SINGLE-ENSEMBLE MULTILEVEL MONTE CARLO FOR DISCRETE INTERACTING-PARTICLE METHODS[*]

ARNE BOUILLON[†], TOON INGELAERE[†], AND GIOVANNI SAMAEY[†]

Abstract. To solve problems in domains such as filtering, optimization, and posterior sampling, *interacting-particle methods* have recently received much attention. These parallelizable and often gradient-free algorithms use an *ensemble* of particles that evolve in time, based on a combination of well-chosen dynamics and interaction between the particles. For computationally expensive dynamics – for example, dynamics that solve inverse problems with an expensive forward model – the cost of attaining a high accuracy quickly becomes prohibitive. We exploit a hierarchy of approximations to this forward model and apply multilevel Monte Carlo (MLMC) techniques, improving the asymptotic cost-to-error relation. More specifically, we use MLMC *at each time step* to estimate the interaction term within a single, globally-coupled ensemble. This technique was proposed by Hoel et al. in the context of the ensemble Kalman filter; the goal of the present paper is to study its applicability to a general framework of interacting-particle methods. After extending the algorithm and its analysis to a broad set of methods with fixed numbers of time steps, we motivate the application of the method to the class of algorithms with an infinite time horizon, which includes popular methods such as ensemble Kalman algorithms for optimization and sampling. Numerical tests confirm the improved asymptotic scaling of the multilevel approach.

Key words. Monte Carlo, multilevel, interacting-particle methods, Bayesian inverse problems, ensemble Kalman inversion, ensemble Kalman sampler, ensemble Kalman filter

MSC codes. 65C05, 65C35, 65Y20, 62F15

1. Introduction. This paper studies interacting-particle methods (IPMs), algorithms that solve various problems with an *ensemble* (or collection) of *particles* (or points, evolving over time in, e.g., a state or parameter space). Three problem classes for which IPMs have been particularly successful are filtering [1, 2, 3, 10, 15], optimization [28, 34, 38], and Bayesian-posterior sampling [5, 17, 27]. Simulated with a finite ensemble size J, IPMs can be viewed as Monte Carlo approximations to a corresponding mean-field model. Some example methods are introduced in subsection 1.1.

Our main contribution is comparing this straightforward approximation to a newly proposed multilevel Monte Carlo (MLMC) scheme. MLMC is discussed in subsection 1.2, after which we can state our goal and discuss related literature in subsection 1.3. Subsection 1.4 summarizes the layout of the paper.

1.1. Interacting-particle methods. This subsection discusses the use of IPMs for filtering, optimization in Bayesian inverse problems, and posterior sampling. The first of these domains, filtering, is concerned with reconstructing state variables from noisy observations. Consider the discrete dynamics

$$u_{n+1} = \mathcal{G}_n(u_n), \qquad 0 \le n < N, \tag{1.1}$$

[†]NUMA research group, Department of Computer Science, KU Leuven, Leuven, Belgium ({arne.bouillon,toon.ingelaere,giovanni.samaey}@kuleuven.be).

where $u_n \in \mathbb{R}^{d_u}$ denotes the state at time step n and $\mathcal{G}_n$ is the potentially stochastic *forward model* at that same time step. Now assume access to noisy observations

$$(1.2) \qquad y_n = H u_n + \eta_n, \qquad 0 < n \leq N,$$

with H a linear state-to-observation map and η_n a noise term. The filtering problem is to estimate the states $\{u_n\}_{n=1}^N$ based on these noisy observations. One popular filtering algorithm is the ensemble Kalman filter (EnKF).

Example 1.1 (Ensemble Kalman filter). The EnKF [15] is an IPM whose J particles $\{u_n^j\}_{j=1}^J$ at time n estimate the expectation and uncertainty on u_n. It assumes that $\eta_n \sim \mathcal{N}(0, \Gamma)$ for some positive definite Γ. An ensemble member u_n^j evolves as

$$(1.3) \qquad u_{n+1}^j = (I - K_n H)\mathcal{G}_n(u_n^j) + K_n(y_{n+1} + \sqrt{\Gamma}\eta_n^j),$$

where $K_n = C_n^{gg} H^\top (H C_n^{gg} H^\top + \Gamma)^{-1}$ (with C_n^{gg} the sample covariance of $\{\mathcal{G}_n(u_n^j)\}_j$) is called the *Kalman gain*, and where $\eta_n^j \sim \mathcal{N}(0, I)$. Theoretical studies of the EnKF and its properties often consider its *mean-field limit* (the limit for an infinite number of particles); convergence to that limit for finite J is quantified in [31, 32].

A second problem class is that of *Bayesian inverse problems*. Here, we assume to have a single, deterministic forward map $\mathcal{G} : \mathbb{R}^{d_u} \to \mathbb{R}^{d_g}$ and an unknown parameter $u \in \mathbb{R}^{d_u}$. We are given a prior distribution $\pi_{\mathrm{prior}}(u)$ on u and an observation

$$(1.4) \qquad y = \mathcal{G}(u) + \eta,$$

in which $\eta \in \mathbb{R}^{d_g}$ follows a known noise model π_η, such that we can define the *likelihood* $\pi_{\mathrm{li}}(y \mid u) := \pi_\eta(y - \mathcal{G}(u))$. Bayes' formula results in the posterior distribution

$$(1.5) \qquad \pi_{\mathrm{post}}(u \mid y) = \frac{\pi_{\mathrm{li}}(y \mid u)\pi_{\mathrm{prior}}(u)}{\pi_{\mathrm{ev}}(y)}.$$

The *evidence* $\pi_{\mathrm{ev}}(y)$ in (1.5) is usually intractable to compute, as it involves integration over the entire parameter domain. Hence, we work under the assumption that only an unnormalized posterior $\pi_{\mathrm{post}}(u \mid y) \propto \pi_{\mathrm{li}}(y \mid u)\pi_{\mathrm{prior}}(u)$ is available.

The posterior distribution is used in two main ways. Optimization methods can target the *maximum a posteriori* (MAP) parameter, the most likely u value given y and π_{prior}. A more complete view of the posterior and its features is provided by sampling methods, which construct a sequence of u values distributed according to $\pi_{\mathrm{post}}(u \mid y)$. IPMs that perform these tasks include ensemble Kalman inversion (EKI) [28] and ensemble Kalman sampling (EKS) [17], both inspired by the EnKF.

Example 1.2 (Ensemble Kalman inversion). We again assume that $\eta \sim \mathcal{N}(0, \Gamma)$ and now require a uniform prior on u. (General noise distributions are handled in [13]; prior regularization is discussed in, e.g., [8, 26].) EKI was proposed as the iterated application of the EnKF (creating an artificial discrete time dimension) in [28], to which time steps τ_n were added in [38]. With $\eta_n^j \sim \mathcal{N}(0, I)$, the resulting dynamics are

$$(1.6) \quad u_{n+1}^j = u_n^j + \tau_n C_n^{ug}(\tau_n C_n^{gg} + \Gamma)^{-1}(y - \mathcal{G}(u_n^j) + \sqrt{\Gamma/\tau_n}\,\eta_n^j), \qquad 1 \leq j \leq J,$$

where C_n^{ug} is the sample covariance of $\{u_n^j\}_j$ and $\{\mathcal{G}(u_n^j)\}_j$. A continuous-time limit was studied in [38] and re-discretized in a slightly different form in, e.g., [30] – in that work, the (artificial) time steps τ_n are also determined adaptively adaptively. For simplicity, we omit this adaptivity and fix all τ_n in advance. Increasing J improves exploration capabilities and reveals connections to gradient descent [4].

Example 1.3 (Ensemble Kalman sampling). Assume again that $\eta \sim \mathcal{N}(0, \Gamma)$ and consider $\pi_{\mathrm{prior}}(u)$ a zero-centered Gaussian with covariance Γ_0. EKS was proposed and motivated in continuous-time form in [17]. In practice, a discretization should be used instead, and a common one is the linearly implicit split-step scheme (see [17])

$$\text{(1.7a)} \qquad u_{n+1}^{*,j} = u_n^j + \tau_n C_n^{ug} \Gamma^{-1}(y - \mathcal{G}(u_n^j)) - \tau_n C_n^{uu} \Gamma_0^{-1} u_{n+1}^{*,j},$$

$$\text{(1.7b)} \qquad u_{n+1}^j = u_{n+1}^{*,j} + \sqrt{2\tau_n C_n^{uu}}\, \eta_n^j,$$

where C_n^{uu} is the sample covariance of $\{u_n^j\}_j$ and where $\eta_n^j \sim \mathcal{N}(0, I)$. These dynamics estimate the posterior (1.5), based on linear and Gaussian ansatzes, as $n \to \infty$. Like many ensemble Kalman methods, EKS is motivated by its mean-field limit [17], and a larger J implies better performance [17] and convergence to the mean-field limit [12].

Two classes of IPMs. While the above division in terms of problem classes is natural from an application point of view, the example IPMs reveal another distinction, relevant to algorithm development and analysis throughout the paper. We distinguish two classes of discrete-time IPMs based on their time dimension:

- *IPMs with a fixed number of discrete time steps.* The ensemble Kalman filter falls into this category, as do other filtering algorithms such as the deterministic ensemble Kalman filter (DEnKF) [37].
- *IPMs with an infinite number of discrete time steps.* The EKS and EKI algorithms use discrete time-stepping, but their target is the particle distribution for the time step $n \to \infty$. Another example in this class are the consensus-based sampling (CBS) algorithm and its variants [5].

While many of the IPMs in this paper discretize continuous-time equations, refinement to obtain lower errors typically uses more particles over more time steps, while keeping the time step length (e.g., [6, 17, 34]). This justifies classifying them as discrete-time methods. Other IPMs, such as the ensemble Kalman–Bucy filter [2] or EKI's sampling variant (see [11]), are defined in continuous time over a fixed interval – lowering their error involves refinement in time. Those IPMs are outside the scope of this paper.

Advantages of IPMs. Many of these methods require no derivatives of the forward model; instead of gradient information, interaction between the ensemble members drives the particle evolution. This is crucial when gradients are expensive, unavailable, or undefined due to a non-differentiable objective [30], or when they are noisy or highly oscillatory [14]. In addition, these methods allow for straightforward parallelization, as only the interaction term requires information from multiple particles.

Error sources for finite ensembles. When simulating an IPM with a finite number of particles J, the interaction term depends on the stochastic positions of all particles. Hence, simulating IPMs with a finite J should be interpreted as an approximation to a theoretical mean-field model. In fact, when the end goal of IPM simulation is the final particle distribution, three sources of errors compound:

1. The interaction effect at each time is computed based on a finite number of particles. This introduces a bias and a statistical error in the dynamics.
2. The finite ensemble size introduces a statistical error on the final distribution.
3. A bias on the final particle distribution is introduced in infinite-time IPMs, as we must truncate the time dimension in practical simulations.

It is plausible that error sources 1 and 2 decrease when the number of particles J increases. For many methods, it turns out that this happens at the $\mathcal{O}(J^{-1/2})$ rate typical for Monte Carlo [11, 12]. When particle evolution is computationally expensive due to a complex forward model in the underlying problem, attaining a high accuracy through straightforward Monte Carlo simulation of the IPM can become prohibitive.

1.2. Multilevel Monte Carlo. To simulate interacting-particle methods with expensive forward models more quickly, this paper will use multilevel Monte Carlo (MLMC) [18]. The core MLMC idea is as follows. An expectation $\mathbb{E}[X^L]$ of a computationally expensive random variable X^L, to which a hierarchy of cheaper and less accurate approximations $\{X^\ell\}_{\ell=0}^{L-1}$ is available, is rewritten with a telescoping sum:

$$(1.8) \qquad \mathbb{E}[X^L] = \mathbb{E}[X^0] + \sum\nolimits_{\ell=1}^{L} \mathbb{E}[X^\ell - X^{\ell-1}].$$

MLMC uses this representation to approximate $\mathbb{E}[X^L]$. Many cheap realizations of X^0 are sampled, giving an accurate estimate of $\mathbb{E}[X^0]$. The difference terms are then each approximated by sampling *correlated* realizations of X^ℓ and $X^{\ell-1}$. This correlation reduces the variance of the difference estimator, such that fewer samples are needed. The challenge in designing MLMC algorithms is typically to find a way to correlate these realizations, and to analyze the correlation mathematically.

1.3. Related work and objectives. The idea of using MLMC to speed up simulation of IPMs is not new: multiple of these methods have been extended with multilevel variants. Research on multilevel interacting-particle methods (MLIPMs) has focused on algorithms for filtering. A multilevel EnKF was proposed in [23] and extended to spatio-temporal processes in [9]. An alternative to this approach, is developed in [24]; an adaptation to a multi-index scheme was made in [25].

Definition 1.4 (Single- and multiple-ensemble MLIPMs). We will refer to the algorithms in [9, 23] as *single-ensemble* MLIPMs, as they employ one single ensemble of pairwise-coupled particles – each assigned a level, with higher levels having fewer particle pairs – that interact globally. The methods in [24, 25] use *multiple independent ensembles on each level*. Higher levels contain ensembles that use more expensive forward models and more particles per ensemble, but they have far fewer ensembles than the lower levels. The fundamental difference between these paradigms is which quantity is chosen as the target for MLMC. In the former, that target is the EnKF covariance matrix at each time step. Hence all particles interact at each time step as well. The multiple-ensemble case, however, performs a multilevel estimation of particle distributions as a last step, after evolving all ensembles independently.

A multilevel EnKF algorithm has also been proposed that is suited for reservoir history matching in the non-asymptotic regime [16]. Another variant uses reduced-order models [35]. In addition, multiple-ensemble MLMC particle filters have been described [20, 29]. The work [41] does not use MLMC, but is related to our work as it selects an appropriate accuracy level for each time step in EKI (Example 1.2).

In the continuous-time context, a multilevel ensemble Kalman–Bucy filter [7] and many multilevel schemes for general McKean–Vlasov equations (SDEs with interacting particles) [22, 36, 39] are found in the literature. Of these, [36] comes closest to the single-ensemble approach, but uses less coupling between levels and requires the interaction term to be the expectation $\mathbb{E}[f(X)]$ of a function f of the particles X.

