the reference solution. **Top:** Growth function and error w.r.t. the reference solution $\sigma_f^{WS}(\mathbf{v}_f^*)$ over time t . **Center:** $c_s^{(k),\text{fine}}$ and error w.r.t. the reference solution c_s^* over time. **Bottom:** Channel width over time and error $|u_{s,y}(\hat{x}) - u_{s,y}^*(\hat{x})|$ w.r.t. the vertical component of the reference solution \mathbf{u}_s^* at the narrowest point \hat{x} over time.

Table 1

Errors $|c_s^{(k),\text{fine}}(T_{\text{end}}) - c_s^*(T_{\text{end}})|$ for $P = 10, \dots, 50$ for the parareal algorithm (Algorithm 2) in the first numerical example (ODE growth). For comparison, the reference value $c_s^*(T_{\text{end}}) = 0.63831273\dots$ resulting from a serial time-stepping with the same fine-scale time-step size δt is taken. The time measure in terms of the number of serial micro problems (# mp) as well as speedup and efficiency compared to the reference computation (right column) are shown; the stopping criterion $|c_s^{(k+1),\text{fine}}(T_{\text{end}}) - c_s^{(k),\text{fine}}(T_{\text{end}})| < \epsilon_{\text{par}} = 10^{-3}$ is used; best numbers marked in **bold face**.

k	$P = 10$	$P = 20$	$P = 30$	$P = 40$	$P = 50$	Ref. (serial)
1	$2.21 \cdot 10^{-2}$	$1.19 \cdot 10^{-2}$	$8.12 \cdot 10^{-3}$	$5.83 \cdot 10^{-3}$	$5.46 \cdot 10^{-3}$	0
2	$2.24 \cdot 10^{-3}$	$5.63 \cdot 10^{-4}$	$2.62 \cdot 10^{-4}$	$1.37 \cdot 10^{-4}$	$8.61 \cdot 10^{-5}$	-
3	$1.42 \cdot 10^{-4}$	$2.02 \cdot 10^{-5}$	$6.53 \cdot 10^{-6}$	$2.32 \cdot 10^{-6}$	$8.30 \cdot 10^{-7}$	-
4	$5.76 \cdot 10^{-6}$	-	-	-	-	-
# mp	450	230	222	235	260	1 000
speedup	2.2	4.3	4.5	4.3	3.8	1.0
efficiency	22%	22%	15%	11%	8%	100%

Table 2

Errors $|\bar{c}_s^{(k)}(T_{\text{end}}) - c_s^*(T_{\text{end}})|$ for $P = 10, \dots, 50$ for the parareal algorithm (Algorithm 2) in the first numerical example (ODE growth). The time measure in terms of the number of serial micro problems (# mp) as well as speedup and efficiency compared to the reference computation (right column) are shown; the stopping criterion $|\bar{c}_s^{(k+1)}(T_{\text{end}}) - \bar{c}_s^{(k)}(T_{\text{end}})| < \epsilon_{\text{par}} = 10^{-3}$ is used; best numbers marked in **bold face**.

k	P = 10	P = 20	P = 30	P = 40	P = 50	Ref. (serial)
1	$2.59 \cdot 10^{-3}$	$5.97 \cdot 10^{-4}$	$3.29 \cdot 10^{-4}$	$1.42 \cdot 10^{-4}$	$8.95 \cdot 10^{-5}$	0
2	$2.81 \cdot 10^{-3}$	$2.04 \cdot 10^{-5}$	$6.59 \cdot 10^{-6}$	$2.36 \cdot 10^{-6}$	$8.43 \cdot 10^{-7}$	-
3	$6.37 \cdot 10^{-6}$	-	-	-	-	-
4	$1.73 \cdot 10^{-7}$	-	-	-	-	-
# mp	450	160	158	170	190	1 000
speedup	2.2	6.25	6.3	5.9	5.3	1.0
efficiency	22%	31%	21%	15%	11%	100%

As is usual for two-level methods, for larger numbers of processes (P), the cost of the coarse problems gets dominant. For $P = 40$, we solve, for example, $3 \cdot 25 = 75$ micro problems in the fine propagator compared to $4 \cdot 40 = 160$ micro problems for the coarse propagator; this results in a total of 235 micro problems and a speedup of 4.3.

In order to relate the speed-up to the computational resources used, we define the computational efficiency by the ratio $\frac{\text{speedup}}{P}$ as a second measure to quantify the results. By definition, the serial computation always has the best efficiency, as it converges within one sweep through all time steps. Even if the parareal algorithm would also converge within one iteration, the parallel computations come with additional costs since they require the solution of the coarse problem.

In fact, using the serial computation as a reference in the comparison in Table 1, the efficiency is best (22%) for $P = 10$ or $P = 20$; it deteriorates for larger numbers of processes due to the increasing cost of the coarse problems. The low efficiency is a typical observation for parallel time integration methods, such as the parareal method and can be explained as follows: For $P = 10$, for instance, 4 parareal iterations are required and each iteration requires the solution of all 1 000 micro problems on the fine scale on the full time interval $[0, T_{\text{end}}]$. Hence, a total of 4 000 micro problems has to be solved. Assuming that the time to solve a micro problem is approximately constant, we can not expect efficiencies above 25%. Due to the additional cost of the coarse propagator, the efficiency is reduced to 22%. However, since we can always compute 10 micro problems on the fine scale at the same time, we can still obtain a speedup of 2.2.

To increase the efficiency, we could take the value $\bar{c}_s^{(k)}(T_{\text{end}})$ computed in eq. (18) instead of $c_s^{(k),\text{fine}}(T_{\text{end}})$ in the stopping criterion

$$|\bar{c}_s^{(k+1)}(T_{\text{end}}) - \bar{c}_s^{(k)}(T_{\text{end}})| < \epsilon_{\text{par}}. \quad (27)$$

In particular, at the end time T_{end} we expect that this has a higher accuracy compared to the fine-scale variable $c_s^{(k),\text{fine}}(T_{\text{end}})$. However, the values $\bar{c}_s^{(k+1)}(t_i)$ are only available at the coarse grid points $t_i \in \{T_p\}_{p=1}^P$. If one is interested in foam cell concentrations at intermediate time steps, a stopping criterion based on $c_s^{(k),\text{fine}}(t_i)$ needs to be used, as in eq. (19).

The errors concerning $\bar{c}_s^{(k)}(T)$ and the number of micro problems to be solved using the stopping criterion eq. (27) are given in Table 2. We observe that for $P \geq 20$ this stopping criterion was satisfied already after $k_{\text{par}} = 2$ iterations. For $P = 30$ we get again the lowest computational cost with a speed-up of around 6.33 compared to a serial computation. Furthermore, the best efficiency is obtained here with $P = 20$, which is due to the lower iteration count compared to the case $P = 10$.

4.2. Variants with cheaper coarse-scale computations

In the parareal algorithm introduced above, a further improvement of the computational cost is not possible for $P \gtrsim \sqrt{N_f}$ processes in the case $P = N_c$ due to the (increasing) cost of the coarse-scale propagators. In fact, these get dominant compared to the fine-scale contributions for $P \gtrsim \sqrt{N_f}$. In this section, we will discuss approximate coarse-scale propagators, where no additional computations of micro problems are needed; hence, they are computationally cheaper.

4.2.1. Heuristic averaging of the FSI problems

As a first variant, we use the heuristic averaging approach mentioned in the beginning of section 3.1 for the coarse-scale propagation in steps (I) and (II.b)(i) of Algorithm 2. The solution of the stationary FSI problem in variational formulation 2 is much cheaper compared to ≥ 100 time steps of a non-stationary FSI problem (approx. by a factor 100), and thus, its computational cost is neglected in the following discussion. Hence, the computational cost is reduced from $k_{\text{par}} \cdot \lceil N_f/P \rceil + (k_{\text{par}} + 1) \cdot P$ to $k_{\text{par}} \cdot \lceil N_f/P \rceil$ micro problems. In Fig. 7, we illustrate the error in the foam cell concentration for $P = 50$ (left), and give the error $|c_s^{(k),\text{fine}}(T_{\text{end}}) - c_s^*(T_{\text{end}})|$ in each iteration k for $P = 30$ and $P = 50$ (right). We observe a much slower convergence compared to the standard parareal algorithm used above. This could already be expected from the results in Fig. 3, where we saw that the heuristic averaging is not a good approximation. Using the stationary problem for the coarse propagator, 8 parareal iterations are necessary until the stopping criterion is satisfied for $P = 30$ and $P = 50$ compared to 3 iterations in the standard parareal algorithm.

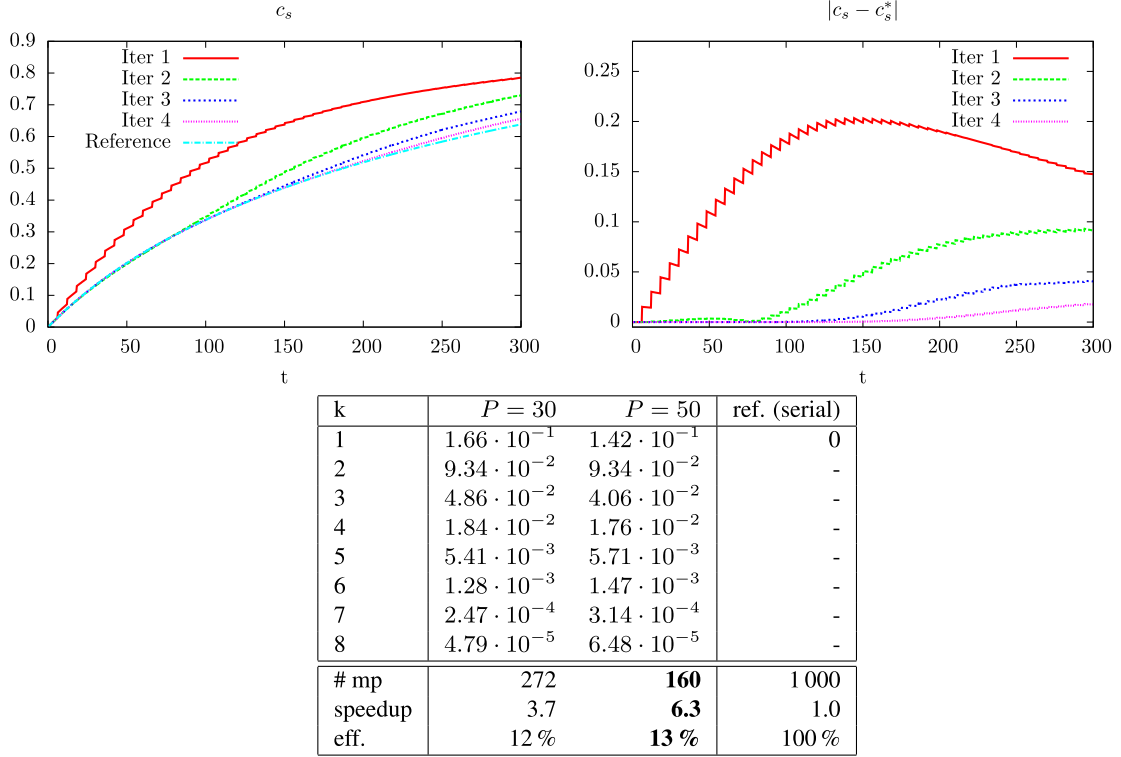


Fig. 7. Behavior of the iterates of the parareal algorithm, when the stationary FSI problem (variational formulation 2) is used in Step (I) and (II.b)(i) of Algorithm 2. A description of the computation is given in section 4.2. Convergence to the reference solution is much slower compared to standard parareal (see Fig. 6). **Top:** Evolution of $c_s^{(k),\text{fine}}(t)$ and error $|c_s^{(k),\text{fine}}(t) - c_s^*(t)|$ of the foam cell concentration over time t for $P = 50$ and $k = 1, 2, 3, 4$. **Right:** Errors $|c_s^{(k),\text{fine}}(T_{\text{end}}) - c_s^*(T_{\text{end}})|$ after each iteration for $P = 30$ and $P = 50$; total number of micro problems (# mp) solved in parallel as well as speedup and efficiency compared to the reference computation (right column); best numbers marked in **bold face**.

