

## STOCHASTIC PARAREAL: AN APPLICATION OF PROBABILISTIC METHODS TO TIME-PARALLELIZATION\*

KAMRAN PENTLAND<sup>†</sup>, MASSIMILIANO TAMBORRINO<sup>‡</sup>, DEBASMITA SAMADDAR<sup>§</sup>,  
AND LYNTON C. APPEL<sup>§</sup>

**Abstract.** Parareal is a well-studied algorithm for numerically integrating systems of time-dependent differential equations by parallelizing the temporal domain. Given approximate initial values at each temporal subinterval, the algorithm locates a solution in a fixed number of iterations using a predictor-corrector, stopping once a tolerance is met. This iterative process combines solutions located by inexpensive (coarse resolution) and expensive (fine resolution) numerical integrators. In this paper, we introduce a stochastic parareal algorithm aimed at accelerating the convergence of the deterministic parareal algorithm. Instead of providing the predictor-corrector with a deterministically located set of initial values, the stochastic algorithm samples initial values from dynamically varying probability distributions in each temporal subinterval. All samples are then propagated in parallel using the expensive integrator. The set of sampled initial values yielding the most continuous (smoothest) trajectory across consecutive subintervals are fed into the predictor-corrector, converging in fewer iterations than the deterministic algorithm with a given probability. The performance of the stochastic algorithm, implemented using various probability distributions, is illustrated on low-dimensional systems of ordinary differential equations (ODEs). We provide numerical evidence that when the number of sampled initial values is large enough, stochastic parareal converges almost certainly in fewer iterations than the deterministic algorithm, maintaining solution accuracy. Given its stochastic nature, we also highlight that multiple simulations of stochastic parareal return a distribution of solutions that can represent a measure of uncertainty over the ODE solution.

**Key words.** parareal, time-parallel integration, probabilistic numerics, sampling-based solver, multivariate copulas

**AMS subject classifications.** 65L05, 65Y05, 65C99

**DOI.** 10.1137/21M1414231

**1. Introduction.** In its most basic form, parallel computing is the process by which an algorithm is partitioned into a number of subproblems that can be solved simultaneously without prior knowledge of each other. More widespread parallelism is becoming increasingly necessary in many different fields to reduce the computational burden and thus overcome the physical limitations (i.e., space, power usage, clock speeds, cooling, and financial costs) arising on machine hardware. Complex models in science often involve solving large systems of ordinary or partial differential equations (ODEs or PDEs) which, in the case of spatio-temporal PDEs, can be parallelized using existing domain decomposition methods. We refer to [31] for an overview. Although

---

\*Received by the editors April 21, 2021; accepted for publication (in revised form) March 30, 2022; published electronically July 7, 2022.

<https://doi.org/10.1137/21M1414231>

**Funding:** This work was partially supported by the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 and 2019-2020 under grant agreement 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. The work of the first author was supported by the Engineering and Physical Sciences Research Council through the MathSys II CDT (grant EP/S022244/1) as well as the Culham Centre for Fusion Energy.

<sup>†</sup>Mathematics Institute, University of Warwick, Coventry, CV4 7AL, United Kingdom (kamran.pentland@warwick.ac.uk).

<sup>‡</sup>Department of Statistics, University of Warwick, Coventry, CV4 7AL, United Kingdom (massimiliano.tamborrino@warwick.ac.uk).

<sup>§</sup>Culham Centre for Fusion Energy, Culham Science Centre, Abingdon, Oxfordshire, OX14 3DB, United Kingdom (Debasmita.Samaddar@ukaea.uk, Lynton.Appel@ukaea.uk).

very efficient for high-dimensional systems, these methods are reaching scale-up limits, and wallclock speeds often bottleneck due to the temporal integration. For instance, modern algorithms used to simulate edge localized modes in turbulent fusion plasmas can take anywhere between 100 and 200 days to integrate over a time interval of just one second [28].

These sequential bottlenecks in time have motivated the research and development of *time-parallel* integration methods for systems of ODEs and PDEs over the last 55 years or so (see [5, 10, 25] for reviews). These methods provide a way to integrate initial value problems (IVPs) over long time intervals where solutions would be unobtainable using existing sequential algorithms. One approach to integrate in a nonsequential manner, similar to spatial parallelization, is to discretize the time interval into  $N$  subintervals upon which  $N$  subproblems are solved in parallel using existing numerical methods. The solution at a given time step, however, depends upon the solution at the previous step(s). This creates a problem prior to the parallel integration as  $N - 1$  of the  $N$  initial values (from which to begin integration in each subinterval) are unknown a priori. Existing algorithms that attempt to locate these  $N - 1$  initial values by direct or iterative means have been collectively referred to as multiple shooting methods [4, 6, 17, 23, 27].

Our focus is on one multiple shooting method in particular: *parareal* [17]. This algorithm uses a combination of coarse and fine grained (in time) numerical integrators to provide parallel speedup and has been tested successfully on problems spanning molecular [1] and fluid dynamics [9, 13, 33] to stochastic differential equations [8, 16] and fusion plasma simulations [28, 29]. The popularity of parareal is due to its relatively straightforward implementation and demonstrable effectiveness in speeding up the integration of IVPs. The algorithm iteratively locates a numerical solution at each subinterval using a predictor-corrector scheme, derived by discretizing the Newton–Raphson method. The algorithm stops once a tolerance is met with convergence guaranteed under certain mathematical conditions—more detail on the algorithm is provided in section 2. Much work has gone into analyzing the numerical convergence of this method [2, 3, 12, 20], combining it with spatial decomposition techniques [21], and developing variants that utilize idle processors [7] and adaptive time-stepping [19]. However, given fixed fine and coarse integrators, the method itself is strictly deterministic and little work has gone into trying to reduce the number of iterations  $k_d \in \{1, \dots, N\}$  that parareal takes to solve a particular IVP. Our primary focus is on reducing  $k_d$  and henceforth we refer to it as the “convergence rate.” In scenarios where parareal struggles to converge in  $k_d \ll N$  iterations, reducing  $k_d$  by even a few iterations, say, to  $k_s < k_d$ , can lead to significant parallel gains (roughly a factor  $k_d/k_s$ ), enabling faster numerical integration.

The purpose of this work is to extend the parareal algorithm using probabilistic methods to converge in fewer than  $k_d$  iterations for a given system of ODEs. To achieve this, we introduce a *stochastic* parareal algorithm. Instead of integrating forward in time from a single initial value (given by the predictor-corrector) in each subinterval, a prespecified probability distribution (see sampling rules in subsection 3.2) is used to generate  $M$  candidate initial values that are *simultaneously* integrated forward in time in parallel. At each subinterval, one optimal initial value (from the  $M$  samples) is selected sequentially such that the most continuous trajectory is chosen across consecutive subintervals. These initial values are then fed directly into the predictor-corrector—the idea being that this stochastically generated set of values provides a better guess to the solution than those found purely deterministically. For example, suppose running the predictor-corrector with initial values  $\mathbf{x}^{(0)}$  yields convergence in

$k_d$  iterations, generating the sequence of solutions  $\{\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \dots, \mathbf{x}^{(k_d)}\}$ . Instead of starting with  $\mathbf{x}^{(0)}$ , suppose we sample initial values from a probability distribution and choose some “better” starting point which is close to, say,  $\mathbf{x}^{(i)}$  for some  $i \in \{1, \dots, k_d - 1\}$ . Then the sequence generated by the predictor-corrector would instead be approximately  $\{\mathbf{x}^{(i)}, \mathbf{x}^{(i+1)}, \dots, \mathbf{x}^{(k_d)}\}$ , converging in  $k_d - i$  iterations. Therefore, given a fixed number of samples  $M$ , the stochastic parareal algorithm will converge in fewer than  $k_d$  iterations with some nonzero probability.

This idea of propagating multiple initial values in each subinterval in parallel is not new. In the first major work proposing the time-parallel integration of IVPs, Nievergelt [23] discussed choosing  $M$  initial values deterministically. The method for determining the solution from this ensemble of trajectories was to combine two of the  $M$  samples in each subinterval using an interpolation coefficient determined from the preceding subinterval. While an excellent first incursion into the field, for nonlinear problems this direct method suffered from interpolation errors and questions remained on how to efficiently scale this up to systems of ODEs. We generalize the original idea of Nievergelt by combining it with the parareal algorithm, generating the  $M$  initial values at each subinterval using probability distributions based on information known about the solutions at the current iteration.

The paper is organized as follows. In section 2, we recall the parareal algorithm and its properties from a multiple shooting viewpoint, including a known modification that will enable us to extend the algorithm to incorporate stochastic sampling. In section 3, we introduce the stochastic parareal algorithm. Following an explanation of its key features, we elucidate how a variety of different sampling rules can be flexibly interchanged in order to carry out the sampling. In section 4, we conduct numerical experiments to illustrate the performance of stochastic parareal against its deterministic counterpart using the different sampling rules. Findings are presented for three ODE systems of increasing complexity, with two additional examples given in the supplementary materials (M141423SupMat.pdf [local/web 1.21MB]). Results indicate that for sufficiently many samples, stochastic parareal almost certainly beats the convergence rate of the deterministic algorithm and generates (stochastic) solutions of comparable numerical accuracy. For the multivariate ODEs, results show that performance is improved by generating correlated, as opposed to uncorrelated, random samples. In section 5, conclusions are drawn and avenues for future research are discussed.

**2. The parareal algorithm.** Following previously outlined descriptions of the parareal algorithm [12, 17], henceforth referred to as “parareal” or  $\mathcal{P}$ , consider a system of  $d \in \mathbb{N}$  ODEs

$$(2.1) \quad \frac{d\mathbf{u}}{dt} = \mathbf{f}(\mathbf{u}(t), t) \quad \text{on } t \in [T_0, T_N], \quad \text{with } \mathbf{u}(T_0) = \mathbf{u}^0,$$

where  $\mathbf{f} : \mathbb{R}^d \times [T_0, T_N] \rightarrow \mathbb{R}^d$  is a smooth multivariate (possibly nonlinear) function,  $\mathbf{u} : [T_0, T_N] \rightarrow \mathbb{R}^d$  the time-dependent vector solution, and  $\mathbf{u}^0 \in \mathbb{R}^d$  the initial values at  $T_0$ . Decompose the time domain into  $N$  subintervals such that  $[T_0, T_N] = [T_0, T_1] \cup \dots \cup [T_{N-1}, T_N]$ , with each subinterval taking fixed length  $\Delta T := T_n - T_{n-1}$  for  $n = 1, \dots, N$  (see Figure 1 for a schematic).