In [24], the multiple-ensemble MLEnKF is compared to the single-ensemble one from [23] for some test problems. This shows the latter approach consistently outperforming the former by a constant factor. Nevertheless, single-ensemble MLIPMs remain restricted to the EnKF. The goal of this paper, then, is twofold.

1. Formulate a broad framework for interacting-particle methods and devise a single-ensemble multilevel simulation algorithm.
2. Analyze the rate at which single- and multilevel simulation algorithms converge to the mean-field model when more particles are added.

1.4. Overview of the paper. The remainder of the paper is laid out as follows. After section 2 introduces our notational conventions, we formulate a general framework for applying single-ensemble MLMC to IPMs with a fixed number of discrete time steps in section 3. This completes our first objective. Then, for our second objective, section 4 states our main theorems on the asymptotic cost-to-error relation of this technique when increasing the number of particles; they are then proven in sections 5 and 6. Section 7 discusses how our algorithms extend to IPMs that adapt the time horizon to the desired accuracy. The performance of our multilevel interacting-particle methods is studied numerically in section 8, after which section 9 concludes the paper. Various proofs are relegated to Appendix A.

2. Notation and prerequisites. We denote the absolute value of a scalar u, the 2-norm of a vector or matrix u, and the sum of the component norms of a tuple u, by $|u|$. We will also make heavy use of the p-norm of a random variable U:

$$(2.1) \qquad \|U\|_p := \mathbb{E}[|U|^p]^{1/p} \quad \text{with } p \geq 2.$$

For a random variable (RV) $U \sim \delta_u$ that only takes one value u, it is straightforward to see that $\|U\|_p = |u|$ for any $p \geq 2$. The following properties will also prove useful.

Property 2.1 (Generalized Hölder's inequality). For any RVs (U, V) and $p \geq 2$,

$$(2.2) \qquad \|UV\|_p \leq \|U\|_q \|V\|_r$$

whenever $1/p = 1/q + 1/r$. In particular,

$$(2.3) \qquad \|UV\|_p \leq \|U\|_{2p} \|V\|_{2p}.$$

Property 2.2 (Norm ordering). For any RV U and any $p \geq 2$,

$$(2.4) \qquad |\mathbb{E}[U]| \leq \mathbb{E}[|U|] \leq \|U\|_p.$$

Property 2.3 (Monotonicity of the p-norm). For any RVs (U, V) and $p \geq 2$,

$$(2.5) \qquad (\forall (u, v) \sim (U, V) \colon |u| \leq |v|) \Rightarrow \|U\|_p \leq \|V\|_p.$$

Property 2.4 (Marcinkiewicz–Zygmund inequality). Let X and $X_1, \ldots, X_J$ be zero-mean i.i.d. RVs such that $\|X\|_p < \infty$ for all $p \geq 2$. Then, for any $p \geq 2$, there exists a constant c_p such that

$$(2.6) \qquad \left\| \frac{1}{J} \sum_{j=1}^{J} X_j \right\|_p \leq c_p J^{-1/2} \|X\|_p.$$

Proof. This result is easily derived from, e.g., [21, Corollary 8.2]. $\qquad \square$

We write $A \succ 0$ (or $A \succeq 0$) to indicate that a matrix A is positive (semi-)definite. The expressions $A \succ B$ and $A \succeq B$ mean $A - B \succ 0$ and $A - B \succeq 0$, respectively. For two positive expressions $f(x)$ and $g(x)$ in a variable x, the notation $f(x) \lesssim_x g(x)$ will denote that there exists a constant c such that $f(x) \leq cg(x)$ for all x. We further write $f(x) \eqsim_x g(x)$ to mean $f(x) \lesssim_x g(x) \lesssim_x f(x)$. Often, the variable x in question is clear from context and we simply write $\lesssim$ or $\eqsim$. For a set Ω and for any $d \in \mathbb{N}$ and $p \geq 2$, we introduce the space

$$(2.7) \qquad L^p(\Omega, \mathbb{R}^d) := \{U \colon \Omega \to \mathbb{R}^d \mid \|U\|_p < \infty\}.$$

We will also use the shorthand notation $L^{\geq 2}(\Omega, \mathbb{R}^d) := \bigcap_{p \geq 2} L^p(\Omega, \mathbb{R}^d)$.

3. Presentation of the core framework. We now present our framework for IPMs. For many practical problems, the potentially stochastic[1] forward model $\mathcal{G}$ is computationally intractable. Instead, approximations

$$(3.1) \qquad \mathcal{G}^\ell(\cdot) \approx \mathcal{G}(\cdot) \in \mathbb{R}^{d_g}, \qquad \ell \geq 0$$

are available, with a higher ℓ indicating a more accurate but more expensive model. We will work in this context. First, subsection 3.1 identifies a common structure to the methods introduced so far. Starting from this unified notation, subsection 3.2 reformulates practical IPMs as approximations to underlying mean-field dynamics. Subsection 3.3 then proposes a multilevel simulation algorithm that approximates the mean-field model by using the hierarchy $\{\mathcal{G}^\ell\}_\ell$.

3.1. Single-level simulation algorithm. The dynamics in Examples 1.1 to 1.3 display many similarities. More broadly, we will study interacting-particle methods where the particles $u_{n+1}^{L,j} \in \mathbb{R}^{d_u}$ at time $n+1$ are defined as

$$(3.2) \qquad u_{n+1}^{L,j} = \Psi_n(u_n^{L,j}, \mathcal{G}^L(u_n^{L,j}), \widehat{\Theta}_J(\boldsymbol{x}_n^L), \eta_n^j), \qquad 1 \leq j \leq J.$$

Here, $L \geq 0$ selects an accuracy level from the hierarchy (3.1). The variables $\eta_n^j \sim \mathcal{N}(0, I)$ are d_η-dimensional random noise, and $\widehat{\Theta}_J(\cdot)$ is a *sample statistic* of the ensemble that constitutes the interaction term between the particles. This may be a scalar, a vector, a matrix, or a tuple $\widehat{\Theta}_J(\cdot) = (\widehat{\Theta}_J^{(1)}(\cdot), \ldots, \widehat{\Theta}_J^{(M)})$ of multiple other sample statistics. For example, $\widehat{\Theta}_J(\cdot)$ might be the mean $E_J(\cdot)$ or a tuple $(E_J(\cdot), C_J(\cdot))$ of the mean and sample covariance. Throughout the paper, a variable using the letter x denotes the concatenation of the corresponding u (particle position) and its image under $\mathcal{G}$, or under $\mathcal{G}^L$ if the particle has a level L. Bold variables denote ensembles. Hence, the argument to $\widehat{\Theta}_J$ contains all particles and their $\mathcal{G}^L$-images:

$$(3.3) \qquad \boldsymbol{x}_n^L := \left\{ \begin{bmatrix} u_n^{L,j} \\ \mathcal{G}^L(u_n^{L,j}) \end{bmatrix} \right\}_{j=1}^J.$$

We also define $d_x := d_u + d_g$.

The functions Ψ_n in (3.2), together with the choice of sample statistic $\widehat{\Theta}_J$ and the distribution from which the initial positions $u_0^{L,j}$ are sampled, define the interacting-particle method. In Appendix A.2, we show that the IPMs from Examples 1.1 to 1.3 fit into the framework (3.2).

In general, the goal of simulating (3.2) is not the positions of the individual particles, but some quantity of interest (QoI) of the entire ensemble. This quantity is another sample statistic, $\widehat{\Theta}_J^{\mathrm{f}}(\cdot)$, of the ensemble at some time index N:

$$(3.4) \qquad q_N^L := \widehat{\Theta}_J^{\mathrm{f}}(\boldsymbol{x}_N^L), \quad \text{a realization of the random variable} \quad Q_N^L.$$

3.2. Mean-field interacting-particle methods. Simulating (3.2) with finite J and L can be seen as a finite-cost approximation of J independent particles following a mean-field equation. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space, such that the single-level particles' initial positions are sampled from the known distribution $U_0 \in L^{\geq 2}(\Omega, \mathbb{R}^{d_u})$. Then, a mean-field particle taking N time steps is a realization of the correlated random variables $\{\bar{u}_n \colon \Omega \to \mathbb{R}^{d_u}\}_{n=0}^N$, with $u_0 = \bar{u}_0$. At further time steps, the particle evolves as

$$(3.5) \qquad \bar{u}_{n+1}(\omega) = \Psi_n(\bar{u}_n(\omega), \mathcal{G}(\bar{u}_n(\omega), \omega), \Theta[\bar{X}_n], \eta_n(\omega)),$$

---

[1]Except when relevant, the random input to $\mathcal{G}$ will be omitted in notation.

where $\Theta[\cdot]$ is a *statistical parameter* corresponding to the finite-sample statistic $\widehat{\Theta}_J(\cdot)$. Capital letters denote random variables and the variable $\bar{X}_n$ follows our established convention, where the letter x indicates the concatenation of the corresponding u and its $\mathcal{G}$-image. For ease of notation, we will write $\bar{u}_n^j := \bar{u}_n(\omega^j)$ and $\eta_n^j := \eta_n(\omega^j)$ with $\omega^j \in \Omega$, and omit $\mathcal{G}$'s second argument. Instead of (3.5), then, we write

$$(3.6) \qquad \bar{u}_{n+1}^j = \Psi_n(\bar{u}_n^j, \mathcal{G}(\bar{u}_n^j), \Theta[\bar{X}_n], \eta_n^j).$$

The similarity to (3.2) is apparent. The difference lies in the use of the exact forward model $\mathcal{G}$ and the statistical parameter $\Theta[\cdot]$ of the entire particle distribution. Due to this latter difference, the particles $\bar{u}_{n+1}^j$ no longer interact; they depend on their own law instead. The mean-field equivalent to the QoI computed by the finite-sample statistic (3.4) is a statistical parameter $\Theta^f[\cdot] \in \mathbb{R}^{d_f}$ of the particle distribution:

$$(3.7) \qquad \bar{q}_N := \Theta^f[\bar{X}_N].$$

3.3. Single-ensemble multilevel simulation algorithm. Subsection 3.1 described the typical, single-level Monte Carlo approximation of the unattainable mean-field dynamics and quantity of interest (3.6)–(3.7): one fixes an accuracy level L and a number of particles J, simulates (3.2), and performs the estimate (3.4). In contrast, our single-ensemble multilevel Monte Carlo approach mixes particles on different levels, generalizing the MLEnKF from [23]. We select a maximum level L and a number J_ℓ of particles or particle pairs for each level $0 \le \ell \le L$. The ensemble then consists of base-level particles and higher-level correction pairs:

$$(3.8) \qquad \boldsymbol{u}_n^{\mathrm{ML}} = (\{u_n^{0,\mathrm{F},j}\}_{j=1}^{J_0}, \{(u_n^{\ell,\mathrm{F},j}, u_n^{\ell,\mathrm{C},j})\}_{j=1,\ell=1}^{J_\ell, L}).$$

We define multilevel x variables similarly to before, now using $\mathcal{G}^\ell$ for fine particles on level ℓ and $\mathcal{G}^{\ell-1}$ for coarse ones. The (F, C) pairs are correlated by setting $u_0^{\ell,\mathrm{F},j} = u_0^{\ell,\mathrm{C},j}$ (for $\ell \ge 1$) and evolving (with the same random input to both $\mathcal{G}$ invocations)

$$(3.9a) \qquad u_{n+1}^{\ell,\mathrm{F},j} = \Psi_n(u_n^{\ell,\mathrm{F},j}, \mathcal{G}^\ell(u_n^{\ell,\mathrm{F},j}), \widehat{\Theta}^{\mathrm{ML}}(\boldsymbol{x}_n^{\mathrm{ML}}), \eta_n^{\ell,j}), \qquad \ell \in \{0,1,\ldots,L\},$$

$$(3.9b) \qquad u_{n+1}^{\ell,\mathrm{C},j} = \Psi_n(u_n^{\ell,\mathrm{C},j}, \mathcal{G}^{\ell-1}(u_n^{\ell,\mathrm{C},j}), \widehat{\Theta}^{\mathrm{ML}}(\boldsymbol{x}_n^{\mathrm{ML}}), \eta_n^{\ell,j}), \qquad \ell \in \{1,\ldots,L\},$$

where $\eta_n^{\ell,j} \sim \mathcal{N}(0, I)$ and, analogously to (1.8), the multilevel sample statistic $\widehat{\Theta}^{\mathrm{ML}}$ is

$$(3.10) \qquad \widehat{\Theta}^{\mathrm{ML}}(\boldsymbol{x}_n^{\mathrm{ML}}) := \widehat{\Theta}_{J_0}(\{x_n^{0,\mathrm{F},j}\}_j) + \sum_{\ell=1}^{L} \left(\widehat{\Theta}_{J_\ell}(\{x_n^{\ell,\mathrm{F},j}\}_j) - \widehat{\Theta}_{J_\ell}(\{x_n^{\ell,\mathrm{C},j}\}_j) \right).$$

After evolving these dynamics, a multilevel estimation of $\bar{q}_N$ is performed as

$$(3.11) \qquad q_N^{\mathrm{ML}} := \widehat{\Theta}^{f,\mathrm{ML}}(\boldsymbol{x}_N^{\mathrm{ML}}), \quad \text{a realization of the random variable} \quad Q_N^{\mathrm{ML}}.$$

Remark 3.1 (Alternative multilevel estimator). As an alternative to (3.9)–(3.10), consider using a different interaction term for each level ℓ', independent of finer levels:

$$(3.12) \qquad \widehat{\Theta}^{\mathrm{ML},\ell'}(\boldsymbol{x}_n^{\mathrm{ML}}) := \widehat{\Theta}_{J_0}(\{x_n^{0,\mathrm{F},j}\}_j) + \sum_{\ell=1}^{\ell'} \left(\widehat{\Theta}_{J_\ell}(\{x_n^{\ell,\mathrm{F},j}\}_j) - \widehat{\Theta}_{J_\ell}(\{x_n^{\ell,\mathrm{C},j}\}_j) \right).$$

The dynamics of $u_n^{\ell,\mathrm{F},j}$ and $u_n^{\ell,\mathrm{C},j}$ can then use $\widehat{\Theta}^{\mathrm{ML},\ell}$ and $\widehat{\Theta}^{\mathrm{ML},\ell-1}$, respectively. Notice that (3.12) is an estimator of the single-level interaction term on level ℓ', rather than on level L. In our tests, this variation on our multilevel scheme performs comparably to (3.9)–(3.10). Hence, we will focus on (3.9)–(3.10), which is a more direct generalization of the methods in [9, 23].