In terms of the computing time, $8 \cdot 34 = 272$ and $8 \cdot 20 = 160$ micro problems are necessary for the cases $P = 30$ and $P = 50$, respectively. While for $P = 30$ this is worse compared to the standard parareal algorithm, this is an improvement of 67% for $P = 50$; cf. Table 1. As for the standard parareal algorithm, convergence is obtained faster if we consider the stopping criterion in eq. (27) based on the parareal iterates $\bar{c}_s^{(k)}(T_{\text{end}})$; however, 7 iterations are still needed. The resulting computational cost is $7 \cdot 34 = 238$ micro problems for $P = 30$ and $7 \cdot 20 = 140$ micro problems for $P = 50$. Compared to 158 micro problems for $P = 30$ and 190 micro problems for $P = 50$ in the standard parareal algorithm (cf. Table 2), we see only a small but no significant improvement. We will thus consider another, more promising approach to compute the required growth values $\gamma(\sigma_f^{WS})$ for the coarse propagator in the following subsection.

Remark 2. As mentioned before, parallelization in space is generally more efficient than parallelization in time. Hence, complex three-dimensional FSI simulations already have a significant potential for parallelization. Additionally using temporal parallelization with large values of P , such as $P > 50$, may only be reasonable on very large supercomputers.

4.2.2. Re-usage of computed growth values

As a second approach we will consider the re-use of growth values $\bar{\gamma}(\sigma_f^{WS}(t_{p,i}))$ computed in the fine-scale propagator on the coarse scale. Therefore, we store all values $\bar{\gamma}_{p,n_p+i} := \bar{\gamma}(\sigma_f^{WS}(t_{p,i}))$ computed on the fine scale on all processes $p = 1, \dots, P$ for all time-steps $i = 1, \dots, n_p$; see step 1.b of Algorithm 1. These can be used in the coarse propagator (step II.b) in the same parareal iteration instead of computing new growth values there.

We introduce the following notations for the coarse resp. fine propagators that start from $c_{s,n-1}$ using certain growth values $\gamma(T_{n-1})$ (that might now differ from $\bar{\gamma}(\sigma_f^{WS}(c_{s,n-1}))$):

$$c_{s,n}^{\text{coarse}} = \mathcal{C}(I_n, c_{s,n-1}, \gamma(T_{n-1})), \quad c_{s,n}^{\text{fine}} = \mathcal{F}(I_n, c_{s,n-1}, \gamma(T_{n-1})),$$

where $I_n = [t_{n-1}, t_n]$. After an initialization step, the modified parareal iteration is defined by the following formula for $n = 1, \dots, P$ and $k \geq 1$:

$$\bar{c}_{s,n}^{(k)} = \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(k)}, \bar{\gamma}(c_{s,n-1}^{(k),\text{fine}})) + \mathcal{F}(I_n, \bar{c}_{s,n-1}^{(k-1)}, \bar{\gamma}(\bar{c}_{s,n-1}^{(k-1)})) - \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(k-1)}, \bar{\gamma}(c_{s,n-1}^{(k-1),\text{fine}})), \tag{28}$$

Algorithm 3: Parallel Time-Stepping with Re-Usage of Growth Values

(I) **Initialization:** Compute $\{\bar{c}_s^{(0)}(T_p), \bar{\mathbf{w}}^{(0)}(T_p)\}_{p=1}^P$ by means of Algorithm 1 with a coarse macro time-step size $\delta T = T_{p+1} - T_p$ on the master process. Set $k \leftarrow 0$

(II) **while** $|c_s^{(k+1),\text{fine}}(T_{\text{end}}) - c_s^{(k),\text{fine}}(T_{\text{end}})| > \epsilon_{\text{par}}$ **do**

(II.a) **Fine problem:**

for $p = 1, \dots, P$ **do**

(i) Initialize $c_s^{(k+1),\text{fine}}(T_p) = \bar{c}_s^{(k)}(T_p)$, $\mathbf{w}^{(k+1)}(T_p) = \mathbf{w}^{(k)}(t_{p-1}, N_p)$

(ii) Compute $\{c_s^{(k+1),\text{fine}}(t_{p,q}), \mathbf{w}^{(k+1)}(t_{p,q})\}_{q=1}^{n_p}$ by Algorithm 1 with fine time step δt

(iii) Store the resulting growth functions $\bar{\gamma}_{p-n_p+q} := \bar{\gamma}(\sigma_f^{WS})(t_{p,q})$ at all fine points $t_{p,q}$, $q = 1, \dots, n_p$ as well as $\mathbf{w}(t_{p,n_p})$ at the last time-step.

(II.b) **Coarse problem:**

(i) **for** $j = 1, \dots, N_f$ **do**

Compute $\bar{c}_s^{(k+1)}(t_j)$ by advancing the ODE eq. (5) resp. solving the PDE in eq. (11) using the growth values $\bar{\gamma}_j$ computed in (II.a).

where

$$c_{s,n-1}^{(k),\text{fine}} := \mathcal{F}(I_{n-1}, \bar{c}_{s,n-2}^{(k-1)}, \bar{\gamma}(\bar{c}_{s,n-2}^{(k-1)})), \tag{29}$$

and $c_{s,n}^{(0),\text{fine}} = \bar{c}_{s,n}^{(0)}$. Moreover, we set $c_{s,0}^{(k),\text{fine}} = \bar{c}_{s,0}^{(k)} = 0$ for all k .

As no new micro problems need to be solved and approximations of the growth values $\bar{\gamma}_j$ are available for all fine time steps $j = 1, \dots, N_f$, the coarse propagator can now even use the fine-scale time-step δt . The only additional cost is to advance the foam cell concentration by eq. (5) resp. eq. (18) on the coarse level. This cost is clearly negligible for the ODE growth model in eq. (5). The additional cost in case of the PDE model eq. (18) will be discussed in section 5. The resulting algorithm is given as Algorithm 3.

The growth values employed for the re-usage are available for all fine-scale time steps, and hence, we are able to perform the coarse propagation efficiently on the fine scale; cf. the discussion on the computational cost later in this section as well as for the PDE growth model in section 5.3. Nonetheless, the expected accuracy of the re-usage approach is still lower compared to the standard parareal algorithm. This is because the growth values on the fine scale have been computed using $c_{s,n-1}^{(k),\text{fine}}$, which depends on the previous iterate $\bar{c}_{s,n-2}^{(k-1)}$ (see eq. (29)), whereas the coarse propagator in the standard parareal algorithm already uses the more accurate $\bar{c}_{s,n-1}^{(k)}$ from the current iteration.

Computational costs Only the coarse step in the initialization (I) has to be carried out without re-usage and comes with a computational cost of N_c micro problems. On the other hand, step (II.b) does not require the solution of any micro problems. In terms of micro problems to be solved, the computational cost of step (II.a) is the same as in the standard parareal algorithm (Algorithm 2). Altogether, the number of micro problems to be solved in k_{par} iterations of Algorithm 3 is

$$k_{\text{par}} \cdot \lceil N_f / P \rceil + N_c.$$

Again, we obtain a saturation of the cost for large P , but the cost for large P is by a factor k_{par} smaller. Moreover, for $N_c = P$, the choice $P \approx \sqrt{k_{\text{par}} N_f}$ would be optimal if the number of iterations k_{par} was independent of the number of processes P .

The speed-up compared to a serial computation is given by

$$\text{speedup}(P) = \frac{N_f}{k_{\text{par}} \cdot \lceil N_f / P \rceil + N_c} \approx \frac{1}{\frac{k_{\text{par}}}{P} + \frac{N_c}{N_f}}, \tag{30}$$

assuming, again, a perfect load balancing of the micro problems.

The communication cost of the approach strongly depends on the employed parallelization scheme; cf. section 4.1.1. In the distributed parallelization approach, the growth values $\bar{\gamma}(\sigma_f^{WS})(t_{p,i})$, which are needed in the coarse propagator, do not have to be communicated. This is because the same intervals in fine and coarse problems are computed on the processes. In the master-slave approach, however, the values have to be communicated from each slave process to the master. Using the ODE growth model, these are N_f scalar values in total (in our example $N_f = 1000$). Even for the master-slave communication scheme, it is thus reasonable to assume that the computational cost of communication is still negligible compared to the solutions of the micro problems. In the case of a PDE growth model, the additional cost for communication will be discussed in section 5.

Theoretical convergence analysis We extend the convergence results discussed in section 4.1.3 for the model problem in eq. (22) for the re-usage algorithm in eq. (28) under the assumptions made in section 4.1.3.

The following result is a direct consequence of Lemma 3, which is shown in the appendix.




Fig. 8. Error in the concentration $|c_s^{(k),\text{fine}}(t) - c_s^*(t)|$ over time t in the first 3 iterates of the parareal algorithm with re-usage of growth values (Algorithm 3) for $P = 10$ (left) and $P = 60$ (right) and the first numerical example (ODE growth model).

Table 3

Errors $|c_s^{(k),\text{fine}}(T_{\text{end}}) - c_s^*(T_{\text{end}})|$ for $P = 10, \dots, 70$ for the parareal Algorithm 3 (re-use of growth values) and the first numerical example (ODE growth model). Value $c_s^*(T_{\text{end}})$ of a serial reference computation for comparison. The time measure in terms of the number of micro problems (# mp) as well as speedup and efficiency compared to the reference computation (right column) are shown; the stopping criterion $|c_s^{(k+1),\text{fine}}(T_{\text{end}}) - c_s^{(k),\text{fine}}(T_{\text{end}})| < \epsilon_{\text{par}} = 10^{-3}$ is used; best numbers marked in **bold face**.

k	$P = 10$	$P = 20$	$P = 30$	$P = 40$	$P = 50$	$P = 60$	$P = 70$	Ref.
1	$2.56 \cdot 10^{-2}$	$1.10 \cdot 10^{-2}$	$7.48 \cdot 10^{-3}$	$5.46 \cdot 10^{-3}$	$4.23 \cdot 10^{-3}$	$3.63 \cdot 10^{-3}$	$3.12 \cdot 10^{-3}$	0
2	$7.78 \cdot 10^{-3}$	$4.48 \cdot 10^{-3}$	$3.16 \cdot 10^{-3}$	$2.36 \cdot 10^{-3}$	$1.99 \cdot 10^{-3}$	$1.61 \cdot 10^{-3}$	$1.41 \cdot 10^{-3}$	-
3	$1.73 \cdot 10^{-3}$	$1.20 \cdot 10^{-3}$	$9.02 \cdot 10^{-4}$	$6.92 \cdot 10^{-4}$	$6.34 \cdot 10^{-4}$	$5.14 \cdot 10^{-4}$	$4.32 \cdot 10^{-4}$	-
4	$2.32 \cdot 10^{-4}$	$2.28 \cdot 10^{-4}$	$1.87 \cdot 10^{-4}$	$1.31 \cdot 10^{-4}$	$1.26 \cdot 10^{-4}$	$1.10 \cdot 10^{-4}$	$9.98 \cdot 10^{-5}$	-
5	$2.19 \cdot 10^{-5}$	$4.71 \cdot 10^{-5}$	$3.05 \cdot 10^{-5}$	-	-	-	-	-
# mp	510	270	200	140	130	128	130	1 000
speedup	2.0	3.7	5.0	7.1	7.7	7.8	7.7	1.0
eff.	20%	19%	17%	18%	15%	13%	11%	100%

Theorem 1. Let $c_s^* \in C^1(0, T_{\text{end}})$ be the exact solution of eq. (22) and let $\{\bar{c}_s^{(k)}(T_n)\}_{n=1}^P$ be the k -th iterate of the re-usage algorithm in eq. (28) using the forward Euler method in eq. (23). Under the assumptions made in section 4.1.3, it holds for the error $e_n^{(k)} = |\bar{c}_s^{(k)}(T_n) - c_s^*(T_n)|$ and $k \in \mathbb{N}_0, n \in \{1, \dots, P\}$ that

$$e_n^{(k)} \leq L\delta T \max\left\{1, T_{n-\frac{k}{2}}\right\}^k \frac{c^k \beta^{n-k}}{[k/2]!} \max_{t \in [0, T_{\text{end}}]} |\partial_t c_s^*(t)|. \tag{31}$$

with a constant $c > 0$ (see Lemma 3) and $\beta = 1 + L\delta T$. This implies $e_n^{(k)} \rightarrow 0$ for $k \rightarrow \infty$.

The convergence $e_n^{(k)} \rightarrow 0$ for $k \rightarrow \infty$ follows due to the factorial $[k/2]!$ in the denominator. Compared to Lemma 1, the estimated convergence in k is much slower. However, we will observe in the numerical examples below that a good accuracy might still be reached within few iterations.

Numerical results In Fig. 8, the evolution of the error in the concentration variable $c_s^{(k)}$ is plotted over time for $P = N_c = 10$ and $P = N_c = 60$. We observe convergence to the reference values c_s^* in both cases. Compared to the results in Fig. 6 for the classical parareal algorithm (Algorithm 2), the convergence is significantly slower. Nevertheless, the stopping criterion is already fulfilled after 4-5 iterations (see Table 3), which is only 1-2 iterations more than in the standard parareal algorithm (see Table 1).

While we had observed reduction factors in the order of $\delta T = \lceil \frac{T_{\text{end}}}{P} \rceil$ between two consecutive parareal iterations for the standard parareal algorithm in Table 1, the reduction factors are slightly worse in Table 3. They are, however, in all cases (besides the last value for $P = 20$) bounded above by $\frac{1}{k}$, where k is the parareal iterate. This indicates a convergence of order $\mathcal{O}(\frac{1}{k!})$, which is related to the factor $\mathcal{O}(\frac{1}{[k/2]!})$ in eq. (31) in Theorem 1. Moreover, the absolute values of the error in each parareal iteration are significantly smaller for $P = 60$ compared to the case of $P = 10$ (Note the different scaling in the vertical axis).

In Table 3, we compare the number of micro problems needed for $P = N_c = 10$ to 70. The optimal number of processes is $P = \sqrt{k_{\text{par}} N_f} \approx 63.25$, which is twice as many processes compared to the parareal algorithm in the previous section. The minimal cost in Table 3 is 128 micro problems for $P = 60$, compared to 222 for standard parareal. Compared to a

Table 4

Errors $|\hat{c}_s^{(k),\text{fine}}(T_{\text{end}}) - \hat{c}_s^*(T_{\text{end}})|$ at the midpoint \hat{x}_m of $\hat{\Gamma}$ for $P = 10$ to 50 for the parareal algorithm (Algorithm 2) and a serial reference computation applied to the second numerical example (PDE growth model, $\hat{c}_s^*(T_{\text{end}}) = 0.5186632$). Estimated parallel runtimes are shown, as well as speedup and efficiency compared to the reference computation (right column). Details on the estimation of the runtimes are given in section 5.3. Best numbers are marked in **bold face**.

k	P = 10	P = 20	P = 30	P = 40	P = 50	Ref. (serial)
1	$2.80 \cdot 10^{-3}$	$1.42 \cdot 10^{-3}$	$7.49 \cdot 10^{-4}$	$6.53 \cdot 10^{-4}$	$5.04 \cdot 10^{-4}$	0
2	$8.73 \cdot 10^{-5}$	$9.46 \cdot 10^{-5}$	$1.27 \cdot 10^{-4}$	$6.22 \cdot 10^{-5}$	$4.78 \cdot 10^{-5}$	
3	$3.45 \cdot 10^{-5}$	$2.90 \cdot 10^{-5}$	$2.97 \cdot 10^{-5}$	$1.36 \cdot 10^{-5}$	$9.19 \cdot 10^{-6}$	
4	-	$1.93 \cdot 10^{-6}$	$4.45 \cdot 10^{-6}$	-	-	
Est. par.	11 347 s	9 661 s	8 692 s	6 914 s	7 491 s	26 840 s
speedup	2.4	2.8	3.1	3.9	3.6	1.0
efficiency	24 %	14%	10%	10%	7%	100%

serial time-stepping scheme, we obtain a maximum estimated speed-up of 7.8. Moreover, the efficiency is also significantly improved for larger numbers of processes compared to the standard approach; cf. Tables 1 and 3.

5. Plaque growth problem with a distributed foam cell concentration

We consider the PDE growth model introduced in section 2.2.2, where the concentration of foam cells $\hat{c}_s = \hat{c}_s(\hat{x}, t)$ is now governed by the non-stationary reaction-diffusion problem in eq. (11). For time discretization, we use a linearized variant of the backward Euler scheme. Starting from $c_{s,0} = 0$, the standard backward Euler scheme writes for $l = 1, \dots, N_f$:

Find $c_{s,l+1} \in \mathcal{C} := H_0^1(\hat{\mathcal{S}}, \hat{\Gamma}_s)$ such that

$$\frac{1}{\delta t} (\hat{c}_{s,l+1} - \hat{c}_{s,l}, \hat{\varphi})_{\hat{\mathcal{S}}} - (D_s \hat{\nabla} \hat{c}_{s,l+1}, \hat{\nabla} \varphi)_{\hat{\mathcal{S}}} + R_s (\hat{c}_{s,l+1} (1 - \hat{c}_{s,l+1}), \hat{\varphi})_{\hat{\mathcal{S}}} + (\overline{\gamma}(\sigma_f^{WS}), \hat{\varphi})_{\hat{\Gamma}} = 0 \quad \forall \hat{\varphi} \in \mathcal{C}. \tag{32}$$

Since eq. (32) is a nonlinear system of equations, several iterations of a nonlinear solver (e.g., Newton’s method) would be necessary to solve it. Thus, we consider the following linearized variant, which can be seen as an implicit-explicit (IMEX) scheme (see, e.g., [3])

$$\frac{1}{\delta t} (\hat{c}_{s,l+1} - \hat{c}_{s,l}, \hat{\varphi})_{\hat{\mathcal{S}}} + (D_s \hat{\nabla} \hat{c}_{s,l+1}, \hat{\nabla} \varphi)_{\hat{\mathcal{S}}} + R_s \theta (\hat{c}_{s,l+1} (1 - \hat{c}_{s,l}), \hat{\varphi})_{\hat{\mathcal{S}}} + R_s (1 - \theta) (\hat{c}_{s,l} (1 - \hat{c}_{s,l+1}), \hat{\varphi})_{\hat{\mathcal{S}}} + (\overline{\gamma}(\sigma_f^{WS}), \hat{\varphi})_{\hat{\Gamma}} = 0, \tag{33}$$

where $\theta \in [0, 1]$ is a weighting parameter. In our numerical experiments, we choose $\theta = 0.7$. For spatial discretization, we use again Q_2 finite elements on the mesh described in section 3.2. If we assume that evolution of the wall shear stress σ_f^{WS} was given exactly, the following bound could be shown for the discretization error

$$\|\hat{c}_s(T_{\text{end}}) - \hat{c}_s^{N_f}\|_{\hat{\mathcal{S}}} + \left(\sum_{l=1}^{N_f} \delta T \|\hat{\nabla}(\hat{c}_s(T_l) - \hat{c}_{s,l})\|_{\hat{\mathcal{S}}}^2 \right)^{1/2} \leq c_1 h^2 + c_2 \delta T. \tag{34}$$

Due to the nonlinear interaction between FSI and reaction-diffusion equation, an analysis of the discretization error of the coupled problem is more involved and not within the scope of this paper. The chosen time-step δT is $\frac{1}{1000}$ times the macro time interval length T_{end} , while the horizontal cell size h is a factor $\frac{1}{20}$ of the length of the channel in horizontal direction. Thus, the errors in eq. (34) should be roughly equilibrated.

The material parameters in the fluid and solid problems are chosen as in the first example. The parameters of the PDE growth model are set to $D_s = 1.2 \cdot 10^{-7} \frac{\text{cm}^2}{\text{s}}$, $R_s = 5 \cdot 10^{-7} \frac{1}{\text{s}}$, $\alpha = 5 \cdot 10^{-8} \frac{\text{cm}}{\text{s}}$, and $\sigma_0 = 30 \frac{\text{g}}{\text{cm}^2 \text{s}^2}$. The inflow profile of velocity is chosen as

$$\mathbf{v}_f^{\text{in}}(t, x) = 30 \left((1 + \sin(\pi \frac{t}{\mathcal{P}}))^2 (1 - x_2^2) \right) \text{cm/s},$$

where $\mathcal{P} = 1$ s. The end time is $T_{\text{end}} = 200$ days, and for the fine time steps, we use $\delta t = 0.2$ days, such that again $N_f = 1000$. In Fig. 9, we visualize the results for horizontal velocity v_x and vertical displacement u_y in the deformed fluid and solid domains in three different time instants, respectively. In Fig. 10 the foam cell concentration c_s on the FSI interface is shown for different times t . First, we observe a dominant growth in the center of the domain due to the penetration of monocytes and the reaction term in eq. (11). After $t > 100$ the diffusion gets more significant, such that foam cells are distributed over the whole interface.




Fig. 9. Visualization of the plaque growth at times $t = 50$ days, $t = 100$ days and $t = 200$ days for the second numerical example (PDE growth model). The horizontal velocity (in cm/s) and the vertical displacement (in cm) are shown on the deformed domain at micro time $\tau = 0.5$ s, i.e., the time of maximum inflow velocity. As the plaque growth evolves, significantly higher velocities arise in the central part.




Fig. 10. Foam cell concentration \hat{c}_s at the FSI interface $\hat{\Gamma}$ at different times t in the second numerical example (PDE growth model). In contrast to the ODE model, the concentration is not symmetric around the center anymore for larger times t ; this is due to the reaction-diffusion equation.

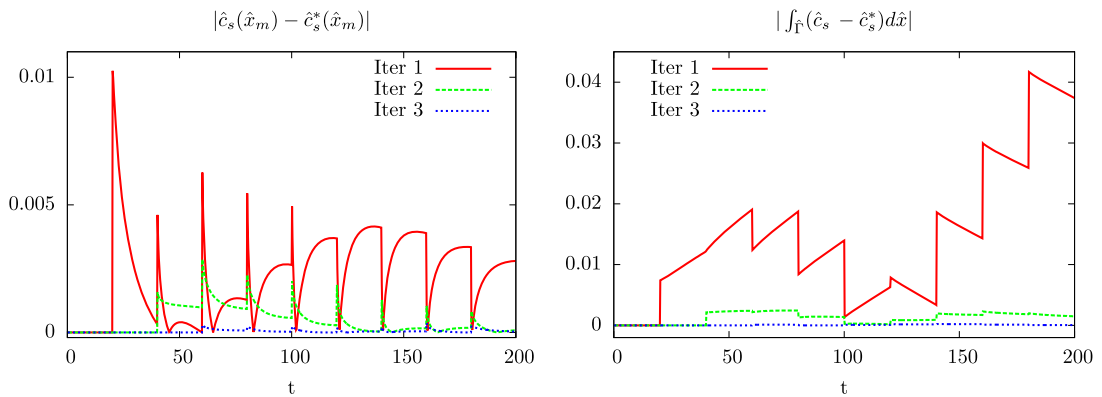


Fig. 11. Errors for the second numerical example (PDE growth model) in different iterations of the parareal algorithm for $P = 10$; the setting is described in section 5. **Left:** Error in the concentration $|\hat{c}_s^{(k),\text{fine}}(\hat{x}_m, t) - \hat{c}_s^*(\hat{x}_m, t)|$ at the midpoint of $\hat{\Gamma}$ over time; see also Fig. 10 for the temporal evolution of \hat{c}_s at the interface. **Right:** Error in the mean value $|\int_{\hat{\Gamma}} \hat{c}_s^{(k),\text{fine}}(\hat{x}, t) - \hat{c}_s^*(\hat{x}, t) d\hat{x}|$ over time.

Table 5

Errors $|\hat{c}_s^{(k),\text{fine}}(T_{\text{end}}) - \hat{c}_s^*(T_{\text{end}})|$ at the midpoint \hat{x}_m of $\hat{\Gamma}$ for a fixed coarse time step ($N_c = 40$), $P \in \{10, 20, 40\}$ and the parareal algorithm (Algorithm 2) applied to the second numerical example (PDE growth model). Estimated parallel runtimes are shown, as well as speedup and efficiency compared to the reference computation (right column). Best numbers are marked in **bold face**.

k	P = 10	P = 20	P = 40	Ref. (serial)
1	$8.51 \cdot 10^{-4}$	$7.60 \cdot 10^{-4}$	$6.53 \cdot 10^{-4}$	0
2	$7.64 \cdot 10^{-6}$	$2.25 \cdot 10^{-5}$	$6.22 \cdot 10^{-5}$	-
3	$4.36 \cdot 10^{-7}$	$3.66 \cdot 10^{-6}$	$1.36 \cdot 10^{-5}$	-
# mp	460	310	235	1 000
speedup	2.2	3.2	4.3	1.0
efficiency	22 %	16%	11%	100%

In the following section 5.1 we investigate numerically the convergence behavior of the standard parareal algorithm (Algorithm 2) and the re-usage variant (Algorithm 3). Then, we discuss the computational costs of the algorithms applied to the PDE growth model in section 5.2. Finally, we compare the algorithms in section 5.3 based on estimated runtimes of a parallel implementation.

5.1. Convergence behavior of the parallel time-stepping algorithms

We consider again the standard parareal algorithm (Algorithm 2) and the modification with re-usage of growth values (Algorithm 3). As stopping criterion, we now choose

$$|\hat{c}_s^{(k+1),\text{fine}}(\hat{x}_m, T_{\text{end}}) - \hat{c}_s^{(k),\text{fine}}(\hat{x}_m, T_{\text{end}})| < \epsilon_{\text{par}} = 10^{-4},$$

where $\hat{x}_m = (0, -1)$ is the center of the FSI interface $\hat{\Gamma}$, that is, a slightly more strict tolerance compared to the ODE model.

Standard parareal In Fig. 11, we show the evolution of the error in the first three iterates of the parareal algorithm over time for $P = N_c = 10$. More precisely, we plot the error in the point functional $\hat{c}_s^{(k),\text{fine}}(\hat{x}_m, t)$, that is, the foam cell concentration at the center of the FSI interface, and the average $\int_{\hat{\Gamma}} \hat{c}_s^{(k),\text{fine}}(\hat{x}, t) dx$ of the foam cell concentration over time. As in the first numerical example, we observe fast convergence towards the reference values. Again, the foam cell concentrations $\hat{c}_s(\hat{x}_m)$ are significantly overestimated in the initialization step, due to the large time-step δT in step (I). These are the starting values for the fine problems in the first parareal iteration.

In Table 4, we compare the convergence of the function values of $\hat{c}_s^{(k),\text{fine}}(\hat{x}_m, T_{\text{end}})$ depending on the number of processes P . The findings are similar to the first numerical example (Table 1). Again, depending on P , three to four parareal iterations were sufficient to reach the stopping criterion, with a slightly faster convergence behavior for larger P .

In Table 5, we show numerical results for a fixed coarse time step $\delta T = \frac{T_{\text{end}}}{40}$, i.e., $N_c = 40$ and varying $P \in \{10, 20, 40\}$. This means that the cost of the coarse propagator (in terms of the number of micro problems to be solved) is equal in all three cases. While the convergence behavior is slightly faster for smaller P , the stopping criterion was satisfied after 3 parareal iterations in all cases. As the fine propagator is cheaper for larger P , the fastest computation is $P = N_c = 40$, with a speed-up of 4.3 in terms of the number of micro problems to be solved.

Parareal with re-usage of growth values In Table 6, we show the results for the parareal algorithm with re-usage of growth values; cf. section 4.2.2. For all tested P , we need $k_{\text{par}} = 5$ iterations to satisfy the stopping criterion. Similar to the first numerical example, these are 1-2 additional iterations compared to the standard parareal algorithm.

5.2. Theoretical discussion of the computational cost

While using the ODE growth model, it was obvious that the solution of the growth model and the communication could be neglected, this requires some discussion for the PDE growth model, as the reaction-diffusion equation (33) needs to be solved to advance the foam cell concentration \hat{c}_s .

Standard parareal In the standard parareal algorithm (Algorithm 2), a time step of the PDE growth model always follows the solution of a micro problem, where the growth values $\gamma(\sigma_f^{WS})$ are computed. Thus, the numbers of time steps of the PDE growth model and micro problems to be solved coincide. Considering that the latter consists of ≥ 100 time steps, each involving the solution of a nonlinear, coupled FSI problem, while the growth model only requires the solution of a single scalar PDE of reaction-diffusion type, it is clear that the computational cost of the growth model is negligible.

Concerning communication, each process $p = 1, \dots, P$ needs to communicate its final foam cell concentration $\mathcal{F}(\bar{c}_s^{(k)}(T_p)) = \hat{c}_s^{(k),\text{fine}}(t_{p,n_p})$, which is a scalar-valued finite element function defined in the solid domain \hat{S} , to the master process if the master-slave approach is used. No communication of $\mathcal{F}(\bar{c}_s^{(k)}(T_p))$ is required in a distributed approach; cf. section 4.1.1. Then, after the coarse propagator, the variables $\bar{c}_s^{(k+1)}(t_{p,n_p})$ have to be communicated: in the master-slave approach, the

data is transferred back to the slaves, and in the distributed case, it is communicated to the next process in line. Using the discretization outlined above, this corresponds to 369 degrees of freedom that need to be communicated each time.

As already discussed in section 4.1.1, in the master-slave case, this communication step is unfavorable because it involves all-to-one and one-to-all communication, whereas the communication pattern for the distributed parallelization only involves neighbor communication, which is more beneficial. However, in both cases, the two communication steps are only performed once during each parareal iteration. In order to investigate if the communication cost is still negligible for the PDE growth model, parallel numerical experiments based on an actual parallel implementation are needed, especially for realistic three-dimensional problems; see also the discussion in [34]. Due to the large computational cost of the micro problems in realistic three-dimensional problems, the communication cost might still be comparably small, which motivates the focus on the number of micro problems to be solved in our discussion; recall that, in our case, the micro problems consist of ≥ 100 time steps of a fully coupled FSI problem. Obviously, this argumentation needs to be confirmed based on an actual parallel implementation; we will investigate this further in our future work.

Parareal with re-usage of growth values The objective of the re-usage algorithm is to reduce the number of micro problems to be solved in the coarse-scale propagator. As this allows for a smaller time-step size δt in the coarse problem, the number of reaction-diffusion equations to be solved, might increase significantly. Instead of N_c of such equations in step (II.b)(i) of Algorithm 2, Algorithm 3 requires N_f reaction-diffusion steps in step (II.b)(ii), where $N_f = 1000$ in the example used in this section. Thus, in k_{par} iterations of the parareal algorithm, the number of such equations to be solved is $k_{\text{par}} \cdot N_f$ in step (II.b)(ii) plus N_c equations in step (I) and $k_{\text{par}} \cdot \lceil \frac{N_f}{P} \rceil$ in step (II.a)(ii) of Algorithm 3. In total, Algorithm 3 requires the solution of

$$k_{\text{par}} \cdot (N_f + \lceil N_f/P \rceil) + N_c$$

reaction-diffusion equations. For a comparison, we note that the number of micro problems to be solved for P processes was $(k_{\text{par}} + 1)N_c + k_{\text{par}} \cdot \lceil N_f/P \rceil$ (see section 4.1.2). This means that the number of reaction-diffusion equations to be solved is by a factor

$$\frac{k_{\text{par}} \cdot (N_f + \lceil N_f/P \rceil) + N_c}{(k_{\text{par}} + 1)N_c + k_{\text{par}} \cdot \lceil N_f/P \rceil}$$

larger compared to the number of micro-problems. A simple calculation yields that

$$\frac{k_{\text{par}} \cdot (N_f + \lceil N_f/P \rceil) + N_c}{(k_{\text{par}} + 1)N_c + k_{\text{par}} \cdot \lceil N_f/P \rceil} \leq \frac{k_{\text{par}} (N_f + \lceil N_f/P \rceil + N_c)}{k_{\text{par}} (N_c + \lceil N_f/P \rceil)} = \frac{N_f}{N_c + \lceil N_f/P \rceil} + 1 \leq \frac{\sqrt{N_f}}{2} + 1.$$

In the last inequality, we have used that $N_c \geq P$. For $N_f = 1000$, the bound on the right-hand side is approximately 16.8.

Of course, another option would be to use coarser time steps for the coarse propagator. However, noting again that a micro problem consists of ≥ 100 FSI steps, the computational cost for ≤ 16.8 scalar reaction-diffusion equations is still much cheaper. This will be confirmed in the following section, where we show estimated runtimes of a parallel implementation and the respective contributions from the micro problems and the solves of reaction-diffusion equations.

For the re-usage of growth values and the distributed parallelization scheme, the communication cost does not change compared to the standard parareal algorithm. This is because the growth values to be re-used are already available on the process; this follows directly from the discussion in section 4.2.2. In the case, of master-slave communication, the communication cost is increased. In particular, N_f growth functions $\bar{\gamma}_i$, $i = 1, \dots, N_f$ need to be transferred from the processes $p = 1, \dots, P$ to the master process. Each $\bar{\gamma}_i$ is a spatially discretized function which is defined on the FSI interface $\hat{\Gamma}$. In the example considered here, it is non-zero only in the area $\hat{x}_1 \in [-1, 1]$ which corresponds to 7 degrees of freedom in our discretization; again, for a realistic three-dimensional problem, the number of degrees of freedom will increase drastically but will still remain small compared to the full problem size. As mentioned before, we assume that the communication cost is negligible compared to the solution of the micro problems in our discussion. Again, this assumption has to be tested in the future based on parallel results.

5.3. Comparison of runtimes

The computational results given in this paper serve as a proof of concept to test the presented algorithms. From the discussion in the previous subsection, we assume that the only relevant computational cost comes from the solution of the micro problems. As mentioned before in our current implementation, we do not solve the fine-scale problems in parallel on different processes. Instead, they are solved sequentially one after the other on a single process. Since the cost for the micro problems dominates the computation times for any serial or parallel simulation of the plaque growth, we can still discuss the parallelization potential of the methods. The parallel implementation of the algorithms itself is subject to future work.

In this section, we will give an estimate of the runtimes that would be required in a parallel implementation. For this purpose, we list the serial part and parallel contributions of the computing times in Table 7. The serial part corresponds to

Table 6

Errors $|c_s^{(k),\text{fine}}(\hat{x}_m, T_{\text{end}}) - \hat{c}_s^*(\hat{x}_m, T_{\text{end}})|$ at the midpoint \hat{x}_m of $\hat{\Gamma}$ for $P = 10$ to 70 for Algorithm 3 (Re-usage of growth values) applied to the second numerical example (PDE growth model). We show estimated parallel runtimes, as well as speedup and efficiency compared to the reference computation (right column). Details on the estimation of runtimes are given in section 5.3. Best numbers are marked in **bold face**.

k	$P = 10$	$P = 20$	$P = 30$	$P = 40$	$P = 50$	$P = 60$	$P = 70$	Ref.
1	$2.80 \cdot 10^{-3}$	$1.42 \cdot 10^{-3}$	$7.49 \cdot 10^{-4}$	$6.53 \cdot 10^{-4}$	$5.04 \cdot 10^{-4}$	$4.00 \cdot 10^{-4}$	$3.32 \cdot 10^{-4}$	0
2	$6.37 \cdot 10^{-4}$	$6.96 \cdot 10^{-4}$	$5.62 \cdot 10^{-4}$	$4.08 \cdot 10^{-4}$	$3.45 \cdot 10^{-4}$	$3.09 \cdot 10^{-4}$	$2.82 \cdot 10^{-4}$	
3	$1.33 \cdot 10^{-4}$	$1.87 \cdot 10^{-4}$	$2.06 \cdot 10^{-4}$	$1.56 \cdot 10^{-4}$	$1.39 \cdot 10^{-4}$	$1.30 \cdot 10^{-4}$	$1.25 \cdot 10^{-4}$	
4	$2.87 \cdot 10^{-5}$	$5.31 \cdot 10^{-5}$	$6.14 \cdot 10^{-5}$	$5.08 \cdot 10^{-5}$	$4.61 \cdot 10^{-5}$	$4.29 \cdot 10^{-5}$	$4.02 \cdot 10^{-5}$	
5	$5.05 \cdot 10^{-6}$	$1.47 \cdot 10^{-5}$	$1.67 \cdot 10^{-5}$	$1.45 \cdot 10^{-5}$	$1.20 \cdot 10^{-5}$	$1.07 \cdot 10^{-5}$	$1.05 \cdot 10^{-5}$	
Est. par.	17 733 s	9 685 s	7 105 s	5 902 s	5 277 s	4 925 s	4 804 s	26 840 s
speedup	1.5	2.8	3.8	4.5	5.1	5.4	5.6	1.0
efficiency	15 %	14%	13%	11%	10%	9%	8%	100%




Fig. 12. Illustration of the computational times spent within the coarse- (in serial) and fine-scale (in parallel) problems in Algorithm 2 and Algorithm 3 for the second numerical example (PDE growth problem)

Table 7

Estimated parallel runtimes (in s) of the parareal algorithm and the variant “Re-usage of growth values” for $P = 10, \dots, 70$ and a PDE growth model. We show the time in seconds spent on the master process and the maximum and average time spent on the slave processes. The estimated parallel runtime is the sum of the time spent in serial (master) and the maximum time needed among the slaves; best numbers marked in **bold face**. A visualization of the runtimes is given in Fig. 13.

Standard parareal (Algorithm 2)					Parareal with re-usage of growth values (Algorithm 3)				
P	Coarse	Fine: max.	(Aver.)	Est. par.	P	Coarse	Fine: max.	(Aver.)	Est. par.
10	1 096 s	10 251 s	(8 009 s)	11 347 s	10	658 s	17 075 s	(13 342 s)	17 733 s
20	2 691 s	6 968 s	(5 318 s)	9 661 s	20	930 s	8 755 s	(6 669 s)	9 685 s
30	3 943 s	4 749 s	(3 581 s)	8 692 s	30	1 171 s	5 934 s	(4 455 s)	7 105 s
40	4 273 s	2 641 s	(1 988 s)	6 914 s	40	1 448 s	4 454 s	(3 339 s)	5 902 s
50	5 361 s	2 130 s	(1 601 s)	7 491 s	50	1 722 s	3 555 s	(2 670 s)	5 277 s
60	6 197 s	1 796 s	(1 331 s)	7 993 s	60	1 933 s	2 992 s	(2 219 s)	4 925 s
70	7 339 s	1 564 s	(1 161 s)	8 903 s	70	2 151 s	2 653 s	(1 899 s)	4 804 s

the coarse propagation, which is performed on a single master process (steps (I) and (II.b) in Algorithm 2 or Algorithm 3). The parallel part corresponds to the solution of the micro problems (step (II.a)). It can be performed concurrently on all processes $p = 1, \dots, P$. The computing times vary slightly across the processes. It increases for larger processes p , as more Newton iterations are required for the FSI problem, when the channel is already significantly narrowed due to the advancing plaque growth. The estimated parallel runtimes given in Tables 4, 6 and 7 are the sum of the serial part, which corresponds to the coarse propagator, and the maximum time needed by one of the processes $p = 1, \dots, P$ in the parallel part.

In Fig. 12, we illustrate the computational times needed to solve the fine problems (in parallel) and coarse-scale propagators (in serial) for different processes and with Algorithm 2 and Algorithm 3. We see that the total computational times decrease for increasing P until $P = 40$ for standard parareal and until $P = 60$ for the re-usage variant. In the standard parareal algorithm (Algorithm 2) the cost of the coarse propagators becomes significant already from $P = 20$, while it is much smaller for the re-usage variant. While for $P = 10$, Algorithm 2 is still faster, this changes for $P > 20$.

The corresponding speed-ups and efficiencies compared to a serial time-stepping are given in the last rows of Tables 4 and 6 for standard parareal algorithm (Algorithm 2) and the re-usage variant (Algorithm 3), respectively. The lowest computing time for standard parareal is 6 914 s for $P = 40$, which corresponds to a speed-up of 3.9. For the re-usage variant, we obtain a maximum speed-up of 5.6 for $P = 40$, which corresponds to an estimated parallel runtime of 4 804 s. The speed-ups are slightly lower than the speed-ups in section 4 (cf. Tables 1 and 3), that were computed based on the number of micro problems to be solved. This is mostly due to the load imbalancing among the slave processes, as - depending on




Fig. 13. Illustration of the computational times spent within micro-scale problems and PDE growth problems on the coarse (in serial) and fine scale (in parallel) in Algorithm 2 and Algorithm 3 for the second numerical example (PDE growth problem). Note that the time needed for the PDE growth model within the fine-scale problems is so small that it is not visible in both plots. The corresponding times can be found in Table 7.

the state of the plaque growth - some micro problems are more costly than others; for instance, due to a higher number of Newton iterations. The differences can be inferred best by comparing the average and the maximum time spent on the slave processes in Table 7. The efficiencies decrease again monotonically for increasing P in both Tables 4 and 6. For the re-usage variant, the efficiencies are again much more stable, due to the reduced cost of the coarse propagator.

Finally, we compare in Fig. 13 the computing times needed for the solution of the FSI micro problems and those needed for the reaction-diffusion equations, for both algorithms and $P = 10, \dots, 60$. We observe that the times needed for the latter are negligible in all cases, which confirms the discussion in section 5.2.

6. Conclusion

We have derived a parareal algorithm for the time parallelization of the macro scale in a two-scale formulation for the simulation of atherosclerotic plaque growth. To reduce the computational cost of the coarse-scale propagators, we have introduced a variant which re-uses growth values that were computed within the fine problems and avoids additional costly micro-scale computations in serial.

The approaches have been tested on two different numerical examples of increasing complexity: first, by means of a simple ODE growth model, and secondly, by a PDE model of reaction-diffusion type. In this proof-of-concept, we analyze the approaches based on results and timings of a serial implementation. Since the number of communication steps is low compared to the computational work, we are still able to provide meaningful results. In the first case, we obtain estimated speedups up to 6.3 (standard parareal) respectively 7.8 (re-usage variant) in terms of the number of micro problems to be solved. In the PDE model, the maximum estimated speedups, now based on estimated parallel runtimes, were 3.9 respectively 5.6. In future work, these results will have to be confirmed using a parallel implementation. Therefore, we have discussed theoretically two parallelization schemes, master-slave and distributed parallelization. The discussion indicated that the distributed scheme might be beneficial, which also needs to be tested in future work.

Additional research should also be invested into further improving the efficiency of the coarse propagator. Therefore, the idea of Algorithm 3 could be extended, for example, by storing and interpolating the computed growth values, instead of simply re-using them. For the ODE case, such an interpolation approach has already been applied in our proceedings paper [53] with promising results. The extension to the PDE model is, however, not straight-forward, as an operator $\mathcal{M} : c_s \rightarrow \overline{\mathcal{P}(\sigma_f^{WS}(c_s))}$ between spatially distributed functions needs to be approximated. Further future investigations include the application of the algorithms in complex three-dimensional geometries, more complex plaque growth, and arterial wall models. Finally, the approaches presented here can be combined with a spatial parallelization of the FSI problems or with adaptive time-stepping on micro and macro scale, as presented by Richter & Lautsch [41].

CRedit authorship contribution statement

All results have been obtained by equal contributions of the authors.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Appendix A

In this section, we show convergence of the re-usage parareal algorithm, Algorithm 3 resp. eq. (28), applied to the ODE model in eq. (22) and its approximation by the explicit Euler method in eq. (23). First, we obtain from eqs. (23) and (24) directly for $c_s^1, c_s^2 \in \mathbb{R}$

$$|\mathcal{C}(I_n, c_s^1(T_{n-1}), \gamma) - \mathcal{C}(I_n, c_s^2(T_{n-1}), \gamma)| = |c_s^1(T_{n-1}) - c_s^2(T_{n-1})|, \tag{35}$$

$$|\mathcal{C}(I_n, c_s, \gamma(c_s^1)) - \mathcal{C}(I_n, c_s, \gamma(c_s^2))| = \delta T |\gamma(c_s^1) - \gamma(c_s^2)| \leq L\delta T |c_s^1 - c_s^2|. \tag{36}$$

By $c_{s,n-1}(T)$ we denote in the following the function that solves the (continuous) ODE eq. (22) with initial value $c_{s,n-1}(T_{n-1}) = c_{s,n-1}$. We have using eqs. (23) to (25)

$$\begin{aligned} & \left| (\mathcal{F}(I_n, c_s(T_{n-1}), \gamma(c_s(T_{n-1}))) - \mathcal{C}(I_n, c_s(T_{n-1}), \gamma(c_s(T_{n-1})))) \right. \\ & \quad \left. - (\mathcal{F}(I_n, c_{s,n-1}, \gamma(c_{s,n-1})) - \mathcal{C}(I_n, c_{s,n-1}, \gamma(c_{s,n-1}))) \right| \\ &= \left| \int_{T_{n-1}}^{T_n} \gamma(c_s(T)) - \gamma(c_{s,n-1}(T)) \, dT - \delta T (\gamma(c_s(T_{n-1})) - \gamma(c_{s,n-1})) \right| \\ &\leq c\delta T^2 \left| \frac{d}{dt} \gamma(c_s(T_{n-1})) - \frac{d}{dt} \gamma(c_{s,n-1}) \right| \leq \alpha_0 \delta T^2 |c_s(T_{n-1}) - c_{s,n-1}| \end{aligned} \tag{37}$$

for some constant $\alpha_0 = cL > 0$; see also [32].

The following recursion builds the basis for the error estimate:

Lemma 2. Let $c_s^* \in C^1(0, T_{end})$ be the exact solution of eq. (22). $\{\bar{c}_s^{(k)}(T_n)\}_{n=1}^P$ be the k -th iterate of the re-usage algorithm in eq. (28) using the forward Euler method in eq. (23) and let $e_n^{(k)} = |c_s^*(T_n) - \bar{c}_{s,n}^{(k)}|$ be the error in the k -th iteration of the re-usage algorithm. Under the assumptions made in section 4.1.3, it holds for $k = 1$ that

$$e_n^{(1)} \leq \alpha_0 \delta T^2 e_{n-1}^{(0)} + \alpha_1 \delta T e_{n-2}^{(0)} + e_{n-1}^{(1)}, \tag{38}$$

and for $k \geq 2$,

$$e_n^{(k)} \leq \alpha_1 \delta T \left(e_{n-1}^{(k-1)} + e_{n-2}^{(k-1)} + e_{n-2}^{(k-2)} \right) + e_{n-1}^{(k)}, \tag{39}$$

with a constant $\alpha_1 \geq \max\{L + \alpha_0 \delta T, L(1 + L\delta T)\}$.

Proof. Consider first the case $k \geq 2$. Using eqs. (25) and (28), we have

$$\begin{aligned} e_n^{(k)} &= |\mathcal{F}(I_n, c_s^*(T_{n-1}), \gamma(c_s^*(T_{n-1}))) - \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(k)}, \gamma(c_{s,n-1}^{(k),\text{fine}})) \\ &\quad - \mathcal{F}(I_n, \bar{c}_{s,n-1}^{(k-1)}, \gamma(\bar{c}_{s,n-1}^{(k-1)})) + \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(k-1)}, \gamma(c_{s,n-1}^{(k-1),\text{fine}}))| \\ &\leq |\mathcal{F}(I_n, c_s^*(T_{n-1}), \gamma(c_s^*(T_{n-1}))) - \mathcal{C}(I_n, c_s^*(T_{n-1}), \gamma(c_s^*(T_{n-1}))) \\ &\quad - \mathcal{F}(I_n, \bar{c}_{s,n-1}^{(k-1)}, \gamma(\bar{c}_{s,n-1}^{(k-1)})) + \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(k-1)}, \gamma(\bar{c}_{s,n-1}^{(k-1)}))| \\ &\quad + |\mathcal{C}(I_n, \bar{c}_{s,n-1}^{(k-1)}, \gamma(\bar{c}_{s,n-1}^{(k-1)})) - \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(k-1)}, \gamma(c_{s,n-1}^{(k-1),\text{fine}}))| \\ &\quad + |\mathcal{C}(I_n, \bar{c}_{s,n-1}^{(k)}, \gamma(c_{s,n-1}^{(k),\text{fine}})) - \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(k)}, \gamma(c_s^*(T_{n-1})))| \\ &\quad + |\mathcal{C}(I_n, \bar{c}_{s,n-1}^{(k)}, \gamma(c_s^*(T_{n-1}))) - \mathcal{C}(I_n, c_s^*(T_{n-1}), \gamma(c_s^*(T_{n-1})))| \end{aligned}$$

From eqs. (35) to (37), we have

$$\begin{aligned} e_n^{(k)} &\leq \alpha_0 \delta T^2 \left| c_s^*(T_{n-1}) - \bar{c}_{s,n-1}^{(k-1)} \right| + L\delta T \left| \bar{c}_{s,n-1}^{(k-1)} - c_{s,n-1}^{(k-1),\text{fine}} \right| \\ &\quad + L\delta T \left| c_{s,n-1}^{(k),\text{fine}} - c_s^*(T_{n-1}) \right| + \left| \bar{c}_{s,n-1}^{(k)} - c_s^*(T_{n-1}) \right| \\ &= \alpha_0 \delta T^2 e_{n-1}^{(k-1)} + L\delta T (e_{n-1}^{(k-1),\text{fine}} + e_{n-1}^{(k-1)}) + L\delta T e_{n-1}^{(k),\text{fine}} + e_{n-1}^{(k)}, \end{aligned} \tag{40}$$

where $e_{n-1}^{(k),\text{fine}} := |c_{s,n-1}^{(k),\text{fine}} - c_s^*(T_n)|$. This error contribution is estimated further by using eqs. (25), (29) and (35)

$$\begin{aligned}
 e_{n-1}^{(k),\text{fine}} &:= |c_{s,n-1}^{(k),\text{fine}} - c_s^*(T_n)| = |\mathcal{F}(I_{n-1}, c_{s,n-2}^{(k-1)}, \gamma(c_{s,n-2}^{(k-1)})) - \mathcal{F}(I_{n-1}, c_s^*(T_{n-2}), \gamma(c_s^*(T_{n-2})))| \\
 &= |c_{s,n-2}^{(k-1)} - c_s^*(T_{n-2}) + \int_{T_{n-2}}^{T_{n-1}} \gamma(c_{s,n-2}^{(k-1)}(t)) - \gamma(c_s^*(t)) dt| \\
 &\leq (1 + L\delta T)e_{n-2}^{(k-1)}.
 \end{aligned}
 \tag{41}$$

Inserting this into eq. (40) yields eq. (39). For $k = 1$, we have $c_{s,n-1}^{(0),\text{fine}} = c_{s,n-1}^{(0)}$. An analogous argumentation yields

$$\begin{aligned}
 e_n^{(1)} &= |\mathcal{F}(I_n, c_s^*(T_{n-1}), \gamma(c_s^*(T_{n-1}))) - \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(1)}, \gamma(c_{s,n-1}^{(1),\text{fine}})) \\
 &\quad - \mathcal{F}(I_n, \bar{c}_{s,n-1}^{(0)}, \gamma(\bar{c}_{s,n-1}^{(0)})) + \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(0)}, \gamma(c_{s,n-1}^{(0)}))| \\
 &\leq |\mathcal{F}(I_n, c_s^*(T_{n-1}), \gamma(c_s^*(T_{n-1}))) - \mathcal{C}(I_n, c_s^*(T_{n-1}), \gamma(c_s^*(T_{n-1}))) \\
 &\quad - \mathcal{F}(I_n, \bar{c}_{s,n-1}^{(0)}, \gamma(\bar{c}_{s,n-1}^{(0)})) + \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(0)}, \gamma(\bar{c}_{s,n-1}^{(0)}))| \\
 &\quad + |\mathcal{C}(I_n, \bar{c}_{s,n-1}^{(1)}, \gamma(c_{s,n-1}^{(1),\text{fine}})) - \mathcal{C}(I_n, \bar{c}_{s,n-1}^{(1)}, \gamma(c_s^*(T_{n-1})))| \\
 &\quad + |\mathcal{C}(I_n, \bar{c}_{s,n-1}^{(1)}, \gamma(c_s^*(T_{n-1}))) - \mathcal{C}(I_n, c_s^*(T_{n-1}), \gamma(c_s^*(T_{n-1})))| \\
 &\leq \alpha_0 \delta T^2 e_{n-1}^{(0)} + L\delta T e_{n-1}^{(1),\text{fine}} + e_{n-1}^{(1)}.
 \end{aligned}$$

Then, eq. (38) follows by using eq. (41). \square

Lemma 3. Let $e_n^{(k)} = |\bar{c}_s^{(k)}(T_n) - c_s^*(T_n)|$ be the error in the k -th iteration of the re-usage algorithm. Under the assumptions made above, it holds for $k \in \mathbb{N}_0$ and $n \in \mathbb{N}$ that

$$e_n^{(k)} \leq L\tilde{\alpha}^k \beta^{n-k} \cdot \left(\sum_{l=\lceil \frac{k}{2} \rceil + 1}^{k+1} \frac{3^{l-1}}{l!} \delta T^{l+1} \left(n - \frac{k}{2} \right)^l \right) \max_{t \in [0, T_{\text{end}}]} |\partial_t c_s^*(t)|
 \tag{42}$$

$$\leq L\tilde{\alpha}^k \beta^{n-k} \delta T \max \left\{ 1, T_{n-\frac{k}{2}} \right\}^k \frac{3^{\lceil \frac{k}{2} \rceil}}{[\lceil k/2 \rceil]!} \max_{t \in [0, T_{\text{end}}]} |\partial_t c_s^*(t)|,
 \tag{43}$$

where $\tilde{\alpha} = \max\{\alpha_0 \delta T, \alpha_1, 1\}$ and $\beta = 1 + L\delta T$.

Proof. We prove the lemma by induction over $k \in \mathbb{N}$. For $k = 0$, a standard estimate of the forward Euler methods gives, using eq. (24),

$$\begin{aligned}
 e_n^{(0)} &:= |c_{s,n}^0 - c_s^*(T_n)| \leq |c_{s,n-1}^0 - c_s^*(T_{n-1})| + \int_{T_{n-1}}^{T_n} |\gamma(c_{s,n-1}^0) - \gamma(c_s^*(t))| dt \\
 &\leq e_{n-1}^{(0)} + L \int_{T_{n-1}}^{T_n} |c_{s,n-1}^0 - \underbrace{c_s^*(t)}_{=c_s^*(T_{n-1}) + \delta T \partial_t c_s^*(\xi)}| dt \\
 &\leq \underbrace{(1 + L\delta T)}_{=\beta} e_{n-1}^{(0)} + L\delta T^2 \max_{t \in [0, T_{\text{end}}]} |\partial_t c_s^*(t)|.
 \end{aligned}
 \tag{44}$$

To abbreviate the notation, we set $c_0 := \max_{t \in [0, T_{\text{end}}]} |\partial_t c_s^*(t)|$. We apply eq. (44) recursively

$$e_n^{(0)} \leq \beta e_{n-1}^{(0)} + c_0 L\delta T^2 \leq c_0 L\delta T^2 \sum_{l=0}^{n-1} \beta^l = c_0 L\delta T^2 \underbrace{\frac{\beta^n - 1}{\beta - 1}}_{=L\delta T} = c_0 \delta T (\beta^n - 1) \leq c_0 L n \delta T^2 \beta^n.
 \tag{45}$$

In the last inequality, we have used that $\beta^n - 1 \leq nL\delta T\beta^n$, which can be shown by induction over n . The estimate in eq. (45) is exactly eq. (42) for $k = 0$.

For $k = 1$, we have using eqs. (38) and (45), the fact that $e_k^{(k)} = 0$, and by the definition of the parareal algorithm

$$\begin{aligned}
 e_n^{(1)} &\leq \alpha_0 \delta T^2 e_{n-1}^{(0)} + \alpha_1 \delta T e_{n-2}^{(0)} + e_{n-1}^{(1)} \\
 &\leq \alpha_0 c_0 L \delta T^4 (n-1) \beta^{n-1} + \alpha_1 c_0 L \delta T^3 (n-2) \beta^{n-2} + e_{n-1}^{(1)}
 \end{aligned}$$

$$\begin{aligned} &\leq 2\tilde{\alpha}c_0L\delta T^3(n-1)\beta^{n-1} + e_{n-1}^{(1)} \\ &\leq 2\tilde{\alpha}c_0L\beta^{n-1}\delta T^3 \sum_{l=1}^{n-1} l + \underbrace{e_1^{(1)}}_{=0} \\ &= 2\tilde{\alpha}c_0L\beta^{n-1}\delta T^3 \frac{n(n-1)}{2} \leq \tilde{\alpha}c_0L\beta^{n-1}\delta T^3 (n-1/2)^2. \end{aligned}$$

This is by a factor $\frac{3}{2}$ smaller compared to eq. (42).

Now, let $k \geq 2$. We assume that the estimate in eq. (42) is true for $k - 1$ and $k - 2$. By eq. (39) and the assumption of the induction, we have

$$\begin{aligned} e_n^{(k)} &\leq \alpha_1 \delta T \left(e_{n-1}^{(k-1)} + e_{n-2}^{(k-1)} + e_{n-2}^{(k-2)} \right) + e_{n-1}^{(k)} \\ &\leq \alpha_1 \delta T \left(c_0 L \tilde{\alpha}^{k-1} \beta^{n-k} \cdot \left(\sum_{l=\lceil \frac{k+1}{2} \rceil}^k \frac{3^{l-1}}{l!} \delta T^{l+1} \left(n - \frac{k+1}{2} \right)^l \right) \right. \\ &\quad \left. + c_0 L \tilde{\alpha}^{k-1} \beta^{n-k-1} \cdot \left(\sum_{l=\lceil \frac{k+1}{2} \rceil}^k \frac{3^{l-1}}{l!} \delta T^{l+1} \left(n - 1 - \frac{k+1}{2} \right)^l \right) \right. \\ &\quad \left. + c_0 L \tilde{\alpha}^{k-2} \beta^{n-k} \cdot \left(\sum_{l=\lceil \frac{k}{2} \rceil}^{k-1} \frac{3^{l-1}}{l!} \delta T^{l+1} \left(n - 1 - \frac{k}{2} \right)^l \right) \right) + e_{n-1}^{(k)} \\ &\leq 3c_0 L \tilde{\alpha}^k \beta^{n-k} \cdot \left(\sum_{l=\lceil \frac{k}{2} \rceil}^k \frac{3^{l-1}}{l!} \delta T^{l+2} \left(n - \frac{k+1}{2} \right)^l \right) + e_{n-1}^{(k)}. \end{aligned} \tag{46}$$

We apply this estimate recursively for the last term $e_{n-1}^{(k)}$ in eq. (46) to get

$$\begin{aligned} e_n^{(k)} &\leq c_0 L \tilde{\alpha}^k \beta^{n-k} \cdot \sum_{m=k+1}^n \left(\sum_{l=\lceil \frac{k}{2} \rceil}^k \frac{3^l}{l!} \delta T^{l+2} \left(m - \frac{k+1}{2} \right)^l \right) + \underbrace{e_k^{(k)}}_{=0} \\ &= c_0 L \tilde{\alpha}^k \beta^{n-k} \cdot \left(\sum_{l=\lceil \frac{k}{2} \rceil}^k \frac{3^l}{l!} \delta T^{l+2} \cdot \underbrace{\left(\sum_{m=\frac{k+1}{2}}^{n-\frac{k+1}{2}} m^l \right)}_{=:s} \right). \end{aligned}$$

In the case that $\frac{k+1}{2}$ and $n - \frac{k+1}{2}$ are no natural number, the sum s is to be understood in such a way that the index m advances iteratively by 1 until reaching the upper limit.

The sum s is an approximation of the integral

$$\int_{k/2}^{n-k/2} x^l dx \leq \frac{1}{l+1} \left(n - \frac{k}{2} \right)^{l+1}$$

with the midpoint rule. As $f(x) = x^l$ is a convex function ($f''(x) > 0$) for a positive x , the integral is an upper bound for s ; see the error representation of the midpoint rule, e.g., in Theorem 8.41 in [49].

We have thus shown that

$$e_n^{(k)} \leq c_0 L \tilde{\alpha}^k \beta^{n-k} \cdot \left(\sum_{l=\lceil \frac{k}{2} \rceil}^k \frac{3^l}{(l+1)!} \delta T^{l+2} \left(n - \frac{k}{2} \right)^{l+1} \right).$$

Shifting the index l by 1 gives eq. (42). The estimate in eq. (43) follows from eq. (42) by noting that

$$(n - \frac{k}{2})\delta T = T_{n-\frac{k}{2}} \leq \max \left\{ 1, T_{n-\frac{k}{2}} \right\}$$

and the fact that the term $\frac{3^{l-1}}{l!}$ is decreasing for $l \geq 2$:

$$\begin{aligned} e_n^{(k)} &\leq c_0 L \tilde{\alpha}^k \beta^{n-k} \cdot \left(\sum_{l=\lceil \frac{k}{2} + 1 \rceil}^{k+1} \frac{3^{l-1}}{l!} \delta T^{l+1} \left(n - \frac{k}{2} \right)^l \right) \\ &\leq c_0 L \tilde{\alpha}^k \beta^{n-k} \delta T \left(\max \left\{ 1, T_{n-\frac{k}{2}} \right\} \right)^{k+1} \cdot \left(\sum_{l=\lceil \frac{k}{2} + 1 \rceil}^{k+1} \frac{3^{l-1}}{l!} \right) \\ &\leq c_0 L \tilde{\alpha}^k \beta^{n-k} \delta T \left(\max \left\{ 1, T_{n-\frac{k}{2}} \right\} \right)^{k+1} \lfloor \frac{k}{2} + 1 \rfloor \frac{3^{\lceil \frac{k}{2} \rceil}}{\lceil \frac{k}{2} + 1 \rceil!} \\ &\leq c_0 L \tilde{\alpha}^k \beta^{n-k} \delta T \left(\max \left\{ 1, T_{n-\frac{k}{2}} \right\} \right)^{k+1} \frac{3^{\lceil \frac{k}{2} \rceil}}{\lceil \frac{k}{2} \rceil!}. \quad \square \end{aligned}$$

References

- [1] A. Abdulle, W. E. B. Engquist, E. Vanden-Eijnden, The heterogeneous multiscale method, *Acta Numer.* (2012) 1–87.
- [2] C. Ager, B. Schott, A.T. Vuong, A. Popp, W.A. Wall, A consistent approach for fluid-structure-contact interaction based on a porous flow model for rough surface contact, *Int. J. Numer. Methods Eng.* 119 (13) (2019) 1345–1378.
- [3] U.M. Ascher, S.J. Ruuth, B.T.R. Wetton, Implicit-explicit methods for time-dependent partial differential equations, *SIAM J. Numer. Anal.* 32 (3) (1995) 797–823.
- [4] Eric Aubanel, Scheduling of tasks in the parareal algorithm, *Parallel Comput.* 37 (3) (March 2011) 172–182.
- [5] D. Balzani, S. Deparis, S. Fausten, D. Forti, A. Heinlein, A. Klawonn, A. Quarteroni, O. Rheinbach, J. Schröder, Numerical modeling of fluid-structure interaction in arteries with anisotropic polyconvex hyperelastic and anisotropic viscoelastic material models at finite strains, *Int. J. Numer. Methods Biomed. Eng.* 32 (10) (2016) e02756.
- [6] D. Balzani, S. Deparis, S. Fausten, D. Forti, A. Heinlein, A. Klawonn, A. Quarteroni, O. Rheinbach, Jörg Schröder, Aspects of arterial wall simulations: nonlinear anisotropic material models and fluid structure interaction, in: *Joint 11th WCCM, 5th ECCM and 6th ECFD 2014*, 2014, pp. 947–958.
- [7] D. Balzani, A. Heinlein, A. Klawonn, O. Rheinbach, J. Schröder, Comparison of arterial wall models in fluid-structure interaction simulations, *Comput. Mech.* (2023), <https://doi.org/10.1007/s00466-023-02321-y>.
- [8] D. Balzani, P. Neff, J. Schröder, G.A. Holzapfel, A polyconvex framework for soft biological tissues. Adjustment to experimental data, *Int. J. Solids Struct.* 43 (20) (2006) 6052–6070.
- [9] A.T. Barker, X.-C. Cai, Scalable parallel methods for monolithic coupling in fluid-structure interaction with application to blood flow modeling, *J. Comp. Physiol.* 229 (3) (2010) 642–659.
- [10] Y. Bazilevs, K. Takizawa, T.E. Tezduyar, *Computational Fluid-Structure Interaction: Methods and Applications*, Wiley Series in Computational Mechanics, Wiley, 2013.
- [11] R. Becker, M. Braack, A finite element pressure gradient stabilization for the Stokes equations based on local projections, *Calcolo* 38 (4) (2001) 173–199.
- [12] R. Becker, M. Braack, D. Meidner, T. Richter, B. Vexler, The finite element toolkit Gascoigne3d, <http://www.gascoigne.de>.
- [13] Dominik Brands, A. Klawonn, O. Rheinbach, Jörg Schröder, Modelling and convergence in arterial wall simulations using a parallel feti solution strategy, *Comput. Methods Biomech. Biomed. Eng.* 11 (5) (2008) 569–583.
- [14] E. Burman, M.A. Fernández, S. Frei, A Nitsche-based formulation for fluid-structure interactions with contact, *ESAIM: M2AN* 54 (2) (2020) 531–564.
- [15] E. Burman, M.A. Fernández, S. Frei, F.M. Gerosa, A mechanically consistent model for fluid-structure interactions with contact including seepage, *Comput. Methods Appl. Mech. Eng.* 392 (2022) 114637.
- [16] L.J. Caucha, S. Frei, O. Rubio, Finite element simulation of fluid dynamics and CO₂ gas exchange in the alveolar sacs of the human lung, *Comput. Appl. Math.* 37 (5) (2018) 6410–6432.
- [17] C.X. Chen, Y. Ding, J.A. Gear, Numerical simulation of atherosclerotic plaque growth using two-way fluid-structural interaction, *ANZIAM J.* 53 (2012) 278–291.
- [18] P. Crosetto, S. Deparis, G. Fourestey, A. Quarteroni, Parallel algorithms for fluid-structure interaction problems in haemodynamics, *SIAM J. Sci. Comput.* 33 (4) (2011) 1598–1622.
- [19] S. Deparis, D. Forti, G. Grandperrin, A. Quarteroni Facsi, A block parallel preconditioner for fluid-structure interaction in hemodynamics, *J. Comp. Physiol.* 327 (2016) 700–718.
- [20] S.S. Dhawan, R.P. Avati Nanjundappa, J.R. Branch, W.R. Taylor, A.A. Quyyumi, H. Jo, M.C. McDaniel, J. Suo, D. Giddens, H. Samady, Shear stress and plaque development, *Expert Rev. Cardiovasc. Ther.* 8 (4) (2010) 545–556.
- [21] J. Donea, A. Huerta, J.-P. Ponthot, A. Rodríguez-Ferran, *Arbitrary Lagrangian-Eulerian Methods*, chapter 14 John Wiley & Sons, Ltd, 2004.
- [22] T. Dunne, An Eulerian approach to fluid-structure interaction and goal-oriented mesh refinement, *Int. J. Numer. Methods Fluids* 51 (2006) 1017–1039.
- [23] W. E., *Principles of Multiscale Modeling*, Cambridge University Press, 2011.
- [24] W. E, B. Engquist, The heterogeneous multiscale method, *Commun. Math. Sci.* 1 (1) (2003) 87–132.
- [25] C.A. Figueroa, S. Baek, C.A. Taylor, J.D. Humphrey, A computational framework for fluid-solid-growth modeling in cardiovascular simulations, *Comput. Methods Appl. Mech. Eng.* 198 (45–46) (2009) 3583–3602.
- [26] L. Formaggia, A. Quarteroni, A. Veneziani, *Cardiovascular Mathematics: Modeling and Simulation of the Circulatory System*, vol. 1, Springer Science & Business Media, 2010.
- [27] S. Frei, Eulerian finite element methods for interface problems and fluid-structure interactions, PhD thesis, Heidelberg University, 2016, <http://www.ub.uni-heidelberg.de/archiv/21590>.
- [28] S. Frei, A. Heinlein, T. Richter, On temporal homogenization in the numerical simulation of atherosclerotic plaque growth, *PAMM* 21 (1) (2021) e202100055.

- [29] S. Frei, T. Richter, Efficient approximation of flow problems with multiple scales in time, *SIAM J. Multiscale Model. Simul.* 18 (2) (2020) 942–969.
- [30] S. Frei, T. Richter, T. Wick, Long-term simulation of large deformation, mechano-chemical fluid-structure interactions in ALE and fully Eulerian coordinates, *J. Comput. Phys.* 321 (2016) 874–891.
- [31] M.J. Gander, 50 years of time parallel time integration, in: *Multiple Shooting and Time Domain Decomposition*, Springer, 2015.
- [32] Martin J. Gander, Ernst Hairer, Nonlinear convergence analysis for the parareal algorithm, in: Ulrich Langer, Marco Discacciati, David E. Keyes, Olof B. Widlund, Walter Zulehner (Eds.), *Domain Decomposition Methods in Science and Engineering XVII*, Springer Berlin Heidelberg, 2008, pp. 45–56.
- [33] M.W. Gee, U. Küttler, W.A. Wall, Truly monolithic algebraic multigrid for fluid–structure interaction, *Int. J. Numer. Methods Eng.* 85 (8) (2011) 987–1016.
- [34] Sebastian Götschel, Michael Minion, Daniel Ruprecht, Robert Speck, Twelve ways to fool the masses when giving parallel-in-time results, in: Benjamin Ong, Jacob Schroder, Gemma Shipton, Stephanie Friedhoff (Eds.), *Parallel-in-Time Integration Methods*, Springer Proceedings in Mathematics & Statistics, Springer International Publishing, Cham, 2021, pp. 81–94.
- [35] A. Heinlein, A. Klawonn, O. Rheinbach, A parallel implementation of a two-level overlapping Schwarz method with energy-minimizing coarse space based on Trilinos, *SIAM J. Sci. Comput.* 38 (6) (2016) C713–C747.
- [36] A. Heinlein, A. Klawonn, O. Rheinbach, Parallel two-level overlapping Schwarz methods in fluid-structure interaction, in: *Proceedings of ENUMATH 2015*, Springer, 2016, pp. 521–530.
- [37] G.A. Holzapfel, T.C. Gasser, R.W. Ogden, A new constitutive framework for arterial wall mechanics and a comparative study of material models, *J. Elast.* 61 (1) (2000) 1–48.
- [38] J. Hron, S. Turek, A monolithic FEM/multigrid solver for an ALE formulation of fluid-structure interaction with applications in biomechanics, in: *Fluid-Structure Interaction*, Springer, 2006, pp. 146–170.
- [39] F. Kong, X.-C. Cai, Scalability study of an implicit solver for coupled fluid-structure interaction problems on unstructured meshes in 3d, *Int. J. High Perform. Comput. Appl.* 32 (2) (2018) 207–219.
- [40] U. Küttler, M. Gee, C. Förster, A. Comerford, W.A. Wall, Coupling strategies for biomedical fluid–structure interaction problems, *Int. J. Numer. Methods Biomed. Eng.* 26 (3–4) (2010) 305–321.
- [41] L. Lautsch, T. Richter, Error estimation and adaptivity for differential equations with multiple scales in time, *Comput. Methods Appl. Math.* (2021).
- [42] J.-L. Lions, Y. Maday, G. Turinici, A “parareal” in time discretization of PDE’s, *C. R. Acad. Sci., Ser. 1 Math.* 332 (2001) 661–668.
- [43] N. Margenberg, T. Richter, Parallel time-stepping for fluid-structure interactions, *Math. Model. Nat. Phenom.* 16 (20) (2021).
- [44] J. Mizerski, T. Richter, The candy wrapper problem—a temporal multiscale approach for pde/pde systems, in: *Proceedings of ENUMATH 2019*, 2019.
- [45] Benjamin W. Ong, Jacob B. Schroder, Applications of time parallelization, *Comput. Vis. Sci.* 23 (1–4) (2020) 11.
- [46] S. Pozzi, A. Redaelli, C. Vergara, E. Votta, P. Zunino, Mathematical modeling and numerical simulation of atherosclerotic plaque progression based on fluid-structure interaction, *J. Math. Fluid Mech.* 23 (3) (2021) 1–21.
- [47] T. Richter, A fully Eulerian formulation for fluid-structure interactions, *J. Comput. Phys.* 233 (2013) 227–240.
- [48] T. Richter, *Finite Elements for Fluid-Structure Interactions. Models, Analysis and Finite Elements*, Lecture Notes in Computational Science and Engineering, vol. 118, Springer, 2017.
- [49] Thomas Richter, Thomas Wick, *Einführung in die Numerische Mathematik: Begriffe, Konzepte und Zahlreiche Anwendungsbeispiele*, Springer-Verlag, 2017.
- [50] E.K. Rodriguez, A. Hoger, A.D. McCulloch, Stress-dependent finite growth in soft elastic tissues, *J. Biomech.* 4 (1994) 455–467.
- [51] D. Ruprecht, Wave propagation characteristics of parareal, *Comput. Vis. Sci.* 19 (1) (2018) 1–17.
- [52] Ruprecht Daniel, Shared memory pipelined parareal, in: *Euro-Par 2017: Parallel Processing: 23rd International Conference on Parallel and Distributed Computing*, Santiago de Compostela, Spain, August 28–September 1, 2017, *Proceedings 23*, Springer, 2017, pp. 669–681.
- [53] A. Heinlein, S. Frei, Efficient coarse correction for parallel time-stepping in plaque growth simulations, in: *Proceedings of the 8th ECCOMAS Congress, 2022*, https://www.scipedia.com/public/Frei_et_al_2022a.
- [54] T. Silva, W. Jäger, M. Neuss-Radu, A. Sequeira, Modeling of the early stage of atherosclerosis with emphasis on the regulation of the endothelial permeability, *J. Theor. Biol.* 496 (2020) 110229.
- [55] F. Sonner, *Temporal Multiscale Methods for a Model of Atherosclerosis*, PhD thesis, FAU Erlangen-Nürnberg, 2021.
- [56] M.P. Thon, A. Hemmler, A. Glinzer, M. Mayr, M. Wildgruber, A. Zerneck-Madsen, M.W. Gee, A multiphysics approach for modeling early atherosclerosis, *Biomech. Model. Mechanobiol.* 17 (3) (2018) 617–644.
- [57] A. Wahle, J.J. Lopez, M.E. Olszewski, S.C. Vigmostad, K.B. Chandran, J.D. Rossen, M. Sonka, Plaque development, vessel curvature, and wall shear stress in coronary arteries assessed by x-r. Angiography and intravascular ultrasound, *Med. Image Anal.* 10 (4) (2006) 615–631.
- [58] Y. Wu, X.-C. Cai, A fully implicit domain decomposition based ale framework for three-dimensional fluid–structure interaction with application in blood flow computation, *J. Comput. Phys.* 258 (2014) 524–537.
- [59] Y. Yang, W. Jäger, M. Neuss-Radu, T. Richter, Mathematical modeling and simulation of the evolution of plaques in blood vessels, *J. Math. Biol.* (2014) 1–24.
- [60] Y. Yang, T. Richter, W. Jäger, M. Neuss-Radu, An ALE approach to mechano-chemical processes in fluid-structure interactions, *Int. J. Numer. Methods Fluids* 84 (4) (2017) 199–220.