

Divide-and-Conquer Fusion

Ryan S.Y. Chan

*The Alan Turing Institute
London, NW1 2DB, UK*

RCHAN@TURING.AC.UK

Murray Pollock

*School of Mathematics, Statistics and Physics
Newcastle University
Newcastle, NE1 7RU, UK*

MURRAY.POLLOCK@NEWCASTLE.AC.UK

Adam M. Johansen

*Department of Statistics
University of Warwick
Coventry, CV4 7AL, UK*

A.M.JOHANSEN@WARWICK.AC.UK

Gareth O. Roberts

*Department of Statistics
University of Warwick
Coventry, CV4 7AL, UK*

GARETH.O.ROBERTS@WARWICK.AC.UK

Editor: Pierre Alquier

Abstract

Combining several (sample approximations of) distributions, which we term *sub-posteriors*, into a single distribution proportional to their product, is a common challenge. Occurring, for instance, in distributed ‘big data’ problems, or when working under multi-party privacy constraints. Many existing approaches resort to approximating the individual sub-posteriors for practical necessity, then find either an analytical approximation or sample approximation of the resulting (product-pooled) posterior. The quality of the posterior approximation for these approaches is poor when the sub-posteriors fall out-with a narrow range of distributional form, such as being approximately Gaussian. Recently, a *Fusion* approach has been proposed which finds an exact Monte Carlo approximation of the posterior, circumventing the drawbacks of approximate approaches. Unfortunately, existing Fusion approaches have a number of computational limitations, particularly when unifying a large number of sub-posteriors. In this paper, we generalise the theory underpinning existing Fusion approaches, and embed the resulting methodology within a recursive divide-and-conquer sequential Monte Carlo paradigm. This ultimately leads to a competitive Fusion approach, which is robust to increasing numbers of sub-posteriors.

Keywords: Distributed computing, importance sampling, Markov chain Monte Carlo, sequential Monte Carlo, stochastic differential equations.

1. Introduction

In this paper, we are interested in the following d -dimensional (*product-pooled*) target density (which we term the *fusion density*),

$$f(\mathbf{x}) \propto f_1(\mathbf{x}) \cdots f_C(\mathbf{x}) = \prod_{c=1}^C f_c(\mathbf{x}), \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^d$, $f_c(\mathbf{x})$ for $c \in \{1, \dots, C\}$ represent the individual densities which we wish to unify (termed *sub-posteriors* in deference to the fact that a major application of this technique will be the setting in which the posterior is proportional to the product of these factors), and C represents the total *number* of sub-posteriors. We assume that we have access to independent realisations from each sub-posterior, and that it is possible to evaluate each sub-posterior pointwise up to its normalising constant. Although typically, one would only have approximate samples from each sub-posterior, we will discuss later that neither of these assumptions form limiting factors for our methodology.

The need to unify several (sample approximations of) distributions, over a common parameter space, into a single sample approximation of the distribution in the manner of (1) is surprisingly common. For instance, it arises classically in expert elicitation (Albert et al., 2012; Berger, 1980; Genest and Zidek, 1986) and meta-analysis (Fleiss, 1993). However, it has proven to be challenging methodologically in a number of modern settings due to problem specific constraints. These include when dealing with the *privacy constraints* of the individual sources (Yıldırım and Ermiş, 2019), in cases where the sheer *number of sources* is overwhelming, or if the *networking constraints* of the sources are truly *distributed* (Scott et al., 2016). This in turn has motivated a range of problem specific and pragmatic *approximations*. These approximations are invariably distributional, and imposed at the level of the individual source (for instance, the sub-posteriors being approximately Gaussian). Such approximations limit the applicability of methodological approaches to particular settings, and outside those settings the unified results can be poorly understood, and even misleading. We instead focus on developing methodology for an *exact* Monte Carlo approximation of the unified distribution (1)—one which provides robust inference in a wide range of practical problems, and yet is amenable to use alongside any problem specific constraints.