Having partitioned the time domain  $N$  smaller subproblems

$$(2.2) \quad \frac{d\mathbf{u}_n}{dt} = \mathbf{f}(\mathbf{u}_n(t|\mathbf{U}_n), t) \quad \text{on } t \in [T_n, T_{n+1}], \quad \text{with } \mathbf{u}_n(T_n) = \mathbf{U}_n,$$

for  $n = 0, \dots, N - 1$  can theoretically be solved in parallel on  $N$  processors, henceforth denoted  $P_1, \dots, P_N$ . Each solution  $\mathbf{u}_n(t|\mathbf{U}_n)$  is defined over  $[T_n, T_{n+1}]$  given the initial



FIG. 1. Schematic of the time domain decomposition. Three levels of temporal discretisation are shown: subintervals (size  $\Delta T$ ), coarse intervals (size  $\delta T$ ), and fine intervals (size  $\delta t$ ). Note how the discretizations align with one another and that  $\delta t \ll \delta T \leq \Delta T$  in our implementation.

values  $\mathbf{U}_n \in \mathbb{R}^d$  at  $t = T_n$ . Note, however, that only the initial value  $\mathbf{U}_0 = \mathbf{u}_0(T_0) = \mathbf{u}^0$  is known, whereas the rest ( $\mathbf{U}_n$  for  $n \geq 1$ ) need to be determined before (2.2) can be solved in parallel. These initial values must satisfy the continuity conditions

$$(2.3) \quad \mathbf{U}_0 = \mathbf{u}^0 \quad \text{and} \quad \mathbf{U}_n = \mathbf{u}_{n-1}(T_n | \mathbf{U}_{n-1}) \quad \text{for} \quad n = 1, \dots, N,$$

which form a (nonlinear) system of  $N + 1$  equations that ensure solutions match at each  $T_n \forall n \geq 1$ . System (2.3) is solved for  $\mathbf{U}_n$  using the Newton–Raphson method to form the iterative system

$$(2.4a) \quad \mathbf{U}_0^{k+1} = \mathbf{u}^0,$$

$$(2.4b) \quad \mathbf{U}_n^{k+1} = \mathbf{u}_{n-1}(T_n | \mathbf{U}_{n-1}^k) + \frac{\partial \mathbf{u}_{n-1}}{\partial \mathbf{U}_{n-1}}(T_n | \mathbf{U}_{n-1}^k) [\mathbf{U}_{n-1}^{k+1} - \mathbf{U}_{n-1}^k]$$

for  $n = 1, \dots, N$ , where  $k = 0, 1, 2, \dots$  is the iteration number. This system contains the unknown solutions  $\mathbf{u}_n$  and their partial derivatives, which even if known would be computationally expensive to calculate.

To solve (2.4) and evaluate  $\mathbf{u}_n$  at discrete times,  $\mathcal{P}$  utilizes two numerical integrators. The first is a numerically fast *coarse integrator*  $\mathcal{G}$  that integrates over  $[T_n, T_{n+1}]$  using initial values  $\mathbf{U}_n$ . The second is a *fine integrator*  $\mathcal{F}$  that runs much slower compared to  $\mathcal{G}$  but has much greater numerical accuracy, chosen such that it integrates over the same interval  $[T_n, T_{n+1}]$  with initial values  $\mathbf{U}_n$ . In our implementation, the difference between fast and slow integration is guaranteed by setting the time steps for  $\mathcal{G}$  and  $\mathcal{F}$  as  $\delta T$  and  $\delta t$ , respectively, with  $\delta t \ll \delta T$  (recall Figure 1). The key principle is that if used to integrate (2.1) over  $[T_0, T_N]$  serially,  $\mathcal{F}$  would take an infeasible amount of computational time, hence the need to use  $\mathcal{P}$  in the first place. Therefore  $\mathcal{G}$  is permitted to run serially across multiple subintervals rapidly, whereas the slower solver  $\mathcal{F}$  is only permitted to run in parallel on subintervals. This is a strict requirement for running  $\mathcal{P}$ , else numerical speedup will not be achieved.

Given that (2.4a) is known a priori  $\forall k$ , Lions, Maday, and Turinici [17] approximate the first term in (2.4b) using the fine solver  $\mathcal{F}(\mathbf{U}_{n-1}^k)$  and the second term by a coarse finite difference of the derivative using  $\mathcal{G}(\mathbf{U}_{n-1}^{k+1}) - \mathcal{G}(\mathbf{U}_{n-1}^k)$ . One could instead approximate the derivative using a fine finite difference  $\mathcal{F}(\mathbf{U}_{n-1}^{k+1}) - \mathcal{F}(\mathbf{U}_{n-1}^k)$ ; however, (2.4b) then becomes a sequential calculation using just  $\mathcal{F}$ —exactly what we wish to avoid. Assuming  $\mathcal{G}$  meets the conditions required for (2.4b) to converge (see subsection 2.1), using the coarse approximation becomes reasonable and enables the parallel computation of (2.4b) [12]. The result is that an initial guess for the initial values  $\mathbf{U}_n^0$  (found using  $\mathcal{G}$ ) is improved at successive parareal iterations  $k$  using the predictor-corrector update rule

$$(2.5) \quad \mathbf{U}_n^{k+1} = \mathcal{G}(\mathbf{U}_{n-1}^{k+1}) + \mathcal{F}(\mathbf{U}_{n-1}^k) - \mathcal{G}(\mathbf{U}_{n-1}^k) \quad \text{for} \quad n = 1, \dots, N.$$

**Algorithm 1** Parareal ( $\mathcal{P}$ )

---

```

1: Set counters  $k = I = 0$  and define  $\mathbf{U}_n^k$ ,  $\hat{\mathbf{U}}_n^k$ , and  $\tilde{\mathbf{U}}_n^k$  as the predictor-corrector,
   coarse, and fine solutions at the  $n$ th time step and  $k$ th iteration, respectively (note
   that  $\mathbf{U}_0^k = \hat{\mathbf{U}}_0^k = \tilde{\mathbf{U}}_0^k = \mathbf{u}^0 \forall k$ ).
2: // Calculate initial values at the start of each subinterval  $T_n$  by running  $\mathcal{G}$  serially
   on processor  $P_1$ .
3: for  $n = 1$  to  $N$  do
4:    $\hat{\mathbf{U}}_n^0 = \mathcal{G}(\hat{\mathbf{U}}_{n-1}^0)$ 
5:    $\mathbf{U}_n^0 = \hat{\mathbf{U}}_n^0$ 
6: end for
7: for  $k = 1$  to  $N$  do
8:   // Propagate the predictor-corrector solutions (from iteration  $k - 1$ ) on each
   subinterval by running  $\mathcal{F}$  in parallel on processors  $P_{I+1}, \dots, P_N$ .
9:   for  $n = I + 1$  to  $N$  do
10:     $\tilde{\mathbf{U}}_n^{k-1} = \mathcal{F}(\mathbf{U}_{n-1}^{k-1})$ 
11:   end for
12:   // Propagate the predictor-corrector solution (at iteration  $k$ ) with  $\mathcal{G}$  on any
   available processor. Then correct this value using coarse and fine solutions
   obtained during iteration  $k - 1$  (this step cannot be carried out in parallel).
13:   for  $n = I + 1$  to  $N$  do
14:     $\hat{\mathbf{U}}_n^k = \mathcal{G}(\mathbf{U}_{n-1}^k)$ 
15:     $\mathbf{U}_n^k = \tilde{\mathbf{U}}_n^{k-1} + \hat{\mathbf{U}}_n^k - \hat{\mathbf{U}}_n^{k-1}$ 
16:   end for
17:   // Check whether the stopping criterion is met, saving all solutions up to time
   step  $T_I$  before the next iteration. If tolerance is met for all time steps, the
   algorithm stops.
18:    $I = \max_{n \in \{I+1, \dots, N\}} \|\mathbf{U}_n^k - \mathbf{U}_n^{k-1}\|_\infty < \varepsilon \forall i < n$ 
19:   if  $I == N$  then
20:     return  $k, \mathbf{U}^k$ 
21:   end if
22: end for

```

---

Pseudocode for an implementation of  $\mathcal{P}$  is given in Algorithm 1 alongside a graphical description of the first iteration in Figure 2. During the “zeroth” iteration,  $\mathcal{G}$  is applied to  $\mathbf{u}^0$ , generating initial values  $\hat{\mathbf{U}}_n^0 \forall n \in \{1, \dots, N\}$  sequentially (lines 1–6). Immediately following,  $\mathcal{F}$  is run in parallel from these initial values to generate a more accurate set of fine solution states  $\tilde{\mathbf{U}}_n^0$  (lines 8–11). Next,  $\mathcal{G}$  is run in the first time subinterval, from the known initial value, to “predict” the solution  $\hat{\mathbf{U}}_1^1$  at  $T_1$ . It is subsequently “corrected” using rule (2.5), and the predictor-corrector solution  $\mathbf{U}_1^1$  is found at  $T_1$ . This prediction and correction process (lines 12–16) repeats sequentially until  $\mathbf{U}_n^1$  is found  $\forall n \in \{1, \dots, N\}$ . After checking if the stopping criteria has been met (lines 17–21), the algorithm either starts the next iteration or stops.

**2.1. Stopping criteria and other properties.** The algorithm is said to have converged up to time  $T_I$  if

$$(2.6) \quad \|\mathbf{U}_n^k - \mathbf{U}_n^{k-1}\|_\infty < \varepsilon \quad \forall n \leq I$$




FIG. 2. First iteration of the parareal algorithm to numerically evaluate the solution of a scalar ODE, either available or obtained via the fine solver (black line). The first runs of  $\mathcal{G}$  and  $\mathcal{F}$  are given in yellow and blue, respectively, and the second run of  $\mathcal{G}$  in red. The red dots represent the solution after applying the predictor-corrector (2.5).

for some small fixed tolerance  $\varepsilon$ —noting that  $\|\cdot\|_\infty$  denotes the usual infinity norm [20, 29]. Taking the relative, instead of absolute, errors in (2.6) would also be appropriate; however, in our numerical experiments the results (not reported) did not change. Once  $I = N$ , we say  $\mathcal{P}$  has taken  $k_d = k$  iterations to converge, yielding a solution with numerical accuracy of the order of the  $\mathcal{F}$  solver. In its original formulation,  $\mathcal{P}$  iteratively improves the solution across *all* time steps, regardless of whether they have converged or not, up until the tolerance has been met  $\forall T_n$ . The modified formulation presented here, however, only iterates on the solutions for the unconverged subintervals [7, 29]. This has no effect on the convergence rate  $k_d$  and becomes especially important once we introduce the stochastic modifications in section 3.

It should be clear that *at least* one subinterval will converge during each iteration  $k$ , as  $\mathcal{F}$  will be run directly from a converged initial value. Therefore it will take at most  $k_d = N$  iterations for  $\mathcal{P}$  to converge, equivalent to running  $\mathcal{F}$  over the  $N$  subintervals serially (yielding zero parallel speedup). To achieve significant parallel speedup, multiple subintervals need to converge during an iteration so that  $k_d \ll N$ . Assuming  $\mathcal{G}$  takes a negligible amount of time to run compared to the parallel components (along with any other serial computations), an approximate upper bound on the speedup achieved by  $\mathcal{P}$  is  $N/k_d$ .

One challenge in attempting to minimize  $k_d$ , hence the overall runtime of  $\mathcal{P}$ , is to identify optimal solvers  $\mathcal{F}$  and  $\mathcal{G}$  for implementation. Whereas  $\mathcal{F}$  is assumed to have high accuracy and be computationally expensive to run,  $\mathcal{G}$  must be chosen such that it runs significantly faster (usually orders of magnitude) than  $\mathcal{F}$  while being sufficiently numerically accurate to converge in as few iterations as possible. Usually  $\mathcal{G}$  is chosen such that its solutions have lower numerical accuracy, coarser temporal resolution, and/or reduced physics/coarser spatial resolution (when solving PDEs). There is currently no rigorous method for choosing the two solvers; however, they should be chosen such that the ratio between the time taken to run  $\mathcal{F}$  over  $[T_n, T_{n+1}]$  and the time taken to run  $\mathcal{G}$  over the same interval is large.

In this paper, we fix both  $\mathcal{F}$  and  $\mathcal{G}$  as explicit<sup>1</sup> fourth-order Runge–Kutta methods (henceforth RK4) for ease of implementation. More detailed analysis on the

<sup>1</sup>When solving the stiff ODE systems in section 4, the initial values at  $t = 0$  are chosen such that explicit methods work. For greater numerical stability, one could alternatively use implicit solvers at extra computational cost.

TABLE 1  
*Additional notation used to describe the stochastic parareal algorithm.*

Notation	Description
$n$	Index of the time subinterval $T_n$ , $n = 0, \dots, N$ .
$k$	Iteration number of $\mathcal{P}$ and $\mathcal{P}_s$ , $k = 1, \dots, N$ .
$k_s$	Total iterations taken for stochastic parareal to stop and return a solution, $k_s \in \{1, \dots, N\}$ .
$M$	Number of random samples taken at each $T_n$ .
$\Phi_n^{k-1}$	The $d$ -dimensional probability distribution used to sample initial values at time $T_n$ and iteration $k$ .
$\alpha_{n,m}^{k-1}$	The $m$ th $d$ -dimensional sample from distribution $\Phi_n^{k-1}$ at time $T_n$ and iteration $k$ , $m = 1, \dots, M$ .
$\hat{\alpha}_n^{k-1}$	The selected $d$ -dimensional sample from $\alpha_{n,1}^{k-1}, \dots, \alpha_{n,M}^{k-1}$ at time $T_n$ and iteration $k$ .
$\mu_n^{k-1}$	The $d$ -dimensional vector of marginal means at time $T_n$ and iteration $k$ .
$\sigma_n^{k-1}$	The $d$ -dimensional vector of marginal standard deviations at time $T_n$ and iteration $k$ .
$\rho_{n,i,j}^{k-1}$	Pairwise correlations coefficients between $i$ th and $j$ th components of the $M$ fine resolution propagated samples at time $T_n$ and iteration $k$ .
$\mathbf{R}_n^{k-1}$	The symmetric positive semidefinite $d \times d$ correlation matrix with elements $\rho_{n,i,j}^{k-1}$ at time $T_n$ and iteration $k$ .
$\mathbb{I}$	The $d \times d$ identity matrix.
$\Sigma_n^{k-1}$	The symmetric positive semidefinite $d \times d$ covariance matrix at time $T_n$ and iteration $k$ .

mathematical conditions required for  $\mathcal{P}$  to converge, i.e., for the numerical solution  $\mathbf{U}_n$  to approach the fine solution, can be found in [11, 12, 17, 21].

**3. A stochastic parareal algorithm.** In this section, we introduce the stochastic parareal algorithm, an extension of parareal that incorporates randomness and utilizes its well-studied deterministic convergence properties to locate a solution in  $k_s < k_d$  iterations. A summary of additional notation required to describe the algorithm, henceforth referred to as  $\mathcal{P}_s$ , is provided in Table 1.

The idea behind  $\mathcal{P}_s$  is to sample  $M$  vectors of initial values  $\alpha_{n,1}^k, \dots, \alpha_{n,M}^k$ , at each unconverged subinterval  $T_n$ , in the neighborhood of the current predictor-corrector solution  $\mathbf{U}_n^k$  from a given probability distribution (with sufficiently broad support) and propagate them all in parallel using  $\mathcal{F}$ . Given a sufficient number of samples is taken, one will be closer (in the Euclidean sense) to the true root that (2.5) is converging toward. Among them, we select an optimal  $\hat{\alpha}_n^k$  by identifying which samples generate the most continuous trajectory, at the fine resolution, in phase space across  $[T_0, T_N]$ . Therefore, at each iteration, we stochastically “jump” toward more accurate initial values and then feed them into the predictor-corrector (2.5). For increasing values of  $M$ , the convergence rate  $k_s$  will decrease, satisfying  $k_s < k_d$  with probability one—shown numerically in section 4.

**3.1. The algorithm.** Suppose again we aim to solve system (2.1), adopting the same conditions, properties, and notation as discussed in section 2.  $\mathcal{P}_s$  follows the first iteration ( $k = 1$ ) of  $\mathcal{P}$  identically—see line 1 of Algorithm 2. This is because information about initial values at the different temporal resolutions (i.e., results from  $\mathcal{F}$  and  $\mathcal{G}$ ) are required to construct the appropriate probability distributions for sampling. Following the convergence check, we assume (for the purposes of explaining the stochastic iterations) that only the first subinterval  $[T_0, T_1]$  converged during

**Algorithm 2** Stochastic parareal ( $\mathcal{P}_s$ )

---

```

1: Run  $\mathcal{P}$  (Algorithm 1) until the end of iteration  $k = 1$ .
2: for  $k = 2$  to  $N$  do
3:   // Calculate correlation matrices (only if  $d > 1$ ).
4:    $\mathbf{R}_n^{k-1} = \mathbb{I} \forall n$ 
5:   if  $k \geq 3$  then
6:     for  $n = I + 1$  to  $N - 1$  do
7:       Calculate  $\mathbf{R}_n^{k-1}$  using  $\mathcal{F}(\boldsymbol{\alpha}_{n-1,1}^{k-2}), \dots, \mathcal{F}(\boldsymbol{\alpha}_{n-1,M}^{k-2})$  (recall (3.1)).
8:     end for
9:   end if
10:  // Initial value sampling and propagation. Both line 11 and the nested for loop
    below (lines 12–22) run in parallel on  $P_1, \dots, P_{M(N-1-I)+1}$ , i.e., all runs of  $\mathcal{F}$ 
    must be in parallel.
11:   $\tilde{\mathbf{U}}_{I+1}^{k-1} = \mathcal{F}(\mathbf{U}_I^{k-1})$  // propagate converged initial value at  $T_I$  on  $P_1$ 
12:  for  $n = I + 1$  to  $N - 1$  do
13:    for  $m = 1$  to  $M$  do
14:      if  $m == 1$  then
15:         $\boldsymbol{\alpha}_{n,1}^{k-1} = \mathbf{U}_n^{k-1}$  // first “sample” is fixed to the predictor-corrector
16:         $\tilde{\mathbf{U}}_{n+1,1} = \mathcal{F}(\boldsymbol{\alpha}_{n,1}^{k-1})$  // temporarily store propagated values
17:      else
18:         $\boldsymbol{\alpha}_{n,m}^{k-1} \sim \Phi_n^{k-1}$  // sample initial value randomly, see subsection 3.2
19:         $\tilde{\mathbf{U}}_{n+1,m} = \mathcal{F}(\boldsymbol{\alpha}_{n,m}^{k-1})$ 
20:      end if
21:    end for
22:  end for
23:  // Select the most continuous fine trajectory from the ensemble sequentially.
24:  for  $n = I + 1$  to  $N - 1$  do
25:     $J = \operatorname{argmin}_{j \in \{1, \dots, M\}} \|\boldsymbol{\alpha}_{n,j}^{k-1} - \tilde{\mathbf{U}}_n^{k-1}\|_2$ 
26:     $\hat{\boldsymbol{\alpha}}_n^{k-1} = \boldsymbol{\alpha}_{n,J}^{k-1}$  // store optimal initial value
27:     $\tilde{\mathbf{U}}_{n+1}^{k-1} = \tilde{\mathbf{U}}_{n+1,J}$  // store most optimal fine trajectories
28:  end for
29:  // Run  $\mathcal{G}$  from the optimal samples (can run in parallel).
30:  for  $n = I + 1$  to  $N - 1$  do
31:     $\hat{\mathbf{U}}_{n+1}^{k-1} = \mathcal{G}(\hat{\boldsymbol{\alpha}}_n^{k-1})$ 
32:  end for
33:  // Predict and correct the initial values.
34:  for  $n = I + 1$  to  $N$  do
35:     $\hat{\mathbf{U}}_n^k = \mathcal{G}(\mathbf{U}_{n-1}^k)$ 
36:     $\mathbf{U}_n^k = \tilde{\mathbf{U}}_n^{k-1} + \hat{\mathbf{U}}_n^k - \hat{\mathbf{U}}_n^{k-1}$ 
37:  end for
38:  // Check whether the stopping criterion is met.
39:   $I = \max_{n \in \{I+1, \dots, N\}} \|\mathbf{U}_n^k - \mathbf{U}_n^{k-1}\|_\infty < \varepsilon \forall i < n$ 
40:  if  $I == N$  then
41:    return  $k, \mathbf{U}^k$ 
42:  end if
43: end for

```

---



$k = 1$ , leaving  $N - 1$  unconverged subintervals. At this point we know the most up-to-date predictor-corrector solutions  $\mathbf{U}_n^1 \forall n \in \{1, \dots, N\}$  and the stochastic iterations can begin (henceforth  $k = 2$ ).

At any unconverged  $T_n$  ( $n > 1$ ), we sample  $M$  vectors of initial values, denoted  $\boldsymbol{\alpha}_{n,m}^{k-1}$  for  $m = 1, \dots, M$ . The first sample is fixed as the predictor-corrector  $\mathbf{U}_n^{k-1}$ , to ensure that  $\mathcal{P}_s = \mathcal{P}$  when  $M = 1$ . The other  $M - 1$  initial values are sampled from a prespecified  $d$ -dimensional probability distribution  $\Phi_n^{k-1}$  with finite marginal means  $\boldsymbol{\mu}_n^{k-1} = (\mu_{n_1}^{k-1}, \dots, \mu_{n_d}^{k-1})^T$ , marginal standard deviations  $\boldsymbol{\sigma}_n^{k-1} = (\sigma_{n_1}^{k-1}, \dots, \sigma_{n_d}^{k-1})^T$ , and correlation structure given by the matrix  $\mathbf{R}_n^{k-1}$ . These quantities depend upon the information available at iteration  $k - 1$ , i.e., a combination of  $\mathbf{U}_n^{k-1}$ ,  $\mathcal{F}(\mathbf{U}_n^{k-1})$ ,  $\mathcal{G}(\mathbf{U}_n^{k-1})$ , and  $\mathcal{G}(\mathbf{U}_n^{k-2})$ ; see, subsection 3.2. The correlation matrix  $\mathbf{R}_n^{k-1}$  is introduced to take into account the dependence between components of the ODE system (lines 3–9). The elements of  $\mathbf{R}_n^{k-1}$ , for  $k \geq 3$ , are defined using the Pearson correlation coefficient

$$(3.1) \quad \rho_{n_{i,j}}^{k-1} = \frac{\sum_{m=1}^M (x_m^{(i)} - \bar{x}^{(i)})(x_m^{(j)} - \bar{x}^{(j)})}{\sqrt{\sum_{m=1}^M (x_m^{(i)} - \bar{x}^{(i)})^2} \sqrt{\sum_{m=1}^M (x_m^{(j)} - \bar{x}^{(j)})^2}}, \quad i, j \in \{1, \dots, d\},$$

where

$$(3.2) \quad x_m^{(i)} = \mathcal{F}(\boldsymbol{\alpha}_{n-1,m}^{k-2})(i), \quad \bar{x}^{(i)} = \frac{1}{M} \sum_{m=1}^M x_m^{(i)},$$

and  $\mathcal{F}(\boldsymbol{\alpha}_{n-1,m}^{k-2})(i)$  denotes the  $i$ th element of  $\mathcal{F}(\boldsymbol{\alpha}_{n-1,m}^{k-2})$ . The coefficients  $\rho_{n_{i,j}}^{k-1}$  in (3.1) are essentially the estimated pairwise correlation coefficients of the  $M$   $d$ -dimensional fine resolution propagations of the sampled initial values at  $T_n$  from the previous iteration, i.e.,  $\mathcal{F}(\boldsymbol{\alpha}_{n-1,1}^{k-2}), \dots, \mathcal{F}(\boldsymbol{\alpha}_{n-1,M}^{k-2})$ . Note that other types of linear correlation coefficient could also be chosen. Since each  $\mathcal{F}(\boldsymbol{\alpha}_{n-1,m}^{k-2})$  is not available at iteration  $k = 2$ , we set  $\mathbf{R}_n^{k-1} = \mathbb{I}$  for  $k = 2$ , i.e., we sample from a multivariate distribution with uncorrelated components.

Following this, the sampling and subsequent propagation using  $\mathcal{F}$  can begin in parallel (lines 10–22). Given the solution between  $[T_0, T_1]$  has converged,  $\mathcal{F}$  will run from the converged initial value at  $T_1$ , with sampling starting from  $T_2$  onward (see Figure 3). All sampled initial values are then propagated forward in parallel using  $\mathcal{F}$ , requiring at least  $M(N - 2) + 1$  processors ( $M$  samples times  $N - 2$  unconverged subintervals plus running  $\mathcal{F}$  once in  $[T_1, T_2]$ ).

Of the  $M$  sampled initial values at each  $T_n$  ( $n > 1$ ), only one is retained, denoted by  $\hat{\boldsymbol{\alpha}}_n^{k-1}$ , chosen such that it minimizes the Euclidean distance between the fine solution and the sampled values (lines 23–28). To do this, start from the converged initial values at  $T_2$  given by the fine solver:  $\mathcal{F}(\mathbf{U}_1^{k-1})$ . Calculate the Euclidean distance between  $\mathcal{F}(\mathbf{U}_1^{k-1})$  and each of the  $M$  samples  $\boldsymbol{\alpha}_{2,1}^{k-1}, \dots, \boldsymbol{\alpha}_{2,M}^{k-1}$ . The sample minimizing this distance is chosen as  $\hat{\boldsymbol{\alpha}}_2^{k-1}$ . Repeat for later  $T_n$ , minimizing the distance between  $\mathcal{F}(\hat{\boldsymbol{\alpha}}_{n-1}^{k-1})$  and one of the samples  $\boldsymbol{\alpha}_{n,1}^{k-1}, \dots, \boldsymbol{\alpha}_{n,M}^{k-1}$ . This process must be run sequentially and relies on the modification to  $\mathcal{P}$  made in section 2—that solutions are not altered once converged. Referring again to Figure 3, the corresponding coarse trajectories of these optimally chosen samples  $\hat{\boldsymbol{\alpha}}_n^{k-1}$  must also be calculated to carry out the predictor-corrector step (lines 29–32).

At this point, the set of initial values  $\{\hat{\boldsymbol{\alpha}}_2^{k-1}, \dots, \hat{\boldsymbol{\alpha}}_{N-1}^{k-1}\}$  has been selected from the ensemble of random samples, effectively replacing the previously found  $\mathbf{U}_n^{k-1}$ . The




FIG. 3. An illustration of the sampling and propagation process within  $\mathcal{P}_s$  following iteration  $k = 1$ . The fine solution is given in black, the  $k = 0$  fine solutions in blue, the  $k = 1$  coarse solutions in red, and the  $k = 1$  predictor-corrector solutions as red dots. With  $M = 5$ , four samples  $\alpha_{n,m}^1$  (green dots) are taken at  $T_2$  and  $T_3$  from distributions with means  $U_2^1$  and  $U_3^1$ , and some finite standard deviations, respectively. These values, along with  $U_2^1$  and  $U_3^1$  themselves, are propagated in parallel forward in time using  $\mathcal{F}$  (green lines). The optimally chosen samples  $\hat{\alpha}_n^1$  (refer to the text for how these are chosen) are then propagated forward in time using  $\mathcal{G}$  (yellow lines).

coarse and fine propagations of these values are now used in the predictor-corrector (lines 33–37) such that

$$(3.3) \quad \mathbf{U}_n^k = \begin{cases} \mathcal{G}(\mathbf{U}_{n-1}^k) + \mathcal{F}(\mathbf{U}_{n-1}^{k-1}) - \mathcal{G}(\mathbf{U}_{n-1}^{k-1}) & \text{for } n = 2, \\ \mathcal{G}(\mathbf{U}_{n-1}^k) + \mathcal{F}(\hat{\alpha}_{n-1}^{k-1}) - \mathcal{G}(\hat{\alpha}_{n-1}^{k-1}) & \text{for } n = 3, \dots, N. \end{cases}$$

Using the same stopping criteria (2.6) from  $\mathcal{P}$  (lines 38–42), the algorithm either stops or runs another stochastic parareal iteration.

As a final remark, instead of minimizing the distance between  $\mathcal{F}(\hat{\alpha}_{n-1}^{k-1})$  and one of the samples  $\alpha_{n,1}^{k-1}, \dots, \alpha_{n,M}^{k-1}$ , one could think about doing some sort of interpolation to choose a more optimal point than the  $M$  samples. In this setting, however, this is not possible because we require not just the exact starting condition, which would be the optimally chosen sample, but also its value having been propagated by  $\mathcal{F}$  (which we only have for the  $M$  samples).

**3.2. Sampling rules.** The probability distributions  $\Phi_n^{k-1}$  incorporate different combinations of available information about the initial values at different temporal resolutions, i.e., the coarse, fine, and predictor-corrector initial values  $\mathcal{G}(\mathbf{U}_{n-1}^{k-1})$ ,  $\mathcal{G}(\mathbf{U}_{n-1}^{k-2})$ ,  $\mathcal{F}(\mathbf{U}_{n-1}^{k-2})$ , and  $\mathbf{U}_{n-1}^{k-1}$ , respectively. This information is used to define the marginal means and standard deviations in the “sampling rules” outlined in the following subsections. Note how the distributions vary with time  $T_n$  and  $k$ , so that the accuracy of the distributions increase as the initial values in  $\mathcal{P}_s$  are iteratively improved. Using Gaussian and copula distributions, we analyze the performance of different sampling rules within  $\mathcal{P}_s$  in section 4. This will give us a more comprehensive understanding on whether the choice of distribution family  $\Phi_n^{k-1}$  or the parameters  $\mu_n^{k-1}$ ,  $\sigma_n^{k-1}$ , and  $\mathbf{R}_n^{k-1}$  have the greatest impact on the convergence rate  $k_s$ . In all tested cases, we observe that taking correlated samples significantly improves the performance of  $\mathcal{P}_s$  compared to the independent setting.

**3.2.1. Multivariate Gaussian.** First, we consider perturbing the initial values using stochastic “noise,” i.e., considering errors compared to the true initial values to be normally distributed, a standard method for modeling uncertainty. The initial values  $\alpha_{n,m}^{k-1}$  are sampled from a multivariate Gaussian distribution  $\mathcal{N}(\mu_n^{k-1}, \Sigma_n^{k-1})$ , where  $(\Sigma_n^{k-1})_{i,j} = \rho_{n_i,j}^{k-1} \sigma_{n_i}^{k-1} \sigma_{n_j}^{k-1}$  is the  $d \times d$  positive semidefinite covariance matrix. As marginal means  $\mu_n^{k-1}$  we choose either the fine resolution values  $\mathcal{F}(\mathbf{U}_{n-1}^{k-2})$  (prior to correction) or the predictor-corrector values  $\mathbf{U}_n^{k-1}$ . For the marginal standard deviations, we choose  $\sigma_n^{k-1} = |\mathcal{G}(\mathbf{U}_{n-1}^{k-1}) - \mathcal{G}(\mathbf{U}_{n-1}^{k-2})|$  as they are of the order of the corrections made by the predictor-corrector and each marginal decreases toward zero as the algorithm converges (as expected). For the correlation coefficients,  $\rho_{n_i,j}^{k-1}$ , we calculate the linear correlation between the  $\mathcal{F}$  propagated samples using Pearson’s method—recall (3.1). Note that  $|\cdot|$  denotes the componentwise absolute value. Testing revealed that alternative marginal standard deviations  $|\mathbf{U}_n^{k-1} - \mathbf{U}_n^{k-2}|$  and  $|\mathcal{F}(\mathbf{U}_{n-1}^{k-2}) - \mathcal{G}(\mathbf{U}_{n-1}^{k-2})|$  did not span sufficiently large distances around  $\mu_n^{k-1}$  in order for sampling to be efficient, i.e., they required much higher sampling to perform as well as  $\sigma_n^{k-1} = |\mathcal{G}(\mathbf{U}_{n-1}^{k-1}) - \mathcal{G}(\mathbf{U}_{n-1}^{k-2})|$  (results not shown). The samples  $\alpha_{n,m}^{k-1} \sim \mathcal{N}(\mu_n^{k-1}, \Sigma_n^{k-1})$  are taken according to the following sampling rules:

**Rule 1:**  $\mu_n^{k-1} = \mathcal{F}(\mathbf{U}_{n-1}^{k-2})$  and  $\sigma_n^{k-1} = |\mathcal{G}(\mathbf{U}_{n-1}^{k-1}) - \mathcal{G}(\mathbf{U}_{n-1}^{k-2})|$ .

**Rule 2:**  $\mu_n^{k-1} = \mathbf{U}_n^{k-1}$  and  $\sigma_n^{k-1} = |\mathcal{G}(\mathbf{U}_{n-1}^{k-1}) - \mathcal{G}(\mathbf{U}_{n-1}^{k-2})|$ .

Note that a linear combination of the rules or taking half the samples from each appears to work well, with performance similar to the individual rules themselves (results not shown).

**3.2.2. Multivariate copula.** Alternatively, we consider errors that may have a different (possibly non-Gaussian) dependency structure. We consider another multivariate distribution, known as a copula, with uniformly distributed marginals scaled such that they have the same value as the marginal means and standard deviations in subsection 3.2.1.

A copula  $\mathcal{C} : [0, 1]^d \rightarrow [0, 1]$  is a joint cumulative distribution function, of a  $d$ -dimensional random vector, with uniform marginal distributions [22]. Sklar’s theorem [30] states that any multivariate cumulative distribution function with continuous marginal distributions can be written in terms of  $d$  uniform marginal distributions and a copula that describes the correlation structure between them. While there are numerous families of copulas, we consider the symmetric  $t$ -copula,  $\mathcal{C}^t$ , the copula underlying the multivariate  $t$ -distribution, which depends on the parameter  $\nu$  and the correlation matrix  $\mathbf{R}_n^{k-1}$ . We fix  $\nu = 1$  so that samples have a higher probability of being drawn toward the edges of the  $[0, 1]^d$  hypercube; see [22]. Note that  $\nu \rightarrow \infty$  can be thought of as sampling from the Gaussian copula.

Correlated samples  $\chi = (\chi_1, \dots, \chi_d)^T \sim \mathcal{C}^t$  that are generated in  $[0, 1]^d$  then need to be rescaled such that each marginal is uniformly distributed in an interval  $[x_i, y_i] \subset \mathbb{R}$ , with mean  $\mu_i$  and standard deviation  $\sigma_i$  for  $i \in \{1, \dots, d\}$ . By definition, a marginal uniform distribution on  $[x_i, y_i]$  has mean  $(x_i + y_i)/2$  and variance  $(y_i - x_i)^2/12$  which we set to  $\mu_i$  and  $\sigma_i^2$ , respectively. Solving these equations, we find that the desired marginals are uniform distributions on  $[\mu_i - \sqrt{3}\sigma_i, \mu_i + \sqrt{3}\sigma_i]$ . Thus, setting  $2\sqrt{3}\sigma_i\chi_i + \mu_i - \sqrt{3}\sigma_i$  guarantees that the generated samples  $\chi \sim \mathcal{C}^t$  have the same marginal means  $\mu_i$  and standard deviations  $\sigma_i$  as the Gaussian distributions. This allows us to compare performance results in section 4. The  $t$ -copula sampling rules (**Rule 3** and **Rule 4**) are therefore defined componentwise as

$\alpha_{n,m}^{k-1}(i) = 2\sqrt{3}\sigma_i\chi_i + \mu_i - \sqrt{3}\sigma_i$ , for  $i \in \{1, \dots, d\}$ , with  $\chi \sim \mathcal{C}^t$  and parameters  $\mu_i$  and  $\sigma_i$  chosen to be identical to Rule 1 and Rule 2, respectively.

Note that this copula construction can be extended to other families of copulas (e.g., the Gaussian copula) and to copulas with different marginal distributions (e.g., Gaussian,  $t$ , or logistic marginals). Therefore, sampling from the Gaussian multivariate distributions in subsection 3.2.1 can be considered a special case of these more general copulas.

**3.3. Convergence.** Independent simulations of  $\mathcal{P}_s$  will return a numerical solution  $U_n^k$  which, along with  $k_s$ , vary stochastically—since the optimal initial values chosen will vary between simulations. Given the randomness in these quantities, we can discuss the convergence of  $\mathcal{P}_s$  in two ways.

First, consider convergence in terms of minimizing the random variable  $k_s$  by studying  $\mathbb{P}(k_s < k_d)$ . Given there are no analytical results for  $\mathcal{P}$  guaranteeing that  $k_d < N$  for any given problem, proving that  $\mathbb{P}(k_s < k_d) = 1$ , or at least  $\mathbb{E}(k_s) = \sum_{k=1}^N k\mathbb{P}(k_s = k) < k_d$ , seems to be analytically intractable. However, we can qualitatively discuss  $\mathbb{P}(k_s < k_d)$  and  $\mathbb{E}(k_s)$  with respect to the number of samples  $M$ . Consider the following cases:

- $M = 1$ .

Running  $\mathcal{P}_s$  is equivalent to running  $\mathcal{P}$ , hence the convergence of  $\mathcal{P}_s$  follows from that of  $\mathcal{P}$  and therefore  $\mathbb{E}(k_s) = k_s = k_d$ .

- $1 < M < \infty$ .

For finitely many samples, we estimate the discrete probability distributions  $\mathbb{P}(k_s = k)$  and observe, in all numerical experiments (see section 4), that  $\mathbb{P}(k_s < k_d) \rightarrow 1$  for  $M \approx 10$ . Moreover, we observe that  $\mathbb{P}(k_s < k_d)$  increases and  $\mathbb{E}(k_s)$  decreases for increasing  $M$ —with  $\mathbb{E}(k_s) < k_d$  for all values of  $M$  tested.

- $M \rightarrow \infty$ .

If we were able to take infinitely many samples,  $\mathcal{P}_s$  effectively samples every possible value in the support of  $\Phi$ , i.e., every initial value that has a nonzero probability of being sampled from  $\Phi$ . Therefore, if  $\Phi$  has infinite support, e.g., the Gaussian distribution, all possible initial values in  $\mathbb{R}^d$  are sampled and propagated, hence the fine solution will be recovered almost surely in  $\mathbb{E}(k_s) = k_s = 2$  iterations. Note this is the smallest value  $k_s$  can take to converge assuming convergence does not occur following the first iteration—although the algorithm could be modified such that convergence occurs almost surely in  $k_s = 1$  iterations as  $M \rightarrow \infty$ . In subsection 4.1, we illustrate this property numerically by taking a large number of samples for a single realization of  $\mathcal{P}_s$ .

In the scenario that  $\mathcal{P}_s$  converges in  $k_s = N$  iterations (irrespective of the value of  $M$ ), it will return the fine solution just as  $\mathcal{P}$  does when  $k_d = N$  (having propagated the exact initial value at  $T_0$  sequentially  $N$  times using  $\mathcal{F}$ ).

Second, consider convergence in terms of the stochastic solution  $U_n^k$  (3.3) approaching (in the mean-square sense) the fine solution  $U_n$  as  $k$  increases. In the numerical experiments in section 4, we did not observe a case where  $\mathcal{P}_s$  fails to converge to the fine solution. In fact, we observe tight confidence intervals on the numerical errors between  $U_n^k$  and  $U_n$  upon multiple realizations of  $\mathcal{P}_s$  (see Figures 7 and 9). We may expect  $\mathcal{P}_s$  to fail to converge (i.e., solutions blow up) in cases in which  $\mathcal{P}$  also fails—typically this means that a more accurate coarse solver is required for both algorithms. It should be noted, however, that because  $\mathcal{P}_s$  samples  $M$  initial values,




FIG. 4. Illustration of a possible processor configuration if two subintervals were to converge between iterations two and three of  $\mathcal{P}_s$ . The letter  $C$  denotes a converged subinterval, the number 1 denotes a subinterval where we propagate the converged initial value from the preceding subinterval using  $\mathcal{F}$ , and the letter  $M$  denotes the number of samples taken in an unconverged subinterval.

only  $M'$  out of the  $M$  propagated solutions may blow up in each time subinterval, due to the poor coarse solver, and so  $\mathcal{P}_s$  could in fact locate a solution in cases where  $\mathcal{P}$  cannot—although this is not tested here.

**3.4. Computational complexity.** At each iteration,  $\mathcal{P}_s$  runs the fine solver more frequently than  $\mathcal{P}$ , albeit still in parallel, and therefore requires a larger number of processors. The first iteration of  $\mathcal{P}_s$  requires  $N$  processors; however, once sampling begins in  $k \geq 2$ , it requires at most  $M(N - I - 1) + 1$  processors – assuming  $I$  subintervals converge during  $k = 1$ . While we assume processors are in abundance, this number scales directly with  $M$  and so it is important to keep  $M$  to a minimum if limited processing power is available in practice.

As the stochastic iterations progress, the number of processors required, i.e.,  $M(N - I - 1) + 1$ , decreases as the number of converged subintervals,  $I$ , increases—leaving a growing number of processors idle. Each additional subinterval that converges leaves  $M$  idle processors that we can reassign to do additional sampling and propagation. We assign each set of  $M$  idle processors to the earliest unconverged subinterval with the least number of samples, ensuring all processors are working at all times to explore the solution space for the true initial values (see Figure 4 for an illustration). We do not explicitly write the pseudocode for reassigning the idle processors in Algorithm 2 (lines 10–22) to avoid additional complexity—the process is, however, implemented in the numerical experiments in section 4.<sup>2</sup>

In terms of timings, an iteration of  $\mathcal{P}_s$  takes approximately the same wallclock time as one of  $\mathcal{P}$ . In this regard, we assume that the extra serial costs in  $\mathcal{P}_s$ , e.g., correlation estimation, selecting optimal samples, and the extra  $\mathcal{G}$  runs, take negligible wallclock time when compared to a single run of  $\mathcal{F}$ . Therefore, the wallclock time for  $\mathcal{P}_s$  will be lower than for  $\mathcal{P}$  if  $k_s < k_d$ . This comes at a cost of requiring  $\mathcal{O}(MN)$  processors rather than  $N$  to solve the problem. However, it should be highlighted that if  $\mathcal{P}_s$  converges in even one less iteration than  $\mathcal{P}$ , we avoid an extra run of  $\mathcal{F}$  which may save a large amount of wallclock time.

**4. Numerical results.** In this section, we compare the numerical performance of  $\mathcal{P}$  and  $\mathcal{P}_s$  on systems of one, two, and three ODEs of varying complexity.<sup>3</sup> Both algorithms use RK4 methods to carry out integration, with  $\mathcal{G}$  using time step  $\delta T$  and  $\mathcal{F}$  using time step  $\delta t$ , the latter at least 75 times smaller than the former. We quantify the performance of  $\mathcal{P}_s$  by estimating the distributions of  $k_s$  for each sampling rule and

<sup>2</sup>To increase efficiency further we also attempted to store previously sampled and propagated fine trajectories to use in future iterations of the algorithm; however, they did not improve performance. This was because only the most recent samples were ever chosen in each iteration (results not shown).

<sup>3</sup>All algorithms were coded in MATLAB and simulations were run using HPC facilities at the University of Warwick. Samples code for both  $\mathcal{P}$  and  $\mathcal{P}_s$  can be found in the public repository at <https://github.com/kpentland/StochasticParareal>.

by measuring the accuracy of the stochastic solutions against those obtained serially with  $\mathcal{F}$ . Since a limited number of processors were available for these experiments, the results are based not on calculating wallclock runtimes but on comparing the convergence rates  $k_d$  and  $k_s$ —which are independent of the number of processors used. Additional results for these test cases, as well as two further test problems, are given in the supplementary materials (M141423SupMat.pdf [local/web 1.21MB]).

**4.1. Scalar nonlinear equation.** First, we consider the nonlinear ODE

$$(4.1) \quad \frac{du_1}{dt} = \sin(u_1) \cos(u_1) - 2u_1 + e^{-t/100} \sin(5t) + \ln(1+t) \cos(t),$$

with initial value  $u_1(0) = 1$  [6]. Discretize the time interval  $t \in [0, 100]$  using  $N = 40$  subintervals, coarse time steps  $\delta T = 100/80$ , and fine time steps  $\delta t = 100/8000$ . Numerical solutions to (4.1) are shown on the interval  $[0, 18]$  in Figure 5(a). Deterministically,  $\mathcal{P}$  locates a solution in  $k_d = 25$  iterations using error tolerance  $\varepsilon = 10^{-10}$ .  $\mathcal{P}_s$  converges in a varying number of iterations  $k_s$  with  $\mathbb{P}(k_s < k_d) = 1$ —see Figure 5(b) for convergence of 10 independent simulations using  $M = 3$ . From this plot, we see that by taking just three samples,  $\mathcal{P}_s$  reduces the number of iterations by almost a factor of two—from 25 to approximately 14 (on average).

When  $M$  is increased above one,  $\mathcal{P}_s$  begins generating stochastic solutions that converge in a varying number of iterations  $k_s$ . In order to accurately compare  $k_d$  with the discrete random variable  $k_s$ , we run 2000 independent simulations of  $\mathcal{P}_s$  to estimate the distribution of  $k_s$  for a given  $M$ . Upon estimating these distributions, it was found that  $\mathbb{P}(k_s < 25) = 1$  for each of the four sampling rules ( $\forall M > 1$ ), meaning that by doubling the number of processors we can beat parareal with probability one. The estimated distributions of  $k_s$ , using sampling rule 1 (the other rules perform similarly), as a function of  $M$  are given in Figure 6(a). The stacked bars represent the estimated discrete probability of a simulation converging in a given number of iterations. The results show  $\mathcal{P}_s$  converging in just five iterations in the best case—demonstrating  $\mathcal{P}_s$  has the potential to yield significant parallel speedup, given sufficiently many samples are drawn. Figure 6(b) emphasizes the power of the stochastic method, showing that




FIG. 5. (a) Numerical solution of (4.1) over  $[0, 18]$  using  $\mathcal{F}$  serially and a single realization of  $\mathcal{P}_s$ . Note that only a subset of the fine time steps of the  $\mathcal{P}_s$  solution are shown for clarity. (b) Errors at successive iterations of  $\mathcal{P}$  (black line) and 10 independent realizations of  $\mathcal{P}_s$  (blue lines). Horizontal dashed red line represents the stopping tolerance  $\varepsilon = 10^{-10}$ . Note that both panels use  $\mathcal{P}_s$  with sampling rule 1 and  $M = 3$ .

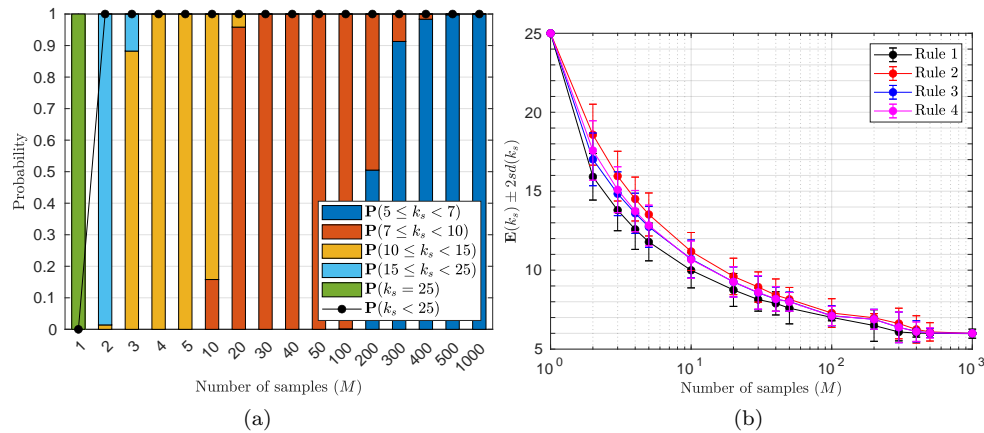


FIG. 6. (a) Estimated discrete distributions of  $k_s$  as a function of  $M$  for sampling rule 1. (b) Estimated expectation of  $k_s$  as a function of  $M$ , calculated using estimated distributions of  $k_s$  for each sampling rule with error bars representing  $\pm$  two standard deviations  $sd(k_s)$ . Distributions in both panels are estimated by simulating 2000 independent realisations of  $\mathcal{P}_s$  for each  $M$ .




FIG. 7. Errors of  $\mathcal{P}$  (red) and mean  $\mathcal{P}_s$  (black) solutions against the serial  $\mathcal{F}$  solution over time. The mean error is obtained by running 2000 independent realizations of  $\mathcal{P}_s$  with sampling rule 2 and  $M = 4$ —the confidence interval representing the mean  $\pm$  two standard deviations is shown in light blue.

the estimated expected value  $\mathbb{E}(k_s)$  decreases as  $M$  increases, with the estimated standard deviation  $sd(k_s) = \sqrt{\sum_{k=1}^N (k - \mathbb{E}(k_s))^2 \mathbb{P}(k_s = k)}$  decreasing too. The improved performance of  $\mathcal{P}_s$  as  $M$  increases reflects what was discussed in subsection 3.3. In particular, we ran a single realization of  $\mathcal{P}_s$  with  $M = 10^6$ , observing that  $\mathcal{P}_s$  converged in four iterations (result not shown), confirming that  $k_s$  continues to decrease for increasing  $M$ . By looking at Figure 6(b), we see that sampling rule 1 yields the lowest expected values of  $k_s$  for small values of  $M$ , with all sampling rules performing similarly for large  $M$ .

To verify accuracy of the stochastic solutions, we plot the difference between the mean of 2000 independent realizations of  $\mathcal{P}_s$  and the serially calculated  $\mathcal{F}$  solution in Figure 7. Also shown is the confidence interval given by two standard deviations of the stochastic solutions (which is at most  $\mathcal{O}(10^{-11})$ ) and the error generated by

$\mathcal{P}$ . Accuracy is maintained with respect to the fine solution across time, even more so than the  $\mathcal{P}$  solution. See SM1 for numerical results of  $\mathcal{P}_s$  applied to a *stiff* scalar nonlinear ODE. In this case, the stiffness of the equation demands a higher value of  $M$  to improve  $k_s$ —something we also observe for the Brusselator in subsection 4.2.

**4.2. The Brusselator system.** Next, consider the Brusselator system

$$(4.2a) \quad \frac{du_1}{dt} = A + u_1^2 u_2 - (B + 1)u_1,$$

$$(4.2b) \quad \frac{du_2}{dt} = B u_1 - u_1^2 u_2,$$

a pair of stiff nonlinear ODEs that model an autocatalytic chemical reaction [15]. Using parameters  $(A, B) = (1, 3)$ , trajectories of the system exhibit oscillatory behaviour in phase space, approaching a limit cycle (as  $t \rightarrow \infty$ ) that contains the unstable fixed point  $(1, 3)^T$ . Now that  $d > 1$ , we use bivariate distributions to sample the initial values—meaning we can compare the effects of including or excluding the correlations between variables. System (4.2) is solved using initial values  $\mathbf{u}(0) = (1, 3.07)^T$  over time interval  $t \in [0, 15.3]$  with  $N = 25$ ,  $\delta T = 15.3/25$ , and  $\delta t = 15.3/2500$  [32]. The numerical solution to (4.2) in phase space and convergence of the successive errors are reported in SM2. With these parameters and a tolerance of  $\varepsilon = 10^{-6}$ ,  $\mathcal{P}$  takes  $k_d = 7$  iterations to stop and return a numerical solution.

The estimated distributions of  $k_s$  for sampling rule 1 are given in Figure 8(a). Even though  $\mathcal{P}$  takes just  $k_d = 7$  iterations to stop, we observe that  $\mathcal{P}_s$  can still reach the desired tolerance in 5 or 6 iterations—albeit requiring larger values of  $M$ . We believe this is due to the stiffness of the system and poor accuracy of the explicit  $\mathcal{G}$  solver—results presented for the stiff ODE in SM1 appear to confirm this. Using adaptive time-stepping methods could be a way to reduce the value of  $M$  needed to converge in fewer  $k_s$  [19]. The solid lines in Figure 8(b) show that, using sampling rules 1 or 3,  $\mathcal{P}_s$  only requires  $M \approx 10$  to beat parareal almost certainly, i.e., to guarantee that  $\mathbb{P}(k_s < 7) \rightarrow 1$ . Sampling rules 1 and 3 outperform 2 and 4 in this particular system. Note, however, the stark decrease in performance if instead uncorrelated

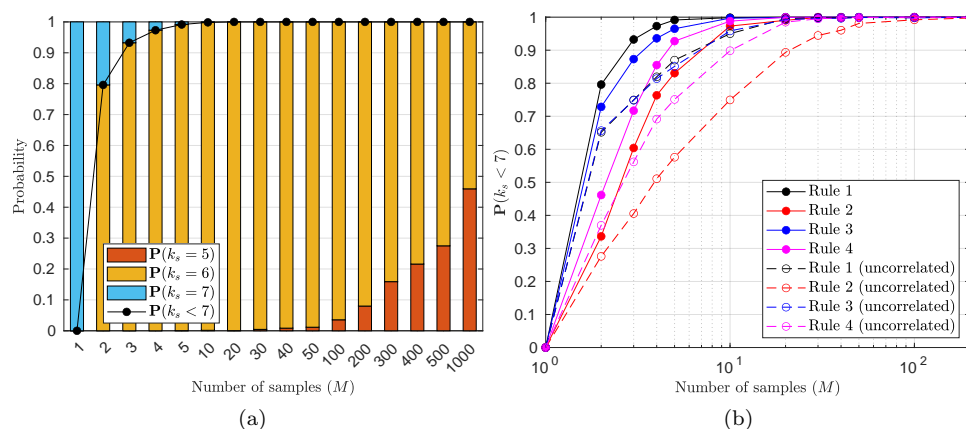


FIG. 8. (a) Estimated discrete probabilities of  $k_s$  as a function of  $M$  for sampling rule 1. (b) Estimated probability that the convergence rate  $k_s$  is smaller than  $k_d = 7$  as a function of  $M$  for the sampling rules with (solid lines) and without (dashed lines) correlations. Distributions were estimated by simulating 2000 independent realisations of  $\mathcal{P}_s$  for each  $M$ .






FIG. 9. Errors of  $\mathcal{P}$  (red) and mean  $\mathcal{P}_s$  (black) solutions against the  $\mathcal{F}$  solution. The mean error is obtained by running 2000 independent realisations of  $\mathcal{P}_s$  with sampling rule 4 with  $M = 200$ —the confidence interval representing the mean  $\pm$  two standard deviations is shown in light blue. Panel (a) displays errors for the  $u_1$  component of the solution while (b) displays the  $u_2$  component.

samples are generated within  $\mathcal{P}_s$  (dashed lines). This demonstrates the importance of accounting for the dependence between variables in nonlinear systems such as (4.2). Observe again, in Figure 9, how the mean  $\mathcal{P}_s$  solutions attain accuracy equivalent to, or better than, the  $\mathcal{P}$  solutions, with standard deviations at most  $\mathcal{O}(10^{-6})$ .

Further results of  $\mathcal{P}_s$  applied to (4.2) and an additional two-dimensional nonlinear system are presented in SM2 and SM3, respectively. Observe again that for the less stiff system in SM3, we require less sampling to improve  $k_s$  compared to the Brusselator—further highlighting the demand for higher  $M$  in stiff systems to improve the convergence rate.

**4.3. The Lorenz system.** Finally, we consider the Lorenz system

$$(4.3a) \quad \frac{du_1}{dt} = \gamma_1(u_2 - u_1),$$

$$(4.3b) \quad \frac{du_2}{dt} = \gamma_2 u_1 - u_1 u_3 - u_2,$$

$$(4.3c) \quad \frac{du_3}{dt} = u_1 u_2 - \gamma_3 u_3,$$

a simplified model for weather prediction [18]. With the parameters  $(\gamma_1, \gamma_2, \gamma_3) = (10, 28, 8/3)$ , (4.3) exhibits chaotic behavior where trajectories with initial values close to one another diverge exponentially. This will test the robustness of  $\mathcal{P}_s$ , as small numerical differences between initial values will mean that errors can grow rapidly as time progresses. We solve (4.3) using initial values  $\mathbf{u}(0) = (-15, -15, 20)^T$  over the interval  $[0, 18]$ , discretized using  $N = 50$  subintervals and time steps  $\delta T = 18/250$  and  $\delta t = 18/18750$ . With a tolerance of  $\varepsilon = 10^{-8}$ ,  $\mathcal{P}$  takes  $k_d = 20$  iterations to converge.

Running  $\mathcal{P}_s$  to compare the performance of the sampling rules, we see again in Figure 10(a) that taking correlated samples is much more efficient than not and that only  $M \approx 10$  samples are required to beat parareal with probability one. For the chaotic trajectories generated by (4.3), sampling close to the predictor-corrector, rules 2 and 4, yields superior performance compared to rules 1 and 3 for small values

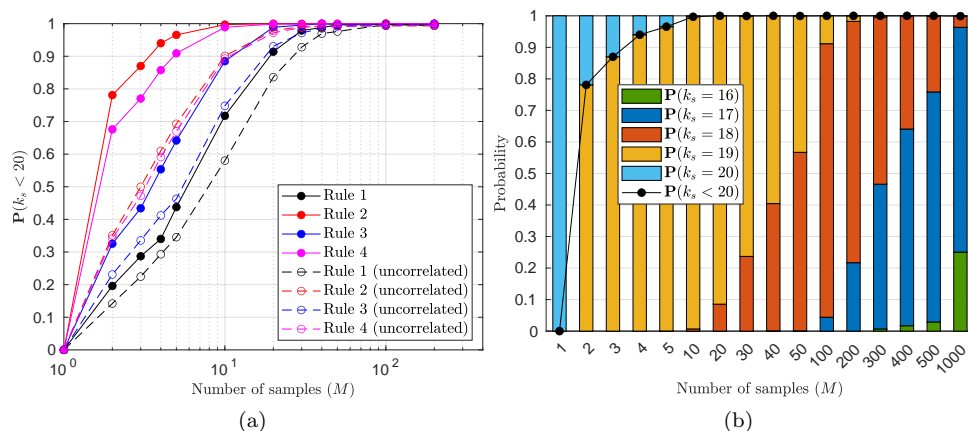


FIG. 10. (a) Estimated probability that the convergence rate  $k_s$  is smaller than  $k_d = 20$  as a function of  $M$  for the sampling rules with (solid lines) and without (dashed lines) correlations. Distributions were estimated by simulating 2000 independent realizations of  $\mathcal{P}_s$  for each  $M$ . (b) Estimated discrete probabilities of  $k_s$  as a function of  $M$  for sampling rule 2.

of  $M$ . Figure 10(b) displays estimated distributions for varying  $M$  using sampling rule 2—yielding a best convergence rate  $k_s = 16$  for approximately 25% of runs with  $M = 1000$ . These results demonstrate the robustness of  $\mathcal{P}_s$  and that the sampling and propagation process is not impeded by the exponential divergence of trajectories. Additional results of  $\mathcal{P}_s$  applied to (4.3) are given in SM4.

**5. Conclusions.** In this paper, we have extended the parareal algorithm using probabilistic methods to develop a stochastic parareal algorithm for solving systems of ODEs in a time-parallel manner. Instead of passing deterministically calculated initial values into parareal’s predictor-corrector, stochastic parareal selects more accurate values from a randomly sampled set, at each temporal subinterval, to converge in fewer iterations. In section 4, we compared performance against the deterministic parareal algorithm on several low-dimensional ODE systems of increasing complexity by calculating the estimated convergence rate distributions (upon multiple independent realization of stochastic parareal) with increasing numbers of random samples  $M$ . By taking just  $M \approx 10$  (correlated) samples, the estimated probability of converging sooner than parareal approached one in all test cases. Similarly, we observed numerical convergence toward the fine (exact) solution with accuracy of similar order to parareal and, in the spirit of probabilistic numerics [14, 24], obtained a measure of uncertainty over the ODE solution upon multiple realizations of the algorithm.

The probability that stochastic parareal converges faster than standard parareal depends on a number of factors: the complexity of the problem being solved, the number of time subintervals ( $N$ ), the accuracy of the coarse integrator ( $\mathcal{G}$ ), the number of random samples ( $M$ ), and the type of sampling rule in use. Sampling rules 1 and 3 (sampling close to the fine solutions) outperformed rules 2 and 4 (sampling close to the predictor-corrector solutions) for the ODE systems in subsection 4.1, subsection 4.2, and SM1. The reverse was true, however, for the systems in subsection 4.3 and SM3, making it difficult to determine an optimal rule for a general ODE system. To overcome having to choose a particular sampling rule, one could linearly combine different rules or even sample from multiple rules simultaneously. We would suggest sampling from probability distributions with infinite support, i.e., the Gaussians (rules

1 and 2), so that samples can be taken anywhere in  $\mathbb{R}^d$  with nonzero probability. Having finite support may have created difficulty for the uniform marginal  $t$ -copulas (rules 3 and 4) because samples could only be taken in a finite hyperrectangle in  $\mathbb{R}^d$ —problematic if the exact solution were to lay outside of this space.

When solving stiff ODEs (see subsection 4.2 and SM1), results indicated that stochastic parareal demanded increasingly high sampling to converge sooner than parareal than for nonstiff systems. For example, we observe that when taking  $M = 100$  samples in the nonstiff scalar ODE in Figure 6, the expected convergence rate decreases from 25 to 7, whereas for the stiff scalar ODE in SM1, the rate only drops from 8 to 6. A similar observation can be made in the two-dimensional test cases in subsection 4.2 (stiff) and SM3 (nonstiff). These results exemplify the role that system complexity, e.g., stiffness or chaos, plays in the performance of both algorithms. In SM1, stochastic parareal was also shown to perform more efficiently for problems that deterministic parareal itself struggles with, i.e., cases in which the accuracy of the coarse integrator  $\mathcal{G}$  is poor. In SM3 it was also observed that, for low sample numbers ( $M = 2$ ), stochastic parareal actually converged in one more iteration than parareal in less than 2.5% of cases. This suggests there may be a minimum number of samples required to beat parareal in some situations—something to be investigated with further experimentation.

In summary, we have demonstrated that probabilistic methods and additional processors can, for low-dimensional ODEs, be used to increase the parallel scalability of existing time-parallel algorithms. Next we need to investigate possible improvements and generalizations, for example, determining whether the algorithm scales for larger systems of equations—essential if it is to be used for solving PDE problems. Moreover, we need to determine whether we can design adaptive sampling rules that do not need to be specified a priori to simulation. Finally, we would aim to make use of the whole ensemble of fine propagated trajectories rather than using only one—avoiding the waste of valuable information about the solution at the coarse and fine resolutions. An approach in this direction has recently been proposed [26], making use of ideas from the field of probabilistic numerics to adopt a more Bayesian approach to this problem.

**Acknowledgments.** The authors would like to acknowledge the University of Warwick Scientific Computing Research Technology Platform for assistance in the research described in this paper and the anonymous reviewers for their insightful comments on improving this manuscript.

#### REFERENCES

- [1] L. BAFFICO, S. BERNARD, Y. MADAY, G. TURINICI, AND G. ZÉRAH, *Parallel-in-time molecular-dynamics simulations*, Phys. Rev. E, 66 (2002), 057701, <https://doi.org/10.1103/PhysRevE.66.057701>.
- [2] G. BAL, *On the convergence and the stability of the parareal algorithm to solve partial differential equations*, in Domain Decomposition Methods in Science and Engineering, Lect. Notes Comput. Sci. Eng. 40, Springer, New York, 2005, pp. 425–432, [https://doi.org/10.1007/3-540-26825-1\\_43](https://doi.org/10.1007/3-540-26825-1_43).
- [3] G. BAL AND Y. MADAY, *A “Parareal” Time Discretization for non-linear PDE’s with application to the pricing of an American put*, in Recent Developments in Domain Decomposition Methods, Springer, Berlin, 2002, pp. 189–202, [https://doi.org/10.1007/978-3-642-56118-4\\_12](https://doi.org/10.1007/978-3-642-56118-4_12).
- [4] A. BELLEN AND M. ZENNARO, *Parallel algorithms for initial-value problems for difference and differential equations*, J. Comput. Appl. Math., 25 (1989), pp. 341–350, [https://doi.org/10.1016/0377-0427\(89\)90037-X](https://doi.org/10.1016/0377-0427(89)90037-X).

- [5] K. BURRAGE, *Parallel and Sequential Methods for Ordinary Differential Equations*, Clarendon Press, Burlington, MA, 1995.
- [6] P. CHARTIER AND B. PHILIPPE, *A parallel shooting technique for solving dissipative ODE's*, *Computing*, 51 (1993), pp. 209–236, <https://doi.org/10.1007/BF02238534>.
- [7] W. R. ELWASIF, S. S. FOLEY, D. E. BERNHOLDT, L. A. BERRY, D. SAMADDAR, D. E. NEWMAN, AND R. SANCHEZ, *A dependency-driven formulation of parareal: Parallel-in-time solution of PDEs as a many-task application*, in *Proceedings of the 2011 ACM International Workshop on Many Task Computing on Grids and Supercomputers*, New York, 2011, ACM Press, pp. 15–24, <https://doi.org/10.1145/2132876.2132883>.
- [8] S. ENGBLOM, *Parallel in time simulation of multiscale stochastic chemical kinetics*, *Multiscale Model. Simul.*, 8 (2009), pp. 46–68, <https://doi.org/10.1137/080733723>.
- [9] P. F. FISCHER, F. HECHT, AND Y. MADAY, *A parareal in time semi-implicit approximation of the Navier-Stokes equations*, in *Domain Decomposition Methods in Science and Engineering*, *Lect. Notes Comput. Sci. Eng.* 40, Springer, New York, 2005, pp. 433–440, <https://doi.org/10.1007/3-540-26825-1.44>.
- [10] M. J. GANDER, *50 years of time parallel time integration*, in *Multiple Shooting and Time Domain Decomposition Methods*, Springer, New York, 2015, pp. 69–113, <https://doi.org/10.1007/978-3-319-23321-5.3>.
- [11] M. J. GANDER AND E. HAIRER, *Nonlinear convergence analysis for the parareal algorithm*, in *Domain Decomposition Methods in Science and Engineering*, *Lect. Notes Comput. Sci. Eng.* 60, Springer, New York, 2008, pp. 45–56, <https://doi.org/10.1007/978-3-540-75199-1.4>.
- [12] M. J. GANDER AND S. VANDEWALLE, *Analysis of the parareal time-parallel time-integration method*, *SIAM J. Sci. Comput.*, 29 (2007), pp. 556–578, <https://doi.org/10.1137/05064607X>.
- [13] I. GARRIDO, M. S. ESPEDAL, AND G. E. FLADMARK, *A convergent algorithm for time parallelization applied to reservoir simulation*, in *Domain Decomposition Methods in Science and Engineering*, *Lect. Notes Comput. Sci. Eng.* 40, Springer, New York, 2005, pp. 469–476, <https://doi.org/10.1007/3-540-26825-1.48>.
- [14] P. HENNIG, M. A. OSBORNE, AND M. GIROLAMI, *Probabilistic numerics and uncertainty in computations*, *Proc. A* 471 (2015), 20150142, <https://doi.org/10.1098/rspa.2015.0142>.
- [15] R. LEFEVER AND G. NICOLIS, *Chemical instabilities and sustained oscillations*, *J. Theoret. Biol.*, 30 (1971), pp. 267–284, [https://doi.org/10.1016/0022-5193\(71\)90054-3](https://doi.org/10.1016/0022-5193(71)90054-3).
- [16] F. LEGOLL, T. LEIÈVRE, K. MYERSCOUGH, AND G. SAMAËY, *Parareal computation of stochastic differential equations with time-scale separation: A numerical convergence study*, *Comput. Vis. Sci.*, 23 (2020), pp. 1–18, <https://doi.org/10.1007/s00791-020-00329-y>.
- [17] J. L. LIONS, Y. MADAY, AND G. TURINICI, *Résolution d'EDP par un schéma en temps pararéel*, *C. R. Math. Acad. Sci. Paris Ser. I Math.*, 332 (2001), pp. 661–668, [https://doi.org/10.1016/S0764-4442\(00\)01793-6](https://doi.org/10.1016/S0764-4442(00)01793-6).
- [18] E. N. LORENZ, *Deterministic Nonperiodic Flow*, *J. Atmos. Sci.*, 20 (1963), pp. 130–141, [https://doi.org/10.1175/1520-0469\(1963\)020<0130:dnf>2.0.co;2](https://doi.org/10.1175/1520-0469(1963)020<0130:dnf>2.0.co;2).
- [19] Y. MADAY AND O. MULA, *An adaptive parareal algorithm*, *J. Comput. Appl. Math.*, 377 (2020), 112915, <https://doi.org/10.1016/j.cam.2020.112915>.
- [20] Y. MADAY AND G. TURINICI, *A parareal in time procedure for the control of partial differential equations*, *C. R. Math. Acad. Sci. Paris*, 335 (2002), pp. 387–392, [https://doi.org/10.1016/S1631-073X\(02\)02467-6](https://doi.org/10.1016/S1631-073X(02)02467-6).
- [21] Y. MADAY AND G. TURINICI, *The parareal in time iterative solver: A further direction to parallel implementation*, in *Domain Decomposition Methods in Science and Engineering*, *Lect. Notes Comput. Sci. Eng.* 40, Springer, New York, 2005, pp. 441–448, <https://doi.org/10.1007/3-540-26825-1.45>.
- [22] R. B. NELSEN, *An Introduction to Copulas*, Springer, New York, 2006.
- [23] J. NIEVERGELT, *Parallel methods for integrating ordinary differential equations*, *Commun. ACM*, 7 (1964), pp. 731–733, <https://doi.org/10.1145/355588.365137>.
- [24] C. J. OATES AND T. J. SULLIVAN, *A modern retrospective on probabilistic numerics*, *Statist. Comput.*, 29 (2019), pp. 1335–1351, <https://doi.org/10.1007/s11222-019-09902-z>.
- [25] B. W. ONG AND J. B. SCHRODER, *Applications of time parallelization*, *Comput. Vis. Sci.*, 23 (2020), pp. 1–11, <https://doi.org/10.1007/s00791-020-00331-4>.
- [26] K. PENTLAND, M. TAMBORRINO, T. J. SULLIVAN, J. BUCHANAN, AND L. C. APPEL, *GParareal: A Time-Parallel ODE Solver Using Gaussian Process Emulation*, arXiv:2201.13418, 2022.
- [27] P. SAHA, J. STADEL, AND S. TREMAINE, *A parallel integration method for solar system dynamics*, *Astronom. J.*, 114 (1997), pp. 409–415, <https://doi.org/10.1086/118485>.
- [28] D. SAMADDAR, D. P. COSTER, X. BONNIN, L. A. BERRY, W. R. ELWASIF, AND D. B. BATCHLOR, *Application of the parareal algorithm to simulations of ELMs in ITER plasma*, *Comput. Phys. Commun.*, 235 (2019), pp. 246–257, <https://doi.org/10.1016/j.cpc.2018.08.007>.

- [29] D. SAMADDAR, D. E. NEWMAN, AND R. SÁNCHEZ, *Parallelization in time of numerical simulations of fully-developed plasma turbulence using the parareal algorithm*, J. Comput. Phys., 229 (2010), pp. 6558–6573, <https://doi.org/10.1016/j.jcp.2010.05.012>.
- [30] A. SKLAR, *Fonctions de Répartition à n Dimensions et Leurs Marges*, Publications de L'Institut de Statistique de L'Université de Paris, 8 (1959), pp. 229–231.
- [31] A. TOSELLI AND O. WIDLUND, *Domain Decomposition Methods—Algorithms and Theory*, Springer, New York, 2005.
- [32] L. TREFETHEN, A. BIRKISSON, AND T. DRISCOLL, *Exploring ODEs*, SIAM, Philadelphia, 2017.
- [33] J. M. TRINDADE AND J. C. PEREIRA, *Parallel-in-time simulation of two-dimensional, unsteady, incompressible laminar flows*, Numerical Heat Transfer Part B Fundamentals, 50 (2006), pp. 25–40, <https://doi.org/10.1080/10407790500459379>.