4. Theoretical properties and convergence. In subsection 4.1, we will formulate assumptions on the ingredients of the framework outlined in section 3. Under these assumptions, subsections 4.2 and 4.3 will prove convergence rates to the mean-field model for our single- and multilevel simulation algorithms, respectively.

4.1. Assumptions. To prove convergence of the single- and multilevel algorithms from subsections 3.1 and 3.3, we need some assumptions on the ingredients of IPMs in the framework: (i) the approximations $\mathcal{G}^\ell$ to the exact forward model $\mathcal{G}$, (ii) the functions Ψ_n defined in (3.2), and (iii) the parameter Θ and its estimator $\widehat{\Theta}$. These conditions are *local*, and hence contain *locality conditions* such as $\|U_1 - U_0\|_r \leq d$.

Assumption 4.1. There exist constants ρ, β, and γ such that, for any $p \geq 2$ and $U_0 \in L^{\geq 2}(\Omega, \mathbb{R}^{d_u})$, there exist constants $d, c_{g,\{1,2,3\}} > 0$ and $r \geq 2$ such that for any $U_{\{1,2\}} \in L^{\geq 2}(\Omega, \mathbb{R}^{d_u})$ with $\|U_{\{1,2\}} - U_0\|_r \leq d$, the following hold for $\ell \geq 0$.
  (i) The forward models $\mathcal{G}^\ell$ satisfy a Lipschitz bound:

$$(4.1a) \qquad \|\mathcal{G}^\ell(U_1) - \mathcal{G}^\ell(U_2)\|_p \leq c_{g,1}\|U_1 - U_2\|_r.$$

  (ii) All $\mathcal{G}^\ell$ are bounded:

$$(4.1b) \qquad \|\mathcal{G}^\ell(U_1)\|_p \leq c_{g,2}.$$

  (iii) The rate at which $\mathcal{G}^\ell$ approximates $\mathcal{G}$ is described by β:

$$(4.1c) \qquad \|\mathcal{G}^\ell(U_1) - \mathcal{G}(U_1)\|_p \leq c_{g,3}\rho^{-\beta\ell/2}.$$

  Convergence as $\ell \to \infty$ justifies the notation $\mathcal{G}^\infty := \mathcal{G}$.
  (iv) The rate at which $\mathcal{G}^\ell$ increases in cost is described by γ:

$$(4.1d) \qquad \mathrm{Cost}(\mathcal{G}^\ell) \lesssim_\ell \rho^{\gamma\ell}.$$

Assumption 4.2. For any $X_0 \in L^{\geq 2}(\Omega, \mathbb{R}^{d_x})$, $\Theta_0 \in L^{\geq 2}(\Omega, \mathbb{R}^{d_\theta})$, and $p \geq 2$, and with $\eta \sim \mathcal{N}(0, I)$, there exist constants $d, c_\psi > 0$ and $r \geq 2$ such that, for any $X_{\{1,2\}} \in L^{\geq 2}(\Omega, \mathbb{R}^{d_x})$ and $\Theta_2 \in L^{\geq 2}(\Omega, \mathbb{R}^{d_\theta})$, and for any fixed θ_1, the following property holds if $\|X_{\{1,2\}} - X_0\|_r \leq d$, $\|\theta_1 - \Theta_0\|_r \leq d$, and $\|\Theta_2 - \Theta_0\|_r \leq d$.
  (i) All functions Ψ_n satisfy a local Lipschitz bound:

$$(4.2)$$
$$\|\Psi_n(U_1, G_1, \theta_1, \eta) - \Psi_n(U_2, G_2, \Theta_2, \eta)\|_p \leq c_\psi(\|X_1 - X_2\|_r + \|\theta_1 - \Theta_2\|_r).$$

Assumption 4.3. For any $X_0 \in L^{\geq 2}(\Omega, \mathbb{R}^{d_x})$ and $p \geq 2$, there exist constants $r \geq 2$ and $d, c_{\theta,\{1,2,3,4\}} > 0$ such that, for any $X_{\{1,2\}} \in L^{\geq 2}(\Omega, \mathbb{R}^{d_x})$ with $\|X_{\{1,2\}} - X_0\|_r \leq d$, the following hold (with $\boldsymbol{X_i}$ an ensemble, each distributed as X_i).
  (i) The summary statistic $\widehat{\Theta}$ satisfies a local Lipschitz bound:

$$(4.3a) \qquad \|\widehat{\Theta}_J(\boldsymbol{X_1}) - \widehat{\Theta}_J(\boldsymbol{X_2})\|_p \leq c_{\theta,1}\|X_1 - X_2\|_r.$$

  (ii) With $\boldsymbol{X_i}$ i.i.d., a difference in Θ is efficiently estimated by a difference in $\widehat{\Theta}$:

$$(4.3b)$$
$$\|(\widehat{\Theta}_J(\boldsymbol{X_1}) - \widehat{\Theta}_J(\boldsymbol{X_2})) - (\Theta[X_1] - \Theta[X_2])\|_p \leq c_{\theta,2}J^{-1/2}\|X_1 - X_2\|_r.$$

  (iii) With $\boldsymbol{X_i}$ i.i.d., the parameter Θ is efficiently estimated by $\widehat{\Theta}$:

$$(4.3c) \qquad \|\widehat{\Theta}_J(\boldsymbol{X_1}) - \Theta[X_1]\|_p \leq c_{\theta,3}J^{-1/2}.$$

  (iv) The parameter Θ and the statistic $\widehat{\Theta}$ are bounded:

$$(4.3d) \qquad |\Theta[X_1]| \leq c_{\theta,4} \qquad \text{and} \qquad \|\widehat{\Theta}_J(\boldsymbol{X_1})\|_p \leq c_{\theta,4}.$$

The properties formulated here for Θ and $\widehat{\Theta}$ must be satisfied by Θ^{f} and $\widehat{\Theta}^{\mathrm{f}}$ as well.

Remark 4.4. By setting $J = 1$ in Assumption 4.3(i–ii), another property emerges:

$$(4.4) \qquad |\Theta[X_1] - \Theta[X_2]| \le (c_{\theta,1} + c_{\theta,2})\|X_1 - X_2\|_r.$$

This inequality, then, does not need to be proven separately.

Assumption 4.1 pertains to $\mathcal{G}$ and hence must be checked on a case-by-case basis. Assumptions 4.2 and 4.3 are discussed for the IPMs used in this paper in Appendix A.

4.2. Single-level convergence. This subsection studies the properties of (3.2) as an approximation to the intractable (3.6). More specifically, we propose a choice of L and J and discuss the resulting cost-to-error relation.

THEOREM 4.5. *Let $\epsilon > 0$. If Assumptions 4.1 to 4.3 are satisfied and*

$$(4.5) \qquad L = \left\lfloor \frac{2\log_\rho(\epsilon^{-1})}{\beta} \right\rfloor \qquad and \qquad J \eqsim_\epsilon \epsilon^{-2},$$

then for every $p \ge 2$, there exists an $\epsilon_0 > 0$ such that we have that

$$(4.6) \qquad \|Q_N^L - \bar{q}_N\|_p \lesssim_\epsilon \epsilon \qquad when \qquad \epsilon \le \epsilon_0$$

when running the single-level simulation algorithm from subsection 3.1, *with a cost*

$$(4.7) \qquad \mathrm{Cost} \eqsim_\epsilon \epsilon^{-(2+2\gamma/\beta)}.$$

Proof. The proof is given in section 5. □

4.3. Multilevel convergence. Our main theorem bounds the asymptotic cost-to-error relation of the single-ensemble MLIPM algorithm from subsection 3.3, given that the number of levels and of particles on each level are chosen in a specified way. It generalizes [23, Theorem 3.2] from the EnKF case to our framework.

THEOREM 4.6. *Let $\epsilon > 0$. If Assumptions 4.1 to 4.3 are satisfied and*

$$(4.8) \qquad L = \left\lfloor \frac{2\log_\rho(\epsilon^{-1})}{\beta} \right\rfloor \qquad and \qquad J_\ell \eqsim_{\{\ell,\epsilon\}} \rho^{-\frac{\beta+2\gamma}{3}\ell} \begin{cases} \rho^{\beta L} & if\ \beta > \gamma, \\ L^2 \rho^{\beta L} & if\ \beta = \gamma, \\ \rho^{\frac{\beta+2\gamma}{3}L} & if\ \beta < \gamma, \end{cases}$$

then for every $p \ge 2$, there exists an $\epsilon_0 > 0$ such that we have that

$$(4.9) \qquad \|Q_N^{\mathrm{ML}} - \bar{q}_N\|_p \lesssim_\epsilon \epsilon \log_\rho(\epsilon^{-1})^N \qquad when\ \epsilon \le \epsilon_0$$

when running the multilevel simulation algorithm from subsection 3.3, *with a cost*

$$(4.10) \qquad \mathrm{Cost} \eqsim_\epsilon \begin{cases} \epsilon^{-2} & if\ \beta > \gamma, \\ \epsilon^{-(2+\delta)} & if\ \beta = \gamma, \qquad for\ any\ \delta > 0. \\ \epsilon^{-2\gamma/\beta} & if\ \beta < \gamma, \end{cases}$$

Proof. The proof is given in section 6. □

Remark 4.7 (On the extra factor in (4.9)). The factor $\log_\rho(\epsilon^{-1})^N$ in (4.9) also appears in the bounds of [23] and its follow-up work [9]. Like us, they note that this factor does not manifest in numerical tests. This is important for the feasibility of the method: while the asymptotic effect of the factor is limited since $\log_\rho(\epsilon^{-1})^N \epsilon \lesssim_\epsilon \epsilon^{1-\delta}$ for all $\delta > 0$, it would introduce an enormous constant when N is moderate or large.

5. Proof of Theorem 4.5. To prove the theorem, we study the random variable $\boldsymbol{U}_n^L$ of ensembles of particles following the IPM dynamics. The p-norms in the proof below are with respect to realizations $\boldsymbol{u}_n^L = \{u_n^{L,j}\}_j$ of this variable and to the various sources of noise. For each of those realizations, we designate an arbitrary but fixed particle $u_n^L := u_n^{L,1}$. To this particle, we correlate another particle $\bar{u}_n^\ell$ on every level ℓ, following the dynamics

$$(5.1) \qquad \bar{u}_{n+1}^\ell = \Psi_n(\bar{u}_n^\ell, \mathcal{G}^\ell(\bar{u}_n^\ell), \Theta[\bar{X}_n], \eta_n)$$

with $\bar{u}_0^\ell = u_0^L$ and using the same η_n as u_n^L. We will make use of $\bar{u}_n^L$ and $\bar{u}_n := \bar{u}_n^\infty$.

We choose $\epsilon_0 > 0$ as some ϵ such that, whenever we use Assumptions 4.1 to 4.3 in this section, their locality conditions can be satisfied with $U_0 = \bar{U}_n$, $X_0 = \bar{X}_n$, and $\Theta_0 = \Theta[\bar{X}_n]$. Such ϵ_0 must exist, as our proof starts at time 0 (where all particles coincide for any ϵ) and continues by induction. Hence, we can use Assumptions 4.1 to 4.3 throughout this section. In addition, Assumptions 4.1(ii) and 4.3(iv) will ensure that X_1, X_2, and Θ_2 in Assumption 4.2 will always be in $L^{\geq 2}(\Omega, \mathbb{R}^{d_x})$, as required.

The main challenge of the proof is to show that, at each time index n,

$$(5.2) \qquad \|\bar{U}_n - U_n^L\|_p \lesssim \epsilon.$$

That is, given the same driving noise η_n, a particle following (3.2) and one following (3.6) are at a distance from each other whose p-norm is at most proportional to ϵ.

To prove (5.2), we can use the triangle inequality to write

$$(5.3) \qquad \|\bar{U}_n - U_n^L\|_p \leq \|\bar{U}_n - \bar{U}_n^L\|_p + \|\bar{U}_n^L - U_n^L\|_p.$$

The right-hand side terms are dealt with in Lemmas 5.1 and 5.4, with Lemma 5.3 providing an auxiliary result for the latter. Corollary 5.2 will be useful in section 6.

LEMMA 5.1. *It holds that*

$$(5.4) \qquad \|\bar{U}_n - \bar{U}_n^L\|_p \lesssim \epsilon.$$

Proof. For $n = 0$, the statement definitely holds, as $\bar{u}_0 = \bar{u}_0^L$. We now proceed by induction. If (5.4) holds for n, then, by applying Assumptions 4.2 and 4.1(i, iii),

$$\begin{aligned}
\|\bar{U}_{n+1} - \bar{U}_{n+1}^L\|_p &= \|\Psi_n(\bar{U}_n, \mathcal{G}(\bar{U}_n), \Theta[\bar{X}_n], \eta_n) - \Psi_n(\bar{U}_n^L, \mathcal{G}^L(\bar{U}_n^L), \Theta[\bar{X}_n], \eta_n)\|_p \\
&\leq c_\psi \|\bar{X}_n - \bar{X}_n^L\|_r \\
&\leq c_\psi\big((1 + c_{g,1})\|\bar{U}_n - \bar{U}_n^L\|_{r'} + c_{g,3}\rho^{-\beta L/2}\big) \lesssim \epsilon.
\end{aligned}$$

The last inequality holds since $\|\bar{U}_n - \bar{U}_n^L\|_{r'} \lesssim \epsilon$ due to the induction hypothesis, and since $\rho^{-\beta L/2} \approx \epsilon$ due to (4.5). All other factors are independent of ϵ. $\qquad\square$

COROLLARY 5.2. *From Lemma* 5.1 *follows that*

$$(5.5) \qquad \|\bar{U}_n^\ell - \bar{U}_n^{\ell-1}\|_p \lesssim \rho^{-\beta\ell/2}.$$

Proof. Let us make the relation between L and ϵ in (5.4) explicit:

$$(5.6) \qquad \|\bar{U}_n - \bar{U}_n^{L(\epsilon)}\|_p \lesssim \epsilon \Leftrightarrow \|\bar{U}_n - \bar{U}_n^L\| \lesssim \epsilon(L).$$

Then we can bound

$$\begin{aligned}
\|\bar{U}_n^\ell - \bar{U}_n^{\ell-1}\|_p &\leq \|\bar{U}_n - \bar{U}_n^\ell\|_p + \|\bar{U}_n - \bar{U}_n^{\ell-1}\| \\
&\lesssim \epsilon(\ell) + \epsilon(\ell-1) \lesssim \rho^{-\beta\ell/2} + \rho^{-\beta(\ell-1)/2} \leq (1 + \rho^{\beta/2})\rho^{-\beta\ell/2} \lesssim \rho^{-\beta\ell/2}\square
\end{aligned}$$

LEMMA 5.3. *It holds that*

(5.7)
$$|\Theta[\bar{X}_n] - \Theta[\bar{X}_n^L]| \lesssim \epsilon.$$

Proof. Application of Assumption 4.3(i) yields

$$|\Theta[\bar{X}_n] - \Theta[\bar{X}_n^L]| \leq c_{\theta,1}\|\bar{X}_n - \bar{X}_n^L\|_r$$
$$\leq c_{\theta,1}\left((1+c_{g,1})\|\bar{U}_n - \bar{U}_n^L\|_{r'} + c_{g,3}\rho^{-\beta L/2}\right) \lesssim \epsilon,$$

where the second inequality holds due to Assumption 4.1(i, iii). The last inequality uses Lemma 5.1 and (4.5). □

LEMMA 5.4. *It holds that*

(5.8a)
$$\|\bar{U}_n^L - U_n^L\|_p \lesssim \epsilon,$$

(5.8b)
$$\|\Theta[\bar{X}_n] - \widehat{\Theta}(\boldsymbol{X}_n^L)\|_p \lesssim \epsilon.$$

Proof. The proof is one by induction. At $n = 0$, (5.8a) is clearly true. We first prove that (5.8b) follows from (5.8a):

$$\|\Theta[\bar{X}_n] - \widehat{\Theta}(\boldsymbol{X}_n^L)\|_p$$
$$\leq |\Theta[\bar{X}_n] - \Theta[\bar{X}_n^L]| + \|\Theta[\bar{X}_n^L] - \widehat{\Theta}(\bar{\boldsymbol{X}}_n^{\boldsymbol{L}})\|_p + \|\widehat{\Theta}(\bar{\boldsymbol{X}}_n^{\boldsymbol{L}}) - \widehat{\Theta}(\boldsymbol{X}_n^L)\|_p$$
$$\leq |\Theta[\bar{X}_n] - \Theta[\bar{X}_n^L]| + c_{\theta,3}J^{-1/2} + c_{\theta,1}\|\bar{X}_n^L - X_n^L\|_r$$
$$\lesssim \epsilon.$$

For the second inequality, we used Assumptions 4.3(iii) and 4.3(i). For the last inequality, the first term was bounded as $\lesssim \epsilon$ due to Lemma 5.3, the second due to (4.5), and the last due to the induction hypothesis combined with Assumption 4.1(i).

If (5.8) is satisfied at time index n, then at $n+1$, by Assumptions 4.1 and 4.2(i),

$$\|\bar{U}_{n+1}^L - U_{n+1}^L\|_p = \|\Psi_n(\bar{U}_n^L, \mathcal{G}^L(\bar{U}_n^L), \Theta[\bar{X}_n], \eta_n) - \Psi_n(U_n^L, \mathcal{G}^L(U_n^L), \widehat{\Theta}(\boldsymbol{X}_n^L), \eta_n)\|_p$$
$$\leq c_\psi\|\bar{X}_n^L - X_n^L\|_r + c_\psi\|\Theta[\bar{X}_n] - \widehat{\Theta}(\boldsymbol{X}_n^L)\|_r$$
$$\lesssim c_\psi(1+c_{g,1})\epsilon + c_\psi\epsilon$$
$$\lesssim \epsilon.$$
□

Lemmas 5.1 and 5.4 together prove (5.2). The proof of (4.6) is then identical to that of (5.8b), but with Θ^{f} and $\widehat{\Theta}^{\mathrm{f}}$ instead of Θ and $\widehat{\Theta}$, respectively.

6. Proof of Theorem 4.6. For the multilevel convergence proof, we will first need additional definitions. The object of study is $\boldsymbol{U}_n^{\mathrm{ML}}$, the random variable of multilevel ensembles of particles generated by our algorithm in subsection 3.3. Similarly to section 5, we associate to realizations $\boldsymbol{u}_n^{\mathrm{ML}}$ of this random variable another ensemble $\bar{\boldsymbol{u}}_n^{\mathrm{ML}}$, which uses the same η noise for each particle, but replaces the multilevel statistic estimates by the mean-field $\Theta[\bar{X}_n]$. Again we designate $u_n^\ell := u_n^{\ell,\mathrm{F}} := u_n^{\ell,\mathrm{F},1}$ and $u_n^{\ell,\mathrm{C}} := u_n^{\ell,\mathrm{C},1}$. Other notations that will be used are $\boldsymbol{u}_n^{\ell,\mathrm{F}} := \{u_n^{\ell,\mathrm{F},j}\}_j$ and $\boldsymbol{u}_n^{\ell,\mathrm{C}} := \{u_n^{\ell,\mathrm{C},j}\}_j$. For notational convenience, we allow ourselves to write $\Theta[\bar{X}_n^{0,\mathrm{C}}]$, $\Theta[X_n^{0,\mathrm{C}}]$, $\widehat{\Theta}(\bar{\boldsymbol{X}}_n^{0,\mathrm{C}})$, and $\widehat{\Theta}(\boldsymbol{X}_n^{0,\mathrm{C}})$, which all equal zero.

We set an ϵ_0 analogously to that in section 5, which means that, in this section, Assumptions 4.1 to 4.3 can always be used.

LEMMA 6.1. *It holds that*

$$\|\widehat{\Theta}^{\mathrm{ML}}(\bar{\boldsymbol{X}}_n^{\mathrm{ML}}) - \Theta[\bar{X}_n^L]\|_p \lesssim \epsilon. \tag{6.1}$$

Proof. We use the definition (3.10) of $\widehat{\Theta}^{\mathrm{ML}}$ to decompose the left-hand side of (6.1), write $\Theta[\bar{X}_n^L]$ as a telescoping sum, and then use the triangle inequality to get

$$\left\|\widehat{\Theta}^{\mathrm{ML}}(\bar{\boldsymbol{X}}_n^{\mathrm{ML}}) - \Theta[\bar{X}_n^L]\right\|_p \leq \sum_{\ell=0}^{L}\left\|\left(\widehat{\Theta}(\bar{\boldsymbol{X}}_n^{\ell,\mathrm{F}}) - \widehat{\Theta}(\bar{\boldsymbol{X}}_n^{\ell,\mathrm{C}})\right) - \left(\Theta[\bar{X}_n^{\ell,\mathrm{F}}] - \Theta[\bar{X}_n^{\ell,\mathrm{C}}]\right)\right\|_p.$$

Then, there exists some $L' \geq 0$ that is independent of ϵ such that Assumption 4.3(ii) with $X_0 = \bar{X}_n^{L',\mathrm{F}}$ bounds each term with $\ell > L'$ as[2]

$$\left\|\left(\widehat{\Theta}(\bar{\boldsymbol{X}}_n^{\ell,\mathrm{F}}) - \widehat{\Theta}(\bar{\boldsymbol{X}}_n^{\ell,\mathrm{C}})\right) - \left(\Theta[\bar{X}_n^{\ell,\mathrm{F}}] - \Theta[\bar{X}_n^{\ell,\mathrm{C}}]\right)\right\|_p \leq c_{\theta,2} J_\ell^{-1/2}\|\bar{X}_n^{\ell,\mathrm{F}} - \bar{X}_n^{\ell,\mathrm{C}}\|_r, \tag{6.2}$$

while Assumption 4.3(iii) with $X_0 \in \{\bar{X}_n^{\ell,\mathrm{F}}, \bar{X}_n^{\ell,\mathrm{C}}\}$ ensures that the others are at most $2c_{\theta,3} J_\ell^{-1/2}$. By Assumption 4.1(i) and Corollary 5.2 with $U_0 = \bar{U}_n^\ell$, we conclude that

$$
\begin{aligned}
\|\widehat{\Theta}^{\mathrm{ML}}(\bar{\boldsymbol{X}}_n^{\mathrm{ML}}) - \Theta[\bar{X}_n^L]\|_p &\lesssim \sum_{\ell=0}^{L} J_\ell^{-1/2}\rho^{-\beta\ell/2} + \sum_{\ell=0}^{L'} J_\ell^{-1/2} \\
&\leq \left(\left(\sum_{\ell=0}^{L}\rho^{\frac{\gamma-\beta}{3}\ell}\right) + L'\rho^{(\beta+2\gamma)L'/6}\right)\begin{cases}\rho^{-\beta L/2} & \text{if } \beta > \gamma \\ L^{-1}\rho^{-\beta L/2} & \text{if } \beta = \gamma \\ \rho^{-(\beta+2\gamma)L/6} & \text{if } \beta < \gamma\end{cases} \\
&\lesssim \epsilon + \epsilon \lesssim \epsilon,
\end{aligned} \tag{6.3}
$$

where the first term of the first inequality starts at $\ell = 0$ instead of $\ell = L' + 1$ to simplify the calculations, while clearly still providing an upper bound. $\square$

LEMMA 6.2. *It holds that*

$$\|\bar{U}_n^{\ell,\{\mathrm{F},\mathrm{C}\}} - U_n^{\ell,\{\mathrm{F},\mathrm{C}\}}\|_p \lesssim \epsilon \log_\rho(\epsilon^{-1})^{n-1}, \tag{6.4a}$$

$$\|\widehat{\Theta}^{\mathrm{ML}}(\bar{\boldsymbol{X}}_n^{\mathrm{ML}}) - \widehat{\Theta}^{\mathrm{ML}}(\boldsymbol{X}_n^{\mathrm{ML}})\|_p \lesssim \epsilon \log_\rho(\epsilon^{-1})^n. \tag{6.4b}$$

Proof. The proof is by induction. When $n = 0$, both left-hand sides are zero. If both inequalities hold at index n, we have (by Assumption 4.2 with $X_0 = \bar{X}_n^{\ell,\{\mathrm{F},\mathrm{C}\}}$ and $\Theta_0 = \Theta[\bar{X}_n]$) that

$$
\begin{aligned}
\|\bar{U}_{n+1}^{\ell,\{\mathrm{F},\mathrm{C}\}} - U_{n+1}^{\ell,\{\mathrm{F},\mathrm{C}\}}\|_p &\leq c_\psi\|\bar{X}_n^{\ell,\{\mathrm{F},\mathrm{C}\}} - X_n^{\ell,\{\mathrm{F},\mathrm{C}\}}\|_r + c_\psi\|\Theta[\bar{X}_n] - \widehat{\Theta}^{\mathrm{ML}}(\boldsymbol{X}_n^{\mathrm{ML}})\|_r \\
&\lesssim \epsilon + \left(|\Theta[\bar{X}_n] - \Theta[\bar{X}_n^L]| + \|\Theta[\bar{X}_n^L] - \widehat{\Theta}^{\mathrm{ML}}(\bar{\boldsymbol{X}}_n^{\mathrm{ML}})\|_r \right. \\
&\qquad\left. + \|\widehat{\Theta}^{\mathrm{ML}}(\bar{\boldsymbol{X}}_n^{\mathrm{ML}}) - \widehat{\Theta}^{\mathrm{ML}}(\boldsymbol{X}_n^{\mathrm{ML}})\|_r\right) \\
&\lesssim \epsilon + \epsilon + \epsilon + \epsilon \log_\rho(\epsilon^{-1})^n \lesssim \epsilon \log_\rho(\epsilon^{-1})^n.
\end{aligned}
$$

The first ϵ appears due to the induction hypothesis and Assumption 4.1(i) with $U_0 = \bar{U}_n^\ell$. The other terms follow from Lemma 5.3, Lemma 6.1, and the induction hypothesis.

---

[2] L' is chosen such that the conditions $\|\bar{X}_n^{\ell,\mathrm{F}} - \bar{X}_n^{L',\mathrm{F}}\|_r \leq d$ and $\|\bar{X}_n^{\ell,\mathrm{C}} - \bar{X}_n^{L',\mathrm{F}}\|_r \leq d$ of Assumption 4.3 are satisfied for all $\ell > L'$. Clearly, L' is independent of ϵ.

For the second induction step, let (6.4b) hold for n and (6.4a) for $n+1$. Then

$$\left\|\widehat{\Theta}^{\mathrm{ML}}(\bar{\boldsymbol{X}}_{n+1}^{\mathrm{ML}}) - \widehat{\Theta}^{\mathrm{ML}}(\boldsymbol{X}_{n+1}^{\mathrm{ML}})\right\|_p$$

$$= \left\|\sum_{\ell=0}^{L} \left(\widehat{\Theta}(\bar{\boldsymbol{X}}_{n+1}^{\ell,\mathrm{F}}) - \widehat{\Theta}(\bar{\boldsymbol{X}}_{n+1}^{\ell,\mathrm{C}})\right) - \left(\widehat{\Theta}(\boldsymbol{X}_{n+1}^{\ell,\mathrm{F}}) - \widehat{\Theta}(\boldsymbol{X}_{n+1}^{\ell,\mathrm{C}})\right)\right\|_p$$

$$= \left\|\sum_{\ell=0}^{L} \left(\widehat{\Theta}(\bar{\boldsymbol{X}}_{n+1}^{\ell,\mathrm{F}}) - \widehat{\Theta}(\boldsymbol{X}_{n+1}^{\ell,\mathrm{F}})\right) - \left(\widehat{\Theta}(\bar{\boldsymbol{X}}_{n+1}^{\ell,\mathrm{C}}) - \widehat{\Theta}(\boldsymbol{X}_{n+1}^{\ell,\mathrm{C}})\right)\right\|_p$$

(6.5)
$$\leq \left|\sum_{\ell=0}^{L} \left(\Theta[\bar{X}_{n+1}^{\ell,\mathrm{F}}] - \Theta[X_{n+1}^{\ell,\mathrm{F}}]\right) - \left(\Theta[\bar{X}_{n+1}^{\ell,\mathrm{C}}] - \Theta[X_{n+1}^{\ell,\mathrm{C}}]\right)\right|$$

$$+ \left\|\sum_{\ell=0}^{L} \left(\widehat{\Theta}(\bar{\boldsymbol{X}}_{n+1}^{\ell,\mathrm{F}}) - \widehat{\Theta}(\boldsymbol{X}_{n+1}^{\ell,\mathrm{F}})\right) - \left(\Theta[\bar{X}_{n+1}^{\ell,\mathrm{F}}] - \Theta[X_{n+1}^{\ell,\mathrm{F}}]\right)\right\|_p$$

$$+ \left\|\sum_{\ell=0}^{L} \left(\widehat{\Theta}(\bar{\boldsymbol{X}}_{n+1}^{\ell,\mathrm{C}}) - \widehat{\Theta}(\boldsymbol{X}_{n+1}^{\ell,\mathrm{C}})\right) - \left(\Theta[\bar{X}_{n+1}^{\ell,\mathrm{C}}] - \Theta[X_{n+1}^{\ell,\mathrm{C}}]\right)\right\|_p.$$

For the first term, we recall that $\Theta[\bar{X}_{n+1}^{\ell,\mathrm{F}}] = \Theta[\bar{X}_{n+1}^{\ell}]$ and $\Theta[\bar{X}_{n+1}^{\ell,\mathrm{C}}] = \Theta[\bar{X}_{n+1}^{\ell-1}]$, and an analogous result holds for the variants without bars. The term is thus a telescoping sum equal to $|\Theta[\bar{X}_{n+1}^{L}] - \Theta[X_{n+1}^{L}]|$. If Assumption 4.3(ii) would apply to the second and third terms, we could bound the norms of the sums by ϵ and avoid the logarithmic term in (4.9). Unfortunately, the particles in the ensembles $\boldsymbol{X}_{n+1}^{\ell,\{\mathrm{F},\mathrm{C}\}}$ are not independent. Hence, we use the third line of (6.5) and bound it as

(6.6)
$$\left\|\sum_{\ell=0}^{L} \left(\widehat{\Theta}(\bar{\boldsymbol{X}}_{n+1}^{\ell,\mathrm{F}}) - \widehat{\Theta}(\boldsymbol{X}_{n+1}^{\ell,\mathrm{F}})\right) - \left(\widehat{\Theta}(\bar{\boldsymbol{X}}_{n+1}^{\ell,\mathrm{C}}) - \widehat{\Theta}(\boldsymbol{X}_{n+1}^{\ell,\mathrm{C}})\right)\right\|_p$$

$$\lesssim \log_\rho(\epsilon^{-1}) \max_{0 \leq \ell \leq L} \left(\|\bar{\boldsymbol{X}}_{n+1}^{\ell,\mathrm{F}} - \boldsymbol{X}_{n+1}^{\ell,\mathrm{F}}\|_r + \|\bar{\boldsymbol{X}}_{n+1}^{\ell,\mathrm{C}} - \boldsymbol{X}_{n+1}^{\ell,\mathrm{C}}\|_r\right)$$

$$\lesssim \epsilon \log_\rho(\epsilon^{-1})^{n+1}.$$

For the first inequality, we used (4.8) and Assumption 4.3(i) with $X_0 = \bar{X}_{n+1}^{\ell,\{\mathrm{F},\mathrm{C}\}}$. The second is due to Assumption 4.1(i) with $U_0 = \bar{U}_{n+1}^{\ell,\{\mathrm{F},\mathrm{C}\}}$ and the induction hypothesis.□

The proof of (4.9) is then equivalent to the combination of those of (6.4b) and (6.1), but with Θ^{f} and $\widehat{\Theta}^{\mathrm{f}}$ instead of Θ and $\widehat{\Theta}$, respectively.

7. Infinite time horizons. For IPMs such as EKI and EKS (Examples 1.2 and 1.3), the goal is to approximate the mean-field particle distribution for $n \to \infty$. In practice, a finite truncation must be made. This can be achieved by simulating the IPM as in section 3, but with an added dependence of the number of time steps N on the error tolerance. We detail the differences in methodology, cost, and accuracy. In this section, we assume that $\|Q_N^{\mathrm{ML}} - \bar{q}_N\|_p \lesssim_\epsilon \epsilon$ without the $\log_\rho(\epsilon^{-1})^N$ factor – see Remark 4.7 for a discussion and section 8 for supporting numerical tests.

Denote by r a bound for the convergence of the QoI of the mean-field in time:

(7.1)
$$|\bar{q}_N - \bar{q}_\infty| \leq r(N).$$

Let Q_N be equal to either Q_N^L or Q_N^{ML}, depending on the context. Suppose we want the error $\|Q_N - \bar{q}_\infty\|_p$ to be at most asymptotically proportional to some $\mathtt{tol} > 0$:

(7.2)
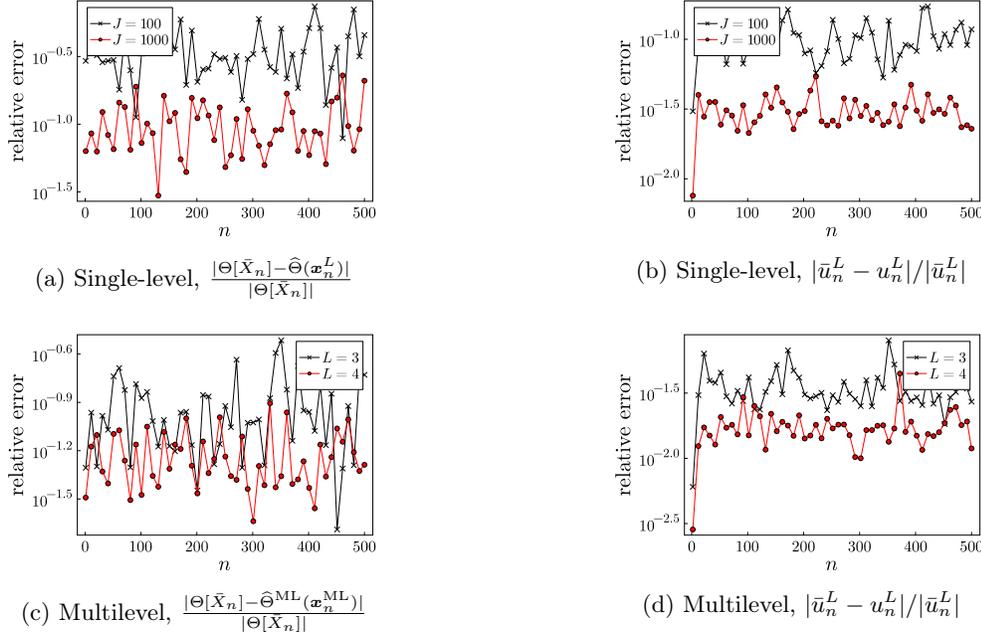$$\|Q_N - \bar{q}_N\|_p + |\bar{q}_N - \bar{q}_\infty| \lesssim \mathtt{tol}.$$

(a) Single-level, $\dfrac{|\Theta[\bar{X}_n] - \widehat{\Theta}(\boldsymbol{x}_n^L)|}{|\Theta[\bar{X}_n]|}$

(b) Single-level, $|\bar{u}_n^L - u_n^L|/|\bar{u}_n^L|$

(c) Multilevel, $\dfrac{|\Theta[\bar{X}_n] - \widehat{\Theta}^{\mathrm{ML}}(\boldsymbol{x}_n^{\mathrm{ML}})|}{|\Theta[\bar{X}_n]|}$

(d) Multilevel, $|\bar{u}_n^L - u_n^L|/|\bar{u}_n^L|$

Fig. 7.1: EKS errors on the statistic and particle paths compared to the mean field.

We then adapt our algorithms from subsections 3.1 and 3.3 to simulate not until a fixed N, but until $N(\mathtt{tol}) = r^{-1}(\mathtt{tol})$. This prompts us to make the N-dependence in the asymptotic inequality (4.6) or (4.9) explicit; define a function s such that

$$(7.3) \qquad \|Q_N - \bar{q}_N\|_p \lesssim s(N)\epsilon.$$

This means that we can choose (with "ih" or "fh" denoting infinite or fixed horizons)

$$(7.4) \qquad \epsilon_{\mathrm{ih}}(\mathtt{tol}) \approx \frac{\mathtt{tol}}{s(r^{-1}(\mathtt{tol}))} \quad \text{with} \quad \mathrm{Cost}_{\mathrm{ih}} \approx r^{-1}(\mathtt{tol})\mathrm{Cost}_{\mathrm{fh}}\left(\frac{\mathtt{tol}}{s(r^{-1}(\mathtt{tol}))}\right)$$

and, through (7.3), achieve the desired $\|Q_N - \bar{q}_N\|_p \lesssim \mathtt{tol}$. Notably, if $s(N) \approx 1$ (i.e., if the constant hidden in (4.6) or (4.9) is N-independent), we can choose $\epsilon_{\mathrm{ih}} = \mathtt{tol}$ and the cost for an infinite-horizon IPM is $r^{-1}(\mathtt{tol})$ times the fixed-horizon cost.

Deriving s for specific IPMs is outside the scope of this paper, but we will study one case numerically. Consider EKS (Example 1.3) applied to the first test problem from subsection 8.2, with $N = 500$ and $J \in \{100, 1000\}$ (single-level) or $L \in \{3, 4\}$ (multilevel). The constant hidden in (4.6) comes from those in (5.8), which pertain to the distance between the exact parameter $\Theta[\bar{X}_n]$ and the sample statistic $\widehat{\Theta}(\boldsymbol{x}_n^L)$, and that between the particle $\bar{u}_n^L$ using the exact parameter and a correlated particle u_n^L using the statistic. An analogous relation exists between (4.9) and the bounds (6.4). After approximating $\Theta[\bar{X}_n]$ by using 10^5 particles, we plot in Figure 7.1 these quantities evolving over time n. For this IPM and this example, all graphs plateau after a short transition period, at a level proportional to ϵ. This suggests that, for EKS and this forward model, $s(N) \approx 1$, which indicates time-uniform errors. This situation is similar to that of multilevel particle filters [29], whose theoretical error bounds are time-dependent, while numerical tests show time-uniformity.

8. Scaling experiments. In this section, we set out to corroborate our single-level and multilevel asymptotic cost-to-error bounds under the assumptions in subsection 4.1. First, in subsection 8.1, we consider state estimation of an Ornstein–Uhlenbeck process in one and in twenty dimensions, by use of a filtering IPM with a fixed number of discrete time steps. Then we move to Bayesian inverse problems in subsection 8.2, and apply various infinite-horizon IPMs to state estimation in an ODE, a set-up with the time-dependent heat equation, and Darcy flow. We show empirically that our analytic cost-to-error rate still holds. Finally, we discuss and interpret our results in subsection 8.3. Our code is open-source and can be found in the repository
https://gitlab.kuleuven.be/numa/public/paper-code-mlipm.

8.1. IPMs with a fixed number of time steps. Our rigorous convergence results in Theorems 4.5 and 4.6 apply to IPMs with a fixed number of time steps. One such method, the ensemble Kalman filter, is extensively studied in [23], whose MLEnKF algorithm matches our method when applied to the EnKF. Hence, we do not repeat their experiments; we instead study an alternative that can now also be simulated in a multilevel way within our framework. It is introduced in Example 8.1.

Example 8.1 (Deterministic ensemble Kalman filter). The work [37] proposes the deterministic ensemble Kalman filter (DEnKF) as an alternative to the EnKF from Example 1.1. It uses the dynamics

$$(8.1) \qquad u_{n+1}^j = (I - \frac{1}{2}K_n H)(\mathcal{G}(u_n^j) - E_J(\boldsymbol{g}_n)) + E_J(\boldsymbol{g}_n) + K_n(y_{n+1} - H\mathcal{G}(u_n^j)),$$

where K_n is the Kalman gain introduced in Example 1.1.

An Ornstein–Uhlenbeck process. We mean to estimate the state of a d_u-dimensional Ornstein–Uhlenbeck process with parameters θ and σ, described by

$$(8.2) \qquad \mathrm{d}u = -\theta u \mathrm{d}t + \sigma \mathrm{d}W, \qquad u(0) = \left[1, 1, \ldots, 1\right]^\top \in \mathbb{R}^{d_u},$$

based on noisy measurements of u. We will consider two cases: $d_u = 1$ and $d_u = 20$. The one-dimensional variant of this problem was also used as an example in [23]. When $d_u = 1$, we choose scalar parameters $\theta = \sigma = 0.3$; when $d_u = 20$, we set

$$(8.3) \qquad \theta = \frac{1}{10}\begin{bmatrix} 3 & 1 & & \\ 1 & \ddots & \ddots & \\ & \ddots & \ddots & 1 \\ & & 1 & 3 \end{bmatrix} \in \mathbb{R}^{20\times20} \quad \text{and} \quad \sigma = \frac{1}{10}\begin{bmatrix} 3 & & \\ & \ddots & \\ & & 3 \end{bmatrix} \in \mathbb{R}^{20\times20}.$$

Noisy observations y_n of the process (8.2) are available at times $t_n = n$. Let us call $\mathcal{G}$ the operator that evolves (8.2) exactly over a time interval of 1. Then,

$$(8.4) \qquad \begin{aligned} u_{n+1} &= \mathcal{G}(u_n), \\ y_n &= u_n + \eta_n, \end{aligned}$$

with measurement noise $\eta_n \sim N(0, \Gamma)$. We assume that computing an exact solution to (8.2) is infeasible and introduce approximations $\{\mathcal{G}^\ell\}_\ell$, which are Milstein discretizations at resolutions $h^\ell = 2^{-\ell}$. In other words, $\mathcal{G}^\ell(u_n) = u_{n,2^\ell}$, where

$$(8.5) \qquad \widehat{u}_{n,k+1}^\ell = (1 - h^\ell\theta)\widehat{u}_{n,k}^\ell + \sigma\sqrt{h^\ell}\Delta W_{n,k}^\ell \qquad \text{and} \qquad \widehat{u}_{n,0}^\ell = u_n.$$

To ensure correlation between a pair of particles on level ℓ, the Brownian increments

$$(8.6) \qquad \Delta W_{n,k}^\ell = W(t_n + (k+1)h^\ell) - W(t_n + kh^\ell)$$

are computed from the same Brownian path W. The cost of (8.5) scales with ℓ as

$$(8.7) \qquad \mathrm{Cost}(\mathcal{G}^\ell) = 2^\ell\mathrm{Cost}(\mathcal{G}^0).$$
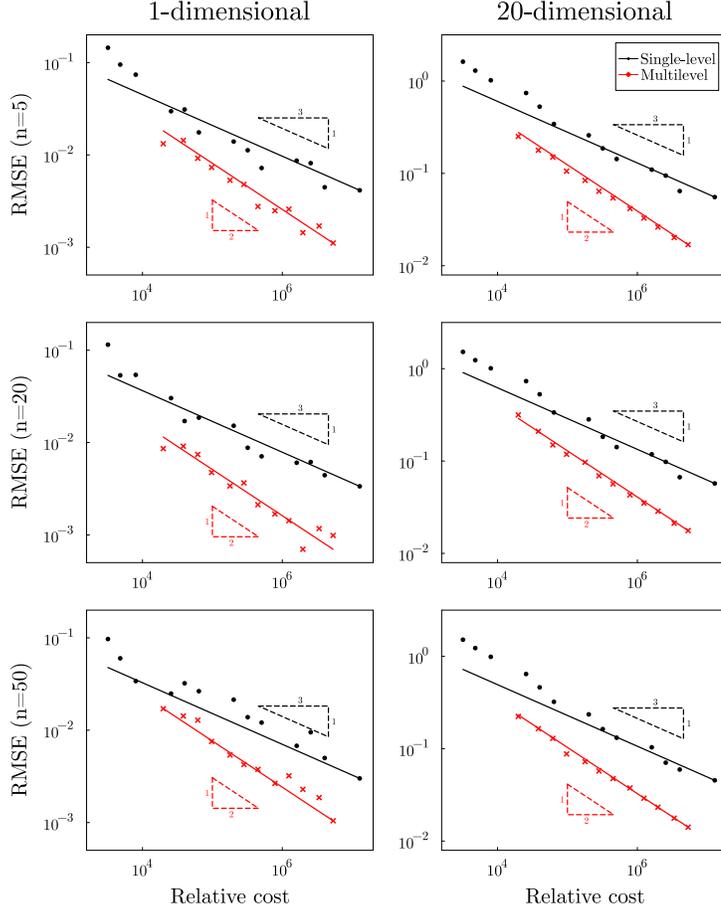
Fig. 8.1: RMSE as a function of relative cost (total cost compared to the cost of $\mathcal{G}^0$) for state estimation of an Ornstein–Uhlenbeck process with DEnKF.

Therefore $\gamma = 1$. We also have $\beta = 2$ (see [19, 23]). For the QoI of the filtering distribution, we choose the particle mean. As a gold standard, a highly accurate approximation of the expected value of the mean-field particle distribution is obtained by averaging the means of 10 large-scale ($J = 10^4$) single-level simulations using a very accurate forward model ($L = 12$). We consider the root-mean-square error (RMSE) between 10 independent ensemble means and this gold standard.

The cost of simulating a single-level or multilevel ensemble can be quantified as

$$(8.8a) \qquad \mathrm{Cost}^{\mathrm{SL}} = J2^{L\gamma}N\mathrm{Cost}(\mathcal{G}^0),$$

$$(8.8b) \qquad \mathrm{Cost}^{\mathrm{ML}} = \Big(J_0 + \sum\nolimits_{\ell=1}^{L} J_\ell(2^{\ell\gamma} + 2^{(\ell-1)\gamma})\Big)N\mathrm{Cost}(\mathcal{G}^0).$$

In Figure 8.1, this cost is plotted against the RMSE for a number of both single-level and multilevel experiments. The values of L, J, and J_ℓ are chosen in accordance with the scaling laws set out in (4.5) and (4.8). From β and γ, by Theorems 4.5 and 4.6, we predict RMSE $\lesssim (\mathrm{Cost}^{\mathrm{SL}})^{-1/3}$ for single-level and RMSE $\lesssim (\mathrm{Cost}^{\mathrm{ML}})^{-1/2}$ for multilevel. The figure shows that the multilevel scheme adheres closely to this rate at all time points $N \in \{5, 20, 50\}$. The single-level method initially converges somewhat faster than its expected rate, but later decreases its slope to $-1/3$.

8.2. IPMs with an infinite number of time steps. Section 7 discussed applying our single- and multilevel simulation algorithms to IPMs whose time horizon grows as the required accuracy increases. In this subsection, we test this idea numerically using the EKI and EKS methods introduced earlier. The first example problem will use the base versions of these methods. For the other two examples, to stabilize and accelerate both the single-level and the multilevel methods, we use the adaptive time steps introduced in [30] and also used in [17]. This step is defined as $\tau_n = \alpha/(|D_n|_F + \varepsilon)$ for some ensemble-dependent matrix D_n encoding the distance from the particles to each other and to the observation. Here, we denoted by $|\cdot|_F$ the Frobenius norm. We compute τ_n for each level separately, and use their minimum.

We will choose the parameters of our multilevel interacting-particle methods as if $s(N) \approx 1$ in (7.3), which is what was observed for EKS in Figure 7.1. Based on the match between experiments and the expected cost-to-error relation (7.4), we can then assess the accuracy of this assumption. The bound $r(N)$ for the convergence of EKI in time is estimated as $r(N) \approx N^{-1}$, motivated by [26]. EKS and adaptive versions of both methods converges exponentially [17], but since the base of this exponential is unknown, we choose the conservative $r(N) \approx N^{-1/\delta}$ (with $\delta = 0.1$).

Low-dimensional Bayesian inversion: A simple ODE. As a first problem, we consider the identification of the parameter $u = [u_1, u_2]^\top$ in the linear ODE

$$(8.9) \qquad\qquad v'(t) = u_2 v(t), \qquad v(0) = u_1$$

given the observation, influenced by measurement noise $\eta \sim \mathcal{N}(0, \Gamma)$,

$$(8.10) \qquad\qquad y = \mathcal{G}(u) + \eta := \begin{bmatrix} v(0.25) \\ v(0.75) \end{bmatrix} + \eta.$$

While the solution $v(t) = \exp(u_2 t) u_1$ is trivial to compute, we will purport for this exploratory example that it must be approximated numerically with a hierarchy of implicit-Euler schemes with resolutions $h^\ell \approx 2^{2+\ell/2}$, such that $\beta = 1$ and $\gamma = 1/2$:

$$(8.11) \qquad \mathcal{G}^\ell(u) = \begin{bmatrix} v^\ell(0.25) \\ v^\ell(0.75) \end{bmatrix} \qquad \text{where} \qquad v^\ell(t) = u_1(1 - h^\ell u_2)^{-t/h^\ell}.$$

We use the observed value $y = [0.85, 0.65]^\top$ and a noise covariance $\Gamma = 0.5^2 I_2$. The prior for EKS is set to a zero-centered Gaussian with covariance $\Gamma_0 = I_2$. The initial ensemble is sampled as $u_1 \sim \mathcal{N}(1, 0.25)$ and $u_2 \sim \mathcal{U}([-1, 0])$. We apply two IPMs to the problem of recovering u given y: EKI with time steps $\tau_n = 1$ and EKS with $\tau_n = 0.001$ for $n \leq 50$ and $\tau_n = 0.1$ otherwise. By (7.4), we expect single- and multilevel cost-to-error relations of $\text{RMSE} \lesssim (\text{Cost}^{\text{SL}})^{-1/4}$ and $\text{RMSE} \lesssim (\text{Cost}^{\text{ML}})^{-1/3}$ for EKI, and of $\text{RMSE} \lesssim (\text{Cost}^{\text{SL}})^{-1/(3+\delta)}$ and $\text{RMSE} \lesssim (\text{Cost}^{\text{ML}})^{-1/(2+\delta)}$ for EKS.

When L, J, $\{J_\ell\}_\ell$, and N are chosen according to our scaling laws (4.5), (4.8), and $N = r^{-1}(\epsilon)$, Figure 8.2 shows that these asymptotic rates are matched in practice. A non-asymptotic comparison of single- and multilevel IPMs should not be made on the basis of these experiments; see the discussion in subsection 8.3.

Higher-dimensional Bayesian inversion: A heat equation. A first higher-dimensional inverse problem we consider is concerned with recovering source terms in a time-dependent heat equation, based on noisy measurements. The physical set-up consists of the one-dimensional domain $x \in [0, 20]$ with boundaries of fixed temperature $\theta_{\text{b}} = 293$. All non-boundary integer positions are fitted with heat point sources or temperature sensors; sensors at the odd positions and sources at the even ones.
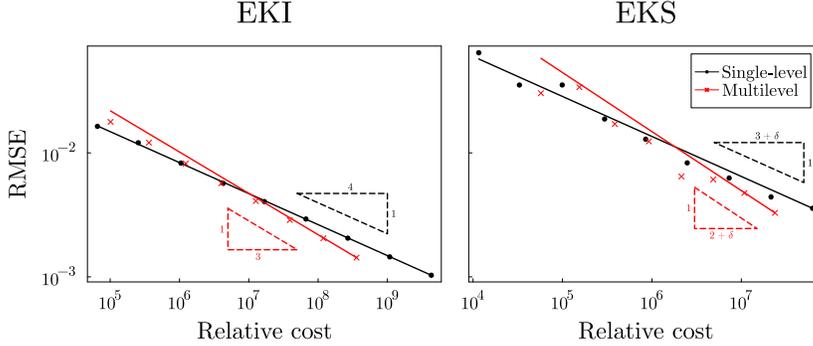
Fig. 8.2: RMSE as a function of relative cost (total cost compared to the cost of $\mathcal{G}^0$) for inversion of a low-dimensional ODE problem. Note that the theoretical multilevel rate hides a logarithmic factor, which is visible in the computational results (which otherwise match the theorized rates well).
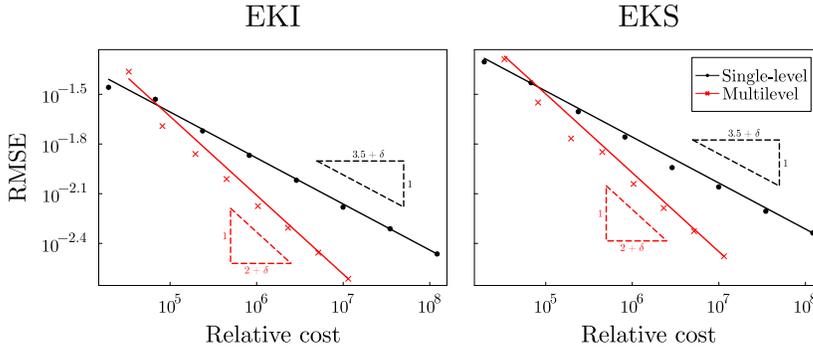

Fig. 8.3: RMSE as a function of relative cost (total cost compared to the cost of $\mathcal{G}^0$) for inversion of a higher-dimensional time-dependent heat problem. Note, as in Figure 8.2, the logarithmic factor visible in the multilevel results.

The sources are modelled as Dirac deltas of unknown but time-invariant magnitude. At time $t = 0$, the entire domain has temperature $\theta(x, 0) = \theta_{\mathrm{b}} = 293$ and the sources are activated. The sensors take measurements at $t = 5$ and at $t = 10$. We thus have $d_u = 9$ unknown sources and $d_g = 20$ measurements.

The forward problem $\mathcal{G}$ is determining $\theta(x^*, t^*)$ at the 10 sensor positions x^* at each of 2 measurement times t^*, given that θ solves the time-dependent heat equation

$$\partial_t \theta = \partial_{xx} \theta + f, \tag{8.12}$$

where the source f is the input. Approximations $\mathcal{G}^\ell$ solve (8.12) with a backward-time centered-space scheme. We ensure that there are always spatial discretization points on the sources and sensors, and temporal ones on the measurement times. With a temporal step $h_t^\ell \approx 2^{(4+\ell)/2}$ and a spatial step $h_x^\ell \approx 2^{(18+\ell)/4}$, we obtain $\beta = 1$ and $\gamma = 3/4$. We model a noise covariance $\Gamma = I$ and, for EKS, use a zero-centered prior with covariance $\Gamma_0 = 10^2 I$. The observation is generated by sampling the prior, solving the forward model at high resolution, and adding noise from $\mathcal{N}(0, \Gamma)$.

We apply EKI and EKS to this problem, both with adaptive time steps; we choose $(\alpha, \varepsilon) = (1, 0.001)$. For this higher-dimensional problem, too, Figure 8.3 reveals good correspondence between the theoretical convergence rates and experimental results.
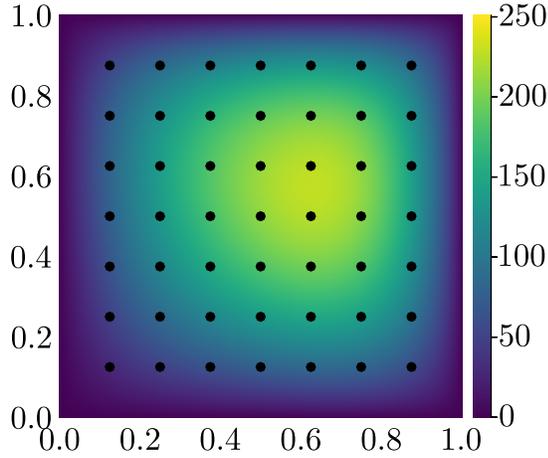
Fig. 8.4: Measurement points in the pressure fields of the Darcy flow problem

Higher-dimensional Bayesian inversion: Darcy flow. Within the literature on Bayesian inversion, the Darcy flow test problem is a classical one (see, e.g., [17, 26]). On the two-dimensional spatial domain $[0, 1]^2$, our forward model computes the map of the positive permeability field $a(x, u)$ of a porous medium to the pressure field $p(x)$ that satisfies

$$(8.13) \qquad \begin{aligned} -\nabla \cdot (a(x, u)\nabla p(x)) &= f(x) \quad \text{if } x \in [0, 1]^2, \\ p(x) &= 0 \qquad \text{if } x \in \partial[0, 1]^2. \end{aligned}$$

We set $f(x_1, x_2) = 1000 \exp(x_1 + x_2)$. The parameter u to the model consists of the first $d_u = 16$ terms in the Karhunen-Loève (KL) expansion

$$(8.14) \qquad \log a(x, u) = \sum_{k \in \mathbb{Z}^2 \setminus \{(0,0)\}} u_k \sqrt{\lambda_k} \phi_k(x),$$

where the eigenpairs are [26]

$$(8.15) \qquad \phi_k(x) = \begin{cases} \sqrt{2} \cos(\pi k_1 x_1) & \text{if } k_2 = 0, \\ \sqrt{2} \cos(\pi k_2 x_2) & \text{if } k_1 = 0, \\ 2 \cos(\pi k_1 x_1) \cos(\pi k_2 x_2) & \text{otherwise,} \end{cases}$$

and

$$(8.16) \qquad \lambda_k = (\pi^2(k_1^2 + k_2^2) + \tau^2)^{-d},$$

and which, for the purpose of truncation, is ordered by descending size of the eigenvalues λ_k. We model $a(x, u)$ as a log-normal random field with covariance $(-\Delta + \tau^2)^{-d}$ that corresponds to a standard normal prior on the coefficients u_k. The output of the forward model consists of p evaluated on 49 equispaced points (as shown in Figure 8.4), and the noise added has covariance $\Gamma = 0.01I$. EKS uses a zero-centered prior with covariance $\Gamma_0 = I$.
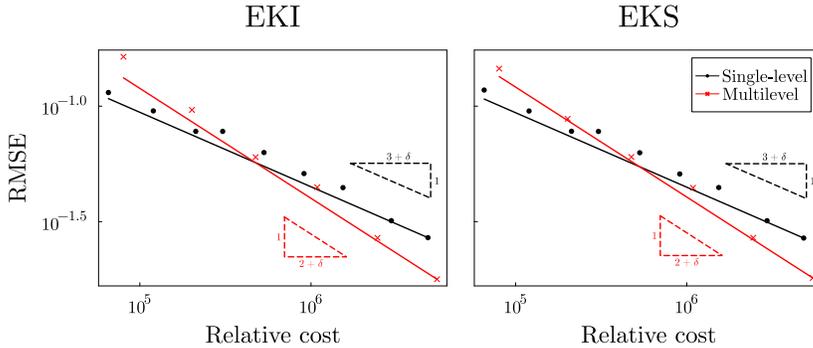
Fig. 8.5: RMSE as a function of relative cost (total cost compared to the cost of $\mathcal{G}^0$) for inversion of a higher-dimensional Darcy flow problem. Note, as in Figure 8.2, the logarithmic factor visible in the multilevel results.

Using central differences for the derivatives in (8.13), with steps $h^\ell \eqsim 2^{(13+\ell)/4}$, yields a forward model with $\beta = 1$ and $\gamma = 1/2$. We show experimental results, next to the theoretically expected rates, in Figure 8.5. As with the other examples, the correspondence between the two is good.

8.3. Discussion. Theorems 4.5 and 4.6 imply that, for IPMs in our framework with a fixed number of time steps, our multilevel simulation algorithm (3.9) asymptotically enjoys faster convergence than the more straightforward single-level algorithm (3.2). The numerical scaling tests conducted in subsections 8.1 and 8.2 demonstrate that these rates are accurate. These conclusions hold for the whole range of IPMs studied, and for both low- and higher-dimensional problems.

We stress that our theory only provides asymptotic rates and does not result in guidelines for selecting the constants in (4.5), (4.8), or – in the case of non-fixed time horizons – the function $r(N)$. We give the following comments.

- Changing these constants should, in general, shift the convergence graphs without changing their slopes.
- Since the constants are not optimized, we cannot compare the performance of single- and multilevel simulation for a fixed cost in a meaningful way. Their relative positions depend on these chosen values for these constants, which can be found in our code repository. However, the convergence rates can be compared, and those are the focus of this article.
- To improve the performance of our multilevel strategy in practice, adaptive or non-adaptive schemes to tune these constants are vital. We comment further on this aspect in the future-work part of section 9.

These limitations are also present in the experiments that are performed in [9, 23]. We intend to tackle this important aspect of multilevel interacting-particle methods in a future paper.

9. Conclusions. We have proposed a general framework for interacting-particle methods, which includes many popular methods used for filtering, optimization, and sampling. Two classes of methods are identified within this framework: those with a fixed and those with an infinite number of discrete time steps. The setting we considered is that where the dynamics of the IPM include an intractable forward model $\mathcal{G}$, to which an approximation hierarchy $\{\mathcal{G}^\ell\}_\ell$ is available.

For IPMs in the framework, we have described and analyzed two methods that allow a numerical approximation of the particle density. The first is a standard Monte Carlo simulation that uses a finite ensemble of particles. As an alternative, we have developed a generalization of the multilevel Monte Carlo approach proposed for the ensemble Kalman filter in [9, 23], which we call *single-ensemble* MLMC. This algorithm uses a single, globally coupled ensemble whose base-level particles and higher-level correction pairs use various different levels ℓ. The MLMC methodology is then used to estimate the interaction term at each time step.

For IPMs with a fixed number of time steps, we formulated Assumptions 4.1 to 4.3, under which rigorous convergence results were shown in Theorems 4.5 and 4.6. These bounds suggest that single-ensemble MLMC always asymptotically outperforms standard Monte Carlo. For a variety of IPMs, we have shown that our assumptions hold, such that the convergence bounds are valid. Experiments suggest that the convergence rates implied by the bounds are accurate. Despite these results, two open questions remain.

- The factor $\log_\rho(\epsilon^{-1})^N$ in the upper bound (4.9) is fortunately absent in practice, as discussed in Remark 4.7. An upper bound for single-ensemble multilevel interacting-particle methods that does not contain this factor would improve the theoretical foundation of the MLMC approach. In this context, we also refer to the discussion in [9, Remark 5].
- Next to this (ϵ, N)-dependent factor, the asymptotic bounds in Theorems 4.5 and 4.6 obscure a constant that is potentially N-dependent. Knowing more about this constant for a given IPM is crucial when the number of time steps depends on the required accuracy (see section 7). Fortunately, our tests in sections 7 and 8 suggest that these effects can be negligible in practice; [23, p. 1833] draws similar conclusions about the EnKF.

There are many interesting directions for further research in three main categories.

- **Theoretical.** The two open questions above require more theoretical work to settle. In addition, theory that suggests good quantitative choices for $\{J_\ell\}_\ell$, as opposed to the proportional rates in (4.8), may enhance performance.
- **Algorithmic.** For IPMs with an ϵ-dependent horizon, it may be feasible to design a single-ensemble MLMC algorithm where only the fine-level particles must take more time steps for increased accuracy. In addition, for problems with high- or infinite-dimensional parameters or data, an extension of our methodology can be developed that uses level-dependent dimensions. The multilevel EnKF in [9] already deals with infinite-dimensional parameters.
- **Numerical.** The MLIPMs developed in this paper should be compared to other, potentially algorithm-specific, MLMC approaches. For an effective comparison, optimized techniques to select the number of particles on each level (either adaptive or theory-based) are first needed.

In addition, this paper lays a clear path for developing multilevel simulation methods for many other interacting-particle methods that have been proposed in the literature. If Assumptions 4.1 to 4.3 hold for the specific forward model, dynamics, and interaction term, then the algorithm (3.9)–(3.11) can be used and Theorem 4.6 applies.

Appendix A. Proofs about specific IPMs. In this appendix, we give proofs about various interacting-particle methods and parameter/statistic pairs. This links the IPMs given throughout the text to the framework detailed in section 3 – and, more specifically, to Assumptions 4.2 and 4.3. Appendix A.1 discusses parameters and statistics, after which Appendix A.2 studies the IPM dynamics themselves.

A.1. Parameters and statistics. We start by proving that Assumption 4.3 holds for two parameter/statistic pairs used in most IPMs mentioned in this text: the (sample) mean and the (sample) covariance. We then discuss whether a similar result holds for self-normalized weighted variants.

LEMMA A.1. *When the parameter $\Theta[\cdot]$ is the expectation $\mathbb{E}[\cdot]$ and the sample statistic $\widehat{\Theta}(\cdot)$ is the sample mean $E(\cdot)$, Assumption 4.3 is satisfied.*

Proof. We first deal with Assumption 4.3(i):

$$(A.1) \qquad \|\widehat{\Theta}_J(\boldsymbol{X}_1) - \widehat{\Theta}_J(\boldsymbol{X}_2)\|_p = \left\|\frac{1}{J}\sum\nolimits_{j=1}^{J}(X_{1,j} - X_{2,j})\right\|_p \leq \|X_1 - X_2\|$$

for any $p \geq 2$. For (ii), we have

$$(A.2) \quad \|E(\boldsymbol{X}_1) - E(\boldsymbol{X}_2) - (\mathbb{E}[X_1] - \mathbb{E}[X_2])\|_p = \|E((\boldsymbol{X}_1 - \boldsymbol{X}_2) - \mathbb{E}[X_1 - X_2])\|_p.$$

Since a sample mean is an unbiased estimator, we use Properties 2.2 and 2.4 to obtain

$$(A.3) \qquad \|E((\boldsymbol{X}_1 - \boldsymbol{X}_2) - \mathbb{E}[X_1 - X_2])\|_p \leq 2c_p J^{-1/2}\|X_1 - X_2\|_p.$$

Then, (iii) follows from (ii) by setting $X_2 = \delta_0$, since $\mathbb{E}[\delta_0] = E(\boldsymbol{0}) = 0$ and $c_{\theta,2}$ is d-independent. Finally, (iv) is easy to check using Properties 2.2 and 2.4. $\square$

LEMMA A.2. *When the parameter $\Theta[\cdot]$ is the covariance $\mathbb{C}[\cdot]$ and the sample statistic $\widehat{\Theta}(\cdot)$ is the sample covariance $C(\cdot)$, Assumption 4.3 is satisfied.*

Proof. The proof of (i) is contained in that of [23, Lemma 3.9]; the proof of (ii) in that of [23, Lemma 3.8]. Then, (iii) and (iv) can be checked as for Lemma A.1. $\square$

A.2. Interacting-particle methods. Next, we discuss Assumption 4.2 for the interacting-particle methods in the paper. After giving Remarks A.3 and A.4, which are relevant for multiple IPMs, we discuss the methods one by one.

Remark A.3 (Covariance matrices). IPMs such as the EnKF and EKS use a sample covariance matrix as statistic. A mean-field covariance or a single-level sample covariance is guaranteed to be positive semi-definite; this is important, as it ensures operations such as matrix square roots and inverses (after addition with a positive definite matrix) are well-defined. A *multilevel* sample covariance matrix computed by (3.10), on the other hand, might have negative eigenvalues. In [9, 23], this is avoided by setting all negative eigenvalues to zero in a preprocessing step. We will follow this approach and, to this end, define the operator

$$(A.4) \qquad I^+(M) = \sum\nolimits_{\lambda_k \geq 0} \lambda_k q_k q_k^\top \qquad \text{with } \{(\lambda_k, q_k)\}_k \text{ the eigenpairs of } M,$$

which can be freely incorporated into existing dynamics where needed. Indeed, when applied to sample covariance matrices in the single-level algorithm, it is the identity operator. In the multilevel context, in ensures positive semi-definiteness.

In future work, it would be interesting to study multilevel covariance estimators that guarantee positive semi-definiteness directly, such as that of [33]. The advantage of the form (A.4) is that it builds on the naive multilevel estimation and thus is compatible with our framework's estimator (3.10) by updating only the dynamics.

Remark A.4 (Matrix square roots). IPMs such as EKS use the square root of the input covariance matrix θ. To prove Assumption 4.2, θ_1 in (4.2) then needs to be positive definite with its smallest eigenvalue bounded from below by some $\mu > 0$, since then Corollary A.6 applies. In sections 5 and 6, θ_1 is always the mean-field parameter $\Theta[\bar{X}_n]$; thus, an additional assumption for IPMs using the square root of the covariance matrix is that the mean-field ensemble does not collapse to zero.

THEOREM A.5 ([40, Corollary 4.2]). *If M_1, $M_2 \succeq 0$ and $M_1 + M_2 \succeq \mu I$ for some $\mu > 0$, then for any symmetric norm Φ, and $0 < m \leq 1$,*

$$(A.5) \qquad |M_1^m - M_2^m|_\Phi \leq (2/\mu)^{1-m}|M_1 - M_2|_\Phi.$$

COROLLARY A.6. *From Theorem A.5 combined with Property 2.3 follows immediately that, if M_1, $M_2 \geq 0$ and $M_1 \geq \mu I$ with $\mu > 0$, then for all $p \geq 2$,*

$$(A.6) \qquad \|\sqrt{M_1} - \sqrt{M_2}\|_p \leq \sqrt{2/\mu}\|M_1 - M_2\|_p.$$

We are now ready to tackle the individual IPMs.

LEMMA A.7. *The dynamics of the ensemble Kalman filter (EnKF; Example 1.1) and of ensemble Kalman inversion (EKI; Example 1.2) fit into the single-level framework (3.2) and satisfy Assumption 4.2.*

Proof. We tackle the ensemble Kalman filter first. In the notation (3.2) of our framework, we have $\widehat{\Theta}_J(\boldsymbol{g}) = C_J(\boldsymbol{g})$ and the EnKF uses the functions

$$(A.7) \qquad \Psi_n(u, g, \theta, \eta) = g - \theta H^\top (HI^+(\theta)H^\top + \Gamma)^{-1}(y_{n+1} - Hg + \sqrt{\Gamma}\eta),$$

with I^+ from Remark A.3. Define $K_i := \theta_i H^\top (HI^+(\theta_i)H^\top + \Gamma)^{-1}$; by Property 2.1,

$$(A.8) \quad \begin{aligned} &\|\Psi_n(U_1, G_1, \theta_1, \eta) - \Psi_n(U_2, G_2, \Theta_2, \eta)\|_p \\ &\leq \|G_1 - G_2\|_p + |K_1||H|\|G_1 - G_2\|_p + \|K_1 - K_2\|_{2p}\|y_{n+1} - HG_2 + \sqrt{\Gamma}\eta\|_{2p}. \end{aligned}$$

Now notice that $|K_1| \leq |\theta_1||H|/\gamma_{\min}$, where $\gamma_{\min} > 0$ is Γ's smallest eigenvalue. In addition, [23, Lemmas 3.3 and 3.4] together prove, in our notation, the bound $\|K_1 - K_2\|_{2p} \leq |H|/\gamma_{\min}(1 + 2|K_1||H|)\|\theta_1 - \Theta_2\|_{2p}$. Thus (A.8) is bounded by

$$\left(1 + \frac{|H|^2}{\gamma_{\min}}|\theta_1|\right)\|G_1 - G_2\|_p + \left(\frac{|H|}{\gamma_{\min}} + 2\frac{|H|^3}{\gamma_{\min}^2}|\theta_1|\right)\|y_{n+1} - HG_2 + \sqrt{\Gamma}\eta\|_{2p}\|\theta_1 - \Theta_2\|_{2p};$$

bounding factors with the locality conditions (e.g., $|\theta_1| \leq \|\Theta_0\| + d$) and Assumption 4.1(ii) concludes the proof for EnKF. The EKI dynamics (1.6) are an instance of the EnKF dynamics (1.3), with an enlarged state space [28]. The proof still applies. □

In the rest of this section, the use of Property 2.1 and the final step of using the locality conditions and Assumption 4.1(ii) will be left implicit.

LEMMA A.8. *The dynamics of the deterministic ensemble Kalman filter (DEnKF; Example 8.1) fit into the single-level framework (3.2) and satisfy Assumption 4.2.*

Proof. With $(\widehat{\Theta}_J^{(1)}(\boldsymbol{g}), \widehat{\Theta}_J^{(2)}(\boldsymbol{g})) = (E_J(\boldsymbol{g}), C_J(\boldsymbol{g}))$ in (3.2) follows from (8.1) that $\Psi_n(u, g, \theta, \eta) = g + K(H\theta^{(1)}/2 + y - 3Hg/2)$ with $K := \theta^{(2)}H^\top (HI^+(\theta^{(2)})H^\top + \Gamma)^{-1}$;

$$\begin{aligned} &\|\Psi_n(U_1, G_1, \theta_1, \eta) - \Psi_n(U_2, G_2, \Theta_2, \eta)\|_p \\ &= \left\|\left(G_1 + K_1\left(H\theta_1^{(1)}/2 + y - 3HG_1/2\right)\right) - \left(G_2 + K_2\left(H\Theta_2^{(1)}/2 + y - 3HG_2/2\right)\right)\right\|_p \\ &\leq \|G_1 - G_2\|_p + |K_1|(\|H(\theta_1^{(1)} - \Theta_2^{(1)})/2\|_p + \|3H(G_1 - G_2)/2\|_p) \\ &\quad + \|K_1 - K_2\|_{2p}\|H\Theta_2^{(1)}/2 + y - 3HG_2/2\|_{2p}. \end{aligned}$$

The proof can then be finished similarly to that of Lemma A.7. □

LEMMA A.9. *The dynamics of ensemble Kalman sampling (EKS; Example* 1.3*) fit into the single-level framework* (3.2) *and, if the mean-field particles do not collapse, satisfy Assumption* 4.2.

Proof. With $\widehat{\Theta}_J(\boldsymbol{u}, \boldsymbol{g}) = (C^{uu}(\boldsymbol{u}), C^{ug}(\boldsymbol{u}, \boldsymbol{g}))$, (1.7) can be manipulated into

$$(A.9) \quad \Psi_n(u, g, \theta, \eta) = (I + \tau_n \theta^{(1)} \Gamma_0^{-1})^{-1} \big[u + \tau_n \theta^{(2)} \Gamma^{-1}(y - g) \big] + \sqrt{2\tau_n I^+(\theta^{(1)})} \, \eta.$$

Then follows (since $\|\Theta_2^{(1)} - I^+(\Theta_2^{(1)})\|_p \le \|\Theta_2^{(1)} - \theta_1^{(1)}\|_p$ similarly to [23, Lemma 3.3]):

$$
\begin{aligned}
\|\Psi_n(U_1, G_1, \theta_1, \eta) &- \Psi_n(U_2, G_2, \Theta_2, \eta)\|_p \le \|\sqrt{2\tau_n}\,\eta\|_{2p} \|\sqrt{\theta_1^{(1)}} - \sqrt{I^+(\Theta_2^{(1)})}\|_{2p} \\
&+ \|(I + \tau_n \Theta_2^{(1)} \Gamma_0^{-1})^{-1}[U_1 - U_2 + \tau_n \theta_1^{(2)} \Gamma^{-1}(y - G_1) - \tau_n \Theta_2^{(2)} \Gamma^{-1}(y - G_2)]\|_p \\
&+ \|((I + \tau_n \theta_1^{(1)} \Gamma_0^{-1})^{-1} - (I + \tau_n \Theta_2^{(1)} \Gamma_0^{-1})^{-1})[U_1 + \tau_n \theta_1^{(2)} \Gamma^{-1}(y - G_1)]\|_p \\
\le 2&\sqrt{2\tau_n/\mu}\|\eta\|_{2p}\|\theta_1^{(1)} - \Theta_2^{(1)}\|_{2p} \\
&+ \|U_1 - U_2\|_p + \tau_n |\theta_1^{(2)} \Gamma^{-1}| \, \|G_1 - G_2\|_p + \tau_n |\Gamma^{-1}| \|(y - G_2)\|_{2p} \|\theta_1^{(2)} - \Theta_2^{(2)}\|_{2p} \\
&+ \|U_1 + \tau_n \theta_1^{(2)} \Gamma^{-1}(y - G_1)\|_{2p} \tau_n |\Gamma_0^{-1}| \|\theta_1^{(1)} - \Theta_2^{(1)}\|_{2p}.
\end{aligned}
$$

The last inequality used Remark A.4 for the first line, bounded $|(I + \tau_n \theta_2^1 \Gamma_0^{-1})^{-1}| \le 1$ for the second, and lastly used the Lipschitz inequality $|(I + A_1)^{-1} - (I + A_2)^{-1}| = |(I + A_1)^{-1}(A_2 - A_1)(I + A_2)^{-1}| \le |A_1 - A_2|$ for positive semi-definite matrices A_i. □

REFERENCES

[1] J. L. ANDERSON, *An Ensemble Adjustment Kalman Filter for Data Assimilation*, Monthly Weather Review, 129 (2001), pp. 2884–2903.

[2] K. BERGEMANN AND S. REICH, *An ensemble Kalman-Bucy filter for continuous data assimilation*, Meteorologische Zeitschrift, (2012), pp. 213–219.

[3] C. H. BISHOP, B. J. ETHERTON, AND S. J. MAJUMDAR, *Adaptive Sampling with the Ensemble Transform Kalman Filter. Part I: Theoretical Aspects*, Monthly Weather Review, 129 (2001), pp. 420–436.

[4] E. CALVELLO, S. REICH, AND A. M. STUART, *Ensemble Kalman Methods: A Mean Field Perspective*, Sept. 2022.

[5] J. A. CARRILLO, F. HOFFMANN, A. M. STUART, AND U. VAES, *Consensus-based sampling*, Stud. Appl. Math., 148 (2022), pp. 1069–1140.

[6] J. A. CARRILLO, S. JIN, L. LI, AND Y. ZHU, *A consensus-based global optimization method for high dimensional machine learning problems*, ESAIM Control Optim. Calc. Var., 27 (2021), p. S5.

[7] N. K. CHADA, A. JASRA, AND F. YU, *Multilevel Ensemble Kalman–Bucy Filters*, SIAM/ASA J. Uncertain. Quantif., 10 (2022), pp. 584–618.

[8] N. K. CHADA, A. M. STUART, AND X. T. TONG, *Tikhonov Regularization within Ensemble Kalman Inversion*, SIAM J. Numer. Anal., 58 (2020), pp. 1263–1294.

[9] A. CHERNOV, H. HOEL, K. J. H. LAW, F. NOBILE, AND R. TEMPONE, *Multilevel ensemble Kalman filtering for spatio-temporal processes*, Numer. Math., 147 (2021), pp. 71–125.

[10] P. DEL MORAL, *Nonlinear filtering: Interacting particle resolution*, Comptes Rendus de l'Académie des Sciences - Series I - Mathematics, 325 (1997), pp. 653–658.

[11] Z. DING AND Q. LI, *Ensemble Kalman inversion: Mean-field limit and convergence analysis*, Stat. Comput., 31 (2021), p. 9.

[12] Z. DING AND Q. LI, *Ensemble Kalman Sampler: Mean-field Limit and Convergence Analysis*, SIAM J. Math. Anal., 53 (2021), pp. 1546–1578.

[13] S. DUFFIELD AND S. S. SINGH, *Ensemble Kalman inversion for general likelihoods*, Statist. Probab. Lett., 187 (2022), p. 109523.

[14] O. R. A. Dunbar, A. B. Duncan, A. M. Stuart, and M.-T. Wolfram, *Ensemble Inference Methods for Models With Noisy and Expensive Likelihoods*, SIAM J. Appl. Dyn. Syst., 21 (2022), pp. 1539–1572.

[15] G. Evensen, *Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics*, J. Geophys. Res. Oceans, 99 (1994), pp. 10143–10162.

[16] K. Fossum, T. Mannseth, and A. S. Stordal, *Assessment of multilevel ensemble-based data assimilation for reservoir history matching*, Comput. Geosci., 24 (2020), pp. 217–239.

[17] A. Garbuno-Inigo, F. Hoffmann, W. Li, and A. M. Stuart, *Interacting Langevin Diffusions: Gradient Structure and Ensemble Kalman Sampler*, SIAM J. Appl. Dyn. Syst., 19 (2020), pp. 412–441.

[18] M. B. Giles, *Multilevel Monte Carlo Path Simulation*, Oper. Res., 56 (2008), pp. 607–617.

[19] C. Graham and D. Talay, *Stochastic Simulation and Monte Carlo Methods: Mathematical Foundations of Stochastic Simulation*, Springer, Berlin, Heidelberg, 2013.

[20] A. Gregory, C. J. Cotter, and S. Reich, *Multilevel Ensemble Transform Particle Filtering*, SIAM J. Sci. Comput., 38 (2016), pp. A1317–A1338.

[21] A. Gut, *Probability: A Graduate Course: A Graduate Course*, vol. 75 of Springer Texts in Statistics, Springer, New York, NY, 2013.

[22] A.-L. Haji-Ali and R. Tempone, *Multilevel and Multi-index Monte Carlo methods for the McKean–Vlasov equation*, Stat. Comput., 28 (2018), pp. 923–935.

[23] H. Hoel, K. J. H. Law, and R. Tempone, *Multilevel ensemble Kalman filtering*, SIAM J. Numer. Anal., 54 (2016), pp. 1813–1839.

[24] H. Hoel, G. Shaimerdenova, and R. Tempone, *Multilevel Ensemble Kalman Filtering based on a sample average of independent EnKF estimators*, Foundations of Data Science, 2 (2020), pp. 351–390.

[25] H. Hoel, G. Shaimerdenova, and R. Tempone, *Multi-index ensemble Kalman filtering*, J. Comput. Phys., 470 (2022), p. 111561.

[26] D. Z. Huang, T. Schneider, and A. M. Stuart, *Iterated Kalman methodology for inverse problems*, J. Comput. Phys., 463 (2022), p. 111262.

[27] M. A. Iglesias, *Iterative regularization for ensemble data assimilation in reservoir models*, Comput. Geosci., 19 (2015), pp. 177–212.

[28] M. A. Iglesias, K. J. H. Law, and A. M. Stuart, *Ensemble Kalman methods for inverse problems*, Inverse Problems, 29 (2013), p. 045001.

[29] A. Jasra, K. Kamatani, K. J. H. Law, and Y. Zhou, *Multilevel Particle Filters*, SIAM J. Numer. Anal., 55 (2017), pp. 3068–3096.

[30] N. B. Kovachki and A. M. Stuart, *Ensemble Kalman inversion: A derivative-free technique for machine learning tasks*, Inverse Problems, 35 (2019), p. 095005.

[31] F. Le Gland, V. Monbet, and V.-D. Tran, *Large sample asymptotics for the ensemble Kalman filter*, in The Oxford Handbook of Nonlinear Filtering, D. Crisan and B. Rozovskii, eds., Oxford University Press, 2011, pp. 598–631.

[32] J. Mandel, L. Cobb, and J. D. Beezley, *On the convergence of the ensemble Kalman filter*, Appl. Math., 56 (2011), pp. 533–541.

[33] A. Maurais, T. Alsup, B. Peherstorfer, and Y. Marzouk, *Multifidelity Covariance Estimation via Regression on the Manifold of Symmetric Positive Definite Matrices*, July 2023, https://arxiv.org/abs/2307.12438.

[34] R. Pinnau, C. Totzeck, O. Tse, and S. Martin, *A consensus-based model for global optimization and its mean-field limit*, Math. Models Methods Appl. Sci., 27 (2017), p. 183.

[35] A. A. Popov, C. Mou, A. Sandu, and T. Iliescu, *A Multifidelity Ensemble Kalman Filter with Reduced Order Control Variates*, SIAM J. Sci. Comput., 43 (2021), pp. A1134–A1162.

[36] L. F. Ricketson, *A multilevel Monte Carlo method for a class of McKean-Vlasov processes*, Aug. 2015, https://arxiv.org/abs/1508.02299.

[37] P. Sakov and P. R. Oke, *A deterministic formulation of the ensemble Kalman filter: An alternative to ensemble square root filters*, Tellus A: Dynamic Meteorology and Oceanography, 60 (2008), pp. 361–371.

[38] C. Schillings and A. M. Stuart, *Analysis of the Ensemble Kalman Filter for Inverse Problems*, SIAM J. Numer. Anal., 55 (2017), pp. 1264–1290.

[39] L. Szpruch, S. Tan, and A. Tse, *Iterative Multilevel Particle Approximation for Mckean–Vlasov Sdes*, Ann. Appl. Probab., 29 (2019), pp. 2230–2265.

[40] J. L. van Hemmen and T. Ando, *An inequality for trace ideals*, Commun. Math. Phys., 76 (1980), pp. 143–148.

[41] S. Weissmann, A. Wilson, and J. Zech, *Multilevel Optimization for Inverse Problems*, in Proceedings of Thirty Fifth Conference on Learning Theory, PMLR, 2022, pp. 5489–5524.