The majority of the recent methodological developments for representing or sampling from (1) have been focused on tackling distributed ‘*big data*’ problems (see for instance Scott et al., 2016; Neiswanger et al., 2014; Wang and Dunson, 2013; Minsker et al., 2014; Srivastava et al., 2015; Nemeth and Sherlock, 2018). In this setting, due to its sheer size, the data is split across a number of cores (say C cores), inference is separately conducted on each core (often using MCMC), and then the respective methodologies attempt to unify the sample approximations of the distribution (as per (1), and typically using a convenient approximation). In this paper, we will compare our methodology with a number of the most popular approaches, and so will briefly describe these here. The *Consensus Monte Carlo (CMC)* approach of Scott et al. (2016) produces approximate samples from (1) by means of a weighted average of sub-posterior samples. It can be shown that CMC is *exact* when

each sub-posterior is Gaussian, and can be useful in settings where each sub-posterior is approximately Gaussian, which is often the case in big data settings (Walker, 1969; Johnson, 1970; Le Cam, 1986; Van der Vaart, 1998; Le Cam and Yang, 2000). However, it has been shown to exhibit large bias in other settings (Wang and Dunson, 2013). Neiswanger et al. (2014) suggest a strategy (which we term the *Kernel Density Estimate Monte Carlo (KDEMC)* approach) based on using a kernel density estimate to approximate the sub-posterior densities, and in effect approximating (1) by implicitly sampling from the product of non-parametric density estimates. Finally, the *Weierstrass sampler* of Wang and Dunson (2013) provides an alternative method for approximating (1) by means of using the product of Weierstrass transforms for each sub-posterior. Interestingly, we find empirically that for a cheap and crude approximation of (1) then the (simplest) CMC approach outperforms all other methodologies, but in cases where accuracy is a concern then our (more computationally expensive) *Fusion* approach should be used.

The *Fusion approach* (Dai et al., 2019, 2023) constructs a direct sample approximation of (1) itself, rather than seeking to obtain an adhoc approximation of f by combining approximations of the sub-posteriors. Underpinning the Fusion approach is the simple observation that if we sampled (independently) $\mathbf{X}^{(c)} \sim f_c$ for $c \in \{1, \dots, C\}$ then conditional on the event that $\mathbf{X}^{(1)} = \dots = \mathbf{X}^{(C)}$, we have that $\mathbf{X}^{(1)}$ has density f given in (1).

Clearly the difficulty with exploiting this observation is that we are conditioning on an event of probability 0. The *Monte Carlo Fusion (MCF)* approach of Dai et al. (2019) provides a framework for practically enforcing this conditioning. This is achieved by initialising C stochastic processes (independently from one another) using a single realisation from each sub-posterior (i.e. $\mathbf{X}_0^{(c)} \sim f_c$ for $c \in \{1, \dots, C\}$ where the subscript is a temporal index, noting that $\mathbf{X}_0^{(1)} \neq \dots \neq \mathbf{X}_0^{(C)}$), evolving the processes in such a manner that (i) these processes *coalesce* at some fixed future time (i.e. $\mathbf{X}_T^{(1)} = \dots = \mathbf{X}_T^{(C)}$); and (ii), the common marginal distribution at the coalescence time, T , is f . By repeating this approach multiple times, MCF provides multiple i.i.d. draws from f .

The *Bayesian Fusion (BF)* approach of Dai et al. (2023) re-examined the theoretical underpinnings of MCF by introducing a *stochastic differential equation (SDE)* describing the coalescence of the C stochastic processes, and exploited this theory together with methodology for *sequential Monte Carlo (SMC)* to *gradually* coalesce the stochastic processes. The resulting output of the BF approach is a number of *correlated* and *weighted* draws from f . BF is a far more practical and robust algorithm than MCF. A key advantage of BF over MCF is that it is possible to give considerable user guidance in its implementation.

Although BF provides significant improvements over MCF, the applicability of the methodology is still limited by factors including: (i) the numbers of sub-posteriors being combined; (ii) the level of sub-posterior correlation; (iii) the dimensionality of the sub-posteriors; (iv) the degree to which the sub-posteriors *conflict*; and (v) the computational cost of the approach even when the user-specified tuning parameters are optimally chosen. In this paper, we make two key contributions to address the limitations of MCF and BF: (i) we significantly improve upon the computational efficiency of BF by allowing the user to incorporate

global information about each sub-posterior within the SDE formulation, and unify *subsets* of the sub-posteriors at any one time—we term this approach *Generalised Bayesian Fusion (GBF)*, and present it in Section 2 and Algorithm 1; (ii) using the flexibility given by (i) in which sub-posteriors can be partially unified, we embed our GBF methodology within the *divide-and-conquer* paradigm of Lindsten et al. (2017), allowing the user to combine sub-posteriors in *stages* to recover the fusion density f . We term this *Divide-and-Conquer Fusion (D&C-Fusion)*, and present it in Section 3 and Algorithm 2.

The remainder of the paper is organised as follows: In Section 4 we present detailed guidance on implementing our GBF and D&C-Fusion approaches, and in particular choosing any tuning parameters. In Section 5 we present applications of our methodology for a variety of models, comparing them to competing approximate methodologies. We conclude by outlining a variety of ways our Fusion approach could be extended, and used in other application settings. All technical proofs and detailed calculations are collated in the appendices.

Statistical computations for this paper were written in **R** (R Core Team, 2022), C++ and **Rcpp** (Eddelbuettel, 2013). All code developed for producing numerical results can be found on GitHub at <https://github.com/rchan26/DCFusion>.

2. A generalisation of the Fusion approach

In this section we develop theory and methodology to generalise and improve upon the BF approach of Dai et al. (2023), by incorporating information about the covariance of the sub-posteriors within the SDE formulation. For completeness in Appendix A we more fully outline the connections of our methodology to the earlier MCF and BF works, highlighting explicitly the advantages of our approach, but for ease of presentation here we instead present our approach directly. In this section we also consider the more abstract problem of sampling from the density $f^{(\mathcal{C})} \propto \prod_{c \in \mathcal{C}} f_c$, where \mathcal{C} is an index set representing the sub-posteriors we want to unify, and we assume we can sample (independently) $\mathbf{X}^{(c)} \sim f_c$ for $c \in \mathcal{C}$. This abstraction is useful for the methodology we develop in Section 3.

For the purposes of simplifying the subsequent notation, we denote by $\vec{\mathbf{x}}_t^{(\mathcal{C})} \in \mathbb{R}^{|\mathcal{C}| \times d}$ a vector composed of $\mathbf{x}_t^{(c)} \in \mathbb{R}^d$ for $c \in \mathcal{C}$ (in particular, we have $\vec{\mathbf{x}}_t^{(\mathcal{C})} := (\mathbf{x}_t^{(c_1)}, \dots, \mathbf{x}_t^{(c_{|\mathcal{C}|})})$, with c_i denoting the i^{th} element of the index set \mathcal{C}). We further assume that for $c \in \{1, \dots, C\}$, f_c is nowhere zero and everywhere differentiable, and that we can compute $A_c(\mathbf{x}) := \log f_c(\mathbf{x})$, $\nabla A_c(\mathbf{x})$, and $\nabla^2 A_c(\mathbf{x})$ pointwise (where ∇ is the gradient operator and ∇^2 is the Hessian). A fuller discussion of these assumptions is given in Appendix A, but note that they match those of the earlier works of Dai et al. (2019, 2023).

We begin by describing the joint distribution of $|\mathcal{C}|$ coalescing stochastic processes on $[0, T]$ that at time T have the common marginal $f^{(\mathcal{C})} \propto \prod_{c \in \mathcal{C}} f_c$. We term this the *fusion measure*, \mathbb{F} . To aid in the development of the subsequent methodology, we require that the stochastic processes can be simulated, and so this is done by considering a Radon-Nikodým correction of the so-called *proposal measure* (\mathbb{P}), which is defined to be the probability law induced

by $|\mathcal{C}|$ interacting d -dimensional parallel continuous-time Markov processes in $[0, T]$ where each process is given by the SDE,

$$d\mathbf{X}_t^{(c)} = \frac{\tilde{\mathbf{X}}_t - \mathbf{X}_t^{(c)}}{T-t} dt + \mathbf{\Lambda}_c^{\frac{1}{2}} d\mathbf{W}_t^{(c)}, \quad \mathbf{X}_0^{(c)} := \mathbf{x}_0^{(c)} \sim f_c, \quad t \in [0, T], \quad (2)$$

where $\mathbf{\Lambda}_c$ are (positive semi-definite) user-specified matrices associated to sub-posterior f_c for $c \in \mathcal{C}$ with $\mathbf{\Lambda}_c^{1/2}$ being the (positive semi-definite) square root of $\mathbf{\Lambda}_c$ where $\mathbf{\Lambda}_c^{1/2} \mathbf{\Lambda}_c^{1/2} = \mathbf{\Lambda}_c$. Note that for the purposes of our numerical simulations later we use the Schur decomposition. Furthermore, $\{\mathbf{W}_t^{(c)}\}_{c \in \mathcal{C}}$ denotes independent Brownian motions, and

$$\tilde{\mathbf{X}}_t^{(c)} := \left(\sum_{c \in \mathcal{C}} \mathbf{\Lambda}_c^{-1} \right)^{-1} \left(\sum_{c \in \mathcal{C}} \mathbf{\Lambda}_c^{-1} \mathbf{X}_t^{(c)} \right),$$

denoting the weighted average of the processes at time t . In practice we typically take $\mathbf{\Lambda}_c$ to be a user estimate of the covariance matrix of the sub-posterior, $\hat{\Sigma}_c$ which can be computed using the available sub-posterior samples for f_c thereby incorporating problem-specific information about covariance structure. We will see that the choice for these matrices influences the efficiency of the algorithm but not the target distribution itself and thus incurs no bias. Realisations of the proposal measure are denoted as $\mathfrak{X} := \{\tilde{\mathbf{x}}_t^{(c)}, t \in [0, T]\}$. For the purposes of exposition, we defer discussion on the practical simulation of \mathbb{P} to Section 2.1.

Now, we let the *Fusion measure* \mathbb{F} be simply the measure induced by the following Radon-Nikodým derivative:

$$\frac{d\mathbb{F}}{d\mathbb{P}}(\mathfrak{X}) \propto \rho_0(\tilde{\mathbf{x}}_0^{(c)}) \cdot \prod_{c \in \mathcal{C}} \left[\exp \left\{ - \int_0^T \phi_c(\mathbf{X}_t^{(c)}) dt \right\} \right], \quad (3)$$

where $\{\mathbf{X}_t^{(c)}, t \in [0, T]\}$ is a Brownian bridge from $\mathbf{X}_0^{(c)} := \mathbf{x}_0^{(c)} \sim f_c$ to $\mathbf{X}_T^{(c)} := \mathbf{x}_T^{(c)}$ with covariance matrix $\mathbf{\Lambda}_c$ and

$$\rho_0(\tilde{\mathbf{x}}_0^{(c)}) := \exp \left\{ - \sum_{c \in \mathcal{C}} \frac{(\tilde{\mathbf{x}}_0^{(c)} - \mathbf{x}_0^{(c)})^\top \mathbf{\Lambda}_c^{-1} (\tilde{\mathbf{x}}_0^{(c)} - \mathbf{x}_0^{(c)})}{2T} \right\}, \quad (4)$$

where

$$\tilde{\mathbf{x}}_t^{(c)} := \left(\sum_{c \in \mathcal{C}} \mathbf{\Lambda}_c^{-1} \right)^{-1} \left(\sum_{c \in \mathcal{C}} \mathbf{\Lambda}_c^{-1} \mathbf{x}_t^{(c)} \right), \quad (5)$$

and

$$\phi_c(\mathbf{x}) := \frac{1}{2} \left(\nabla \log f_c(\mathbf{x})^\top \mathbf{\Lambda}_c \nabla \log f_c(\mathbf{x}) + \text{Tr}(\mathbf{\Lambda}_c \nabla^2 \log f_c(\mathbf{x})) \right). \quad (6)$$

Now, considering the time T marginal of $\mathfrak{X} \sim \mathbb{F}$ we (almost surely) have:

Theorem 1 *Under the fusion measure \mathbb{F} , the ending points of the $|\mathcal{C}|$ interacting, parallel processes have a common value at time T , $\mathbf{y}^{(c)}$ which has density $f^{(c)}$ and $\mathbf{y}^{(c)} = \mathbf{x}_T^{(c_1)} = \dots = \mathbf{x}_T^{(c_{|\mathcal{C}|})}$ almost surely.*

Proof See Appendix B. ■

Theorem 1 suggests that we can simulate from the fusion target density $f^{(C)}$ by simulating $\mathfrak{X} \sim \mathbb{F}$ and retaining the T time marginal, $\mathbf{y}^{(C)}$. As suggested by the theory, we do so by means of simulating a number of proposals $\mathfrak{X} \sim \mathbb{P}$ and accepting (or importance weighting) the terminal time marginal $\mathbf{y}^{(C)}$ with probability proportional to the Radon-Nikodým derivative in (3). As such, we need to consider: (i), how to simulate proposals from $\mathfrak{X} \sim \mathbb{P}$ (outlined in Section 2.1); and (ii), how to compute the Radon-Nikodým correction (3) (outlined in Section 2.2). We then present our proposed complete methodology in Section 2.3. We discuss possible extensions of our approach in Section 2.4.

2.1 Simulating from the Proposal Measure

First, we consider how to simulate proposals from $\mathfrak{X} \sim \mathbb{P}$. We begin by noting that the initialisation of the proposal measure given by (2) at time $t = 0$ only requires independent draws from the $|\mathcal{C}|$ sub-posteriors that we wish to unify, which in this paper we assume we have access to. If independent sampling is not feasible, it is possible to obtain approximate sub-posterior samples using MCMC (see Dai et al. (2023, Section 3.6) for a discussion on the impacts of using approximate sub-posterior samples for Fusion). Further, although paths $\mathfrak{X} \sim \mathbb{P}$ are infinite dimensional random variables (and so we cannot draw entire sample paths from \mathbb{P}), it is sufficient for our needs to simulate (exactly) the paths at a finite collection of times provided we can ensure that we are able to simulate the path (*exactly*) at time T . For clarity, we only consider simulating \mathfrak{X} at times given by the following auxiliary temporal partition,

$$\mathcal{P} = \{t_0, t_1, \dots, t_n : 0 =: t_0 < t_1 < \dots < t_n := T\}. \quad (7)$$

We let $\Delta_j := t_j - t_{j-1}$ and for notational simplicity, subscripts are suppressed when considering the processes at times given in the temporal partition. In particular, let $\mathbf{x}_j^{(c)}$ denote $\mathbf{x}_{t_j}^{(c)}$, and let $\vec{\mathbf{x}}_j^{(C)}$ denote $\vec{\mathbf{x}}_{t_j}^{(C)}$. We will see from the following proposition, that algorithmically, to simulate from \mathbb{P} at the time points in \mathcal{P} , we can simply initialise the $|\mathcal{C}|$ paths with $\mathbf{x}_0^{(c)} \sim f_c$ for $c \in \mathcal{C}$ and sequentially simulate from the Normal distributions given in Proposition 2a (for $j \in \{1, \dots, n-1\}$) and b (for $j = n$). The following proposition tells us how to simulate from the transition density of \mathbb{P} :

Proposition 2 Let $\mathcal{C} := (c_1, \dots, c_{|\mathcal{C}|})$ denote the index set representing the sub-posteriors we wish to unify, then if \mathfrak{X} satisfies (2), then under the proposal measure, \mathbb{P} , we have

(a) For $s < t < T$,

$$\vec{\mathbf{X}}_t^{(C)} \Big| \left(\vec{\mathbf{X}}_s^{(C)} = \vec{\mathbf{x}}_s^{(C)} \right) \sim \mathcal{N}_{|\mathcal{C}|d} \left(\vec{\mathbf{M}}_{s,t}^{(C)}, \mathbf{V}_{s,t} \right), \quad (8)$$

where $\vec{M}_{s,t}^{(c)} \in \mathbb{R}^{|\mathcal{C}| \times d} := \left(M_{s,t}^{(c_1)}, \dots, M_{s,t}^{(c_{|\mathcal{C}|})} \right)$ with

$$M_{s,t}^{(c)} = \frac{T-t}{T-s} \mathbf{x}_s^{(c)} + \frac{t-s}{T-s} \tilde{\mathbf{x}}_s, \quad (9)$$

and

$$\mathbf{V}_{s,t} = \begin{pmatrix} \Gamma_{11} & \Gamma_{12} & \dots & \Gamma_{1|\mathcal{C}|} \\ \Gamma_{21} & \Gamma_{22} & \dots & \Gamma_{2|\mathcal{C}|} \\ \vdots & \vdots & \ddots & \vdots \\ \Gamma_{|\mathcal{C}|1} & \Gamma_{|\mathcal{C}|2} & \dots & \Gamma_{|\mathcal{C}||\mathcal{C}|} \end{pmatrix} \in \mathbb{R}^{|\mathcal{C}|d \times |\mathcal{C}|d}, \quad (10)$$

where for $i, j = 1, \dots, |\mathcal{C}|$,

$$\Gamma_{ii} = \frac{(t-s)(T-t)}{T-s} \Lambda_{c_i} + \frac{(t-s)^2}{T-s} \Lambda_{\mathcal{C}} \in \mathbb{R}^{d \times d}, \quad (11)$$

$$\Gamma_{ij} = \frac{(t-s)^2}{T-s} \Lambda_{\mathcal{C}} \in \mathbb{R}^{d \times d}. \quad (12)$$

(b) For $s < t = T$, $\mathbf{y}^{(c)} := \mathbf{x}_T^{(c_1)} = \dots = \mathbf{x}_T^{(c_{|\mathcal{C}|})} \sim \mathcal{N}_d(\tilde{\mathbf{x}}_s, \Lambda_{\mathcal{C}})$.

(c) For each $c \in \mathcal{C}$, the distribution of $\{\mathbf{X}_q^{(c)}, s \leq q \leq t\}$ given endpoints $\mathbf{X}_s^{(c)} = \mathbf{x}_s^{(c)}$ and $\mathbf{X}_t^{(c)} = \mathbf{x}_t^{(c)}$ is a Brownian bridge with covariance matrix $\Lambda_{\mathcal{C}}$, so

$$\mathbf{X}_u^{(c)} \mid \left(\mathbf{x}_s^{(c)}, \mathbf{x}_t^{(c)} \right) \sim \mathcal{N}_d \left(\frac{(t-q)\mathbf{x}_s^{(c)} + (q-s)\mathbf{x}_t^{(c)}}{t-s}, \frac{(t-q)(q-s)}{t-s} \Lambda_{\mathcal{C}} \right). \quad (13)$$

Proof See Appendix C. ■

As we can initialise a draw from \mathbb{P} , and from Proposition 2 we can simulate from its transition density, we can now explicitly express the $d(n|\mathcal{C}| + 1)$ -dimensional density of the $|\mathcal{C}|d$ -dimensional Markov process at the $(n + 1)$ time marginals given by the temporal partition under \mathbb{P} , by iterative simulation from the transition density:

$$h_{\mathcal{C}} \left(\vec{\mathbf{x}}_0^{(c)}, \dots, \vec{\mathbf{x}}_{n-1}^{(c)}, \mathbf{y}^{(c)} \right) \propto f \left(\vec{\mathbf{x}}_0^{(c)} \right) \cdot \prod_{j=1}^{n-1} \mathcal{N}_{|\mathcal{C}|d} \left(\vec{\mathbf{x}}_j^{(c)} \mid \vec{M}_j^{(c)}, \mathbf{V}_j \right) \cdot \mathcal{N}_d \left(\mathbf{y}^{(c)} \mid \vec{\mathbf{x}}_{n-1}^{(c)}, \Lambda_{\mathcal{C}} \right), \quad (14)$$

where $f \left(\vec{\mathbf{x}}_0^{(c)} \right) \propto \prod_{c \in \mathcal{C}} f_c \left(\mathbf{x}_0^{(c)} \right)$, and $\mathcal{N}_d(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the density of a d -dimensional Normal distribution (evaluated at \mathbf{x}) with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. For notational convenience we let $\vec{M}_j^{(c)} = \vec{M}_{t_{j-1}, t_j}^{(c)}$ and $\mathbf{V}_j = \mathbf{V}_{t_{j-1}, t_j}$.

2.2 Radon-Nikodým correction of the Proposal

Now, we direct our consideration to the second step: computing the Radon-Nikodým correction of (3), given we have drawn our proposal from \mathbb{P} restricted to the times given by the

partition \mathcal{P} . Factorising the Radon-Nikodým derivative in (3) according to the temporal partition \mathcal{P} , the $d(n|\mathcal{C}| + 1)$ -dimensional density under \mathbb{F} is

$$g_{\mathcal{C}}\left(\vec{\mathbf{x}}_0^{(c)}, \dots, \vec{\mathbf{x}}_{n-1}^{(c)}, \mathbf{y}^{(c)}\right) \propto h_{\mathcal{C}}\left(\vec{\mathbf{x}}_0^{(c)}, \dots, \vec{\mathbf{x}}_{n-1}^{(c)}, \mathbf{y}^{(c)}\right) \cdot \prod_{j=0}^n \rho_j, \quad (15)$$

where ρ_0 is given in (4) and for $j \in \{1, \dots, n\}$,

$$\rho_j\left(\vec{\mathbf{x}}_{j-1}^{(c)}, \vec{\mathbf{x}}_j^{(c)}\right) = \prod_{c \in \mathcal{C}} \mathbb{E}_{\mathbb{W}_{\Lambda_c, j}} \left[\exp \left\{ - \int_{t_{j-1}}^{t_j} \left(\phi_c\left(\mathbf{X}_t^{(c)}\right) - \Phi_c \right) \right\} \right] \in (0, 1], \quad (16)$$

and where $\mathbb{W}_{\Lambda_c, j}$ is the law of a Brownian bridge $\{\mathbf{X}_t^{(c)}, t \in (t_{j-1}, t_j)\}$ from $\mathbf{X}_{t_{j-1}} := \mathbf{x}_{j-1}^{(c)}$ to $\mathbf{X}_{t_j} := \mathbf{x}_j^{(c)}$ with covariance Λ_c , and $\Phi_c < \infty$ is a constant such that $\phi_c(\mathbf{x}) \geq \Phi_c$ for all \mathbf{x} and each $c \in \mathcal{C}$. We note that the terms Φ_c for $c \in \mathcal{C}$ (the global lower bounds of the respective ϕ_c for $c \in \mathcal{C}$) in (16) can be absorbed into normalising constants and, hence, as we apply sequential Monte Carlo methodology, they need not be evaluated (as shown more explicitly in Section 2.3).

Whilst ρ_0 (given by (4)) can be computed easily, direct computation of ρ_j in (16) for $j \in \{1, \dots, n\}$ is not possible as it requires evaluation of path integrals of Brownian motion. However, it is possible to construct non-negative unbiased estimators for (16) (with finite variance and computable in finite cost) in a similar fashion to Beskos et al. (2008); Fearnhead et al. (2008); Dai et al. (2019, 2023). To do so, we require for a given sample path $\mathbf{X}_{[t_{j-1}, t_j]}^{(c)} \sim \mathbb{W}_{\Lambda_c, j}$ that we have upper and lower bounds for $\phi_c(\mathbf{X}_t^{(c)})$ for each $c \in \mathcal{C}$. In general, it is not possible to find global bounds for ϕ_c , so we follow the approach of Beskos et al. (2008) and Pollock et al. (2016) who noted that if we can bound a sample path $\mathbf{X}_{[t_{j-1}, t_j]}^{(c)} \sim \mathbb{W}_{\Lambda_c, j}$, then conditional on these *layers* (or bounds) of the sample path, then we will be able to find *local* upper and lower bounds of ϕ_c denoted $U_j^{(c)}$ and $L_j^{(c)}$, respectively, such that $\phi_c(\mathbf{X}_t^{(c)}) \in [L_j^{(c)}, U_j^{(c)}]$ for $t \in [t_{j-1}, t_j]$. In order to practically implement this, we need to simulate Brownian bridges jointly with a compact region which almost surely constrains their path (a mechanism for doing this is described in Pollock et al. (2016, Sections 7 and 8)). We now describe one approach for doing this.

To achieve this, let $R_c := R_c(\mathbf{X}_{[t_{j-1}, t_j]})$ denote the *layer information* (i.e the compact region in which $\mathbf{X}_t^{(c)}$ is constrained in time $[t_{j-1}, t_j]$). We note that it is possible to partition the sample space into disjoint sets and simulate from associated distribution function (without having to sample the underlying path), $R_c \sim \mathcal{R}_c$. If $\Lambda_c = \mathbb{I}_d$, then we can simulate a layer to which $\mathbf{X}_t^{(c)} \in R_c$ for $t \in [t_{j-1}, t_j]$ by using algorithms outlined in Pollock et al. (2016, Sections 7 and 8) (for instance Pollock et al. (2016, Algorithm 14)). In the case where $\Lambda_c \neq \mathbb{I}_d$, we can still simulate R_c by appealing to a suitable transformation (which we detail fully in Appendix F and in Algorithm 5). Furthermore, once we have simulated layer information for $\mathbf{X}_t^{(c)}$ for $t \in [t_{j-1}, t_j]$, we can simulate the path at any required time marginals conditional on the simulated layer, $\mathbf{X}_t^{(c)} \sim \mathbb{W}_{\Lambda_c, j} | R_c$ (via a transformation and applying for instance Pollock et al. (2016, Algorithm 15)).

Although it is possible to find tight local bounds for ϕ_c in a problem specific manner by exploiting specific structure, there are some generic strategies that can be followed. In sufficiently regular settings one might construct the partition necessary by first partitioning the domain of ϕ_c and then looking at the pre-image of that partition under ϕ_c , thereby reducing the problem to a univariate one. Alternatively, it is helpful in practice to note that it is possible to find generic (less tight) bounds given by the following proposition:

Proposition 3 For all $c \in \mathcal{C}$ and $\mathbf{x} \in R_c$, we have $\phi_c(\mathbf{x}) \in [L_j^{(c)}, U_j^{(c)}]$, where

$$L_j^{(c)} := -\frac{1}{2} (d \cdot P^{\Lambda_c}), \quad (17)$$

$$U_j^{(c)} := \frac{1}{2} \left[\left(\left\| \Lambda_c^{\frac{1}{2}} \nabla \log f_c(\hat{\mathbf{x}}^{(c)}) \right\| + \max_{\mathbf{x} \in R_c} \left\| \Lambda_c^{-\frac{1}{2}} (\mathbf{x} - \hat{\mathbf{x}}^{(c)}) \right\| \cdot P^{\Lambda_c} \right)^2 + d \cdot P^{\Lambda_c} \right], \quad (18)$$

where d denotes the dimension of \mathbf{x} , $\|\cdot\|$ is the Euclidean norm, $\hat{\mathbf{x}}^{(c)}$ is a user-specified point central to R_c , and where P^{Λ_c} is a quantity such that

$$P^{\Lambda_c} \geq \max_{\mathbf{x} \in R_c} \gamma(\Lambda_c \nabla^2 \log f_c(\mathbf{x})), \quad (19)$$

with γ denoting the matrix norm, defined as

$$\gamma(A) := \max_{\|\mathbf{x}\| \neq 0} \frac{\|A\mathbf{x}\|}{\|\mathbf{x}\|}. \quad (20)$$

Proof See Appendix D. ■

Once local bounds for ϕ_c are obtained, we can unbiasedly estimate ρ_j (16) for $j \in \{1, \dots, n\}$ by letting $\Delta_j := t_j - t_{j-1}$ and computing $a_j \tilde{\rho}_j$, where $a_j := \exp(\sum_{c \in \mathcal{C}} \Phi_c \Delta_j)$ and

$$\tilde{\rho}_j(\vec{\mathbf{x}}_{j-1}^{(C)}, \vec{\mathbf{x}}_j^{(C)}) := \prod_{c \in \mathcal{C}} \left(\frac{\Delta_j^{\kappa_c} \cdot e^{-U_j^{(c)} \Delta_j}}{\kappa_c! \cdot p(\kappa_c | R_c)} \cdot \prod_{k_c=1}^{\kappa_c} [U_j^{(c)} - \phi_c(\mathbf{X}_{\xi_{c,k_c}}^{(c)})] \right), \quad (21)$$

where R_c is the simulated layer information for the Brownian bridge sample path $\mathbf{X}_t^{(c)} \sim \mathbb{W}_{\Lambda_{c,j}}$ from $\mathbf{x}_{j-1}^{(c)}$ to $\mathbf{x}_j^{(c)}$, $L_j^{(c)}$ and $U_j^{(c)}$ are constants such that $L_j^{(c)} \leq \phi(\mathbf{X}_t^{(c)}) \leq U_j^{(c)}$ for all $\mathbf{X}_t^{(c)} \sim \mathbb{W}_{\Lambda_{c,j}} | R_c$, κ_c is a discrete random variable with conditional probabilities $\mathbb{P}[\kappa_c = k_c | R_c] := p(\kappa_c | R_c)$ (which at this stage we allow to be arbitrary) and $\xi_{c,1}, \dots, \xi_{c,\kappa_c} \stackrel{\text{iid}}{\sim} \mathcal{U}[t_{j-1}, t_j]$ for all $c \in \mathcal{C}$.

Theorem 4 Let $a_j := \exp(\sum_{c=1}^C \Phi_c \Delta_j)$, then for every $j = 1, \dots, n$, $a_j \tilde{\rho}_j$ is an unbiased estimator of ρ_j . In particular, we have

$$\rho_j = \mathbb{E} \left[\mathbb{E} \left[\mathbb{E} \left[\mathbb{E} \left[a_j \tilde{\rho}_j \mid \{\mathcal{R}_c, \mathbf{X}_{[t_{j-1}, t_j]}^{(c)}, \kappa_c\}_{c \in \mathcal{C}} \right] \mid \{\mathcal{R}_c, \mathbf{X}_{[t_{j-1}, t_j]}^{(c)}\}_{c \in \mathcal{C}}, \right] \mid \{\mathcal{R}_c\}_{c \in \mathcal{C}} \right] \right]$$

$$= \mathbb{E}_{\mathcal{R}} \mathbb{E}_{\mathbb{W}|\mathcal{R}} \mathbb{E}_{\mathbb{K}} \mathbb{E}_{\bar{\mathbb{U}}} [a_j \tilde{\rho}_j], \quad (22)$$

where (for readability) the expectation subscript denotes the law with which they are taken. Here, \mathcal{R} denotes the law of $\{R_c \sim \mathcal{R}_c : c = 1, \dots, C\}$, \mathbb{W} denotes the law of the C Brownian bridges $\{\mathbb{W}_{\Lambda_c, j} : c = 1, \dots, C\}$, \mathbb{K} denotes the law of $\{\kappa_c : c = 1, \dots, C\}$ and $\bar{\mathbb{U}}$ denotes the law of $\{\xi_{c,1}, \dots, \xi_{c,\kappa_c} : c = 1, \dots, C\} \stackrel{iid}{\sim} \mathcal{U}[t_{j-1}, t_j]$.

Proof See Appendix E. ■

We note that this unbiased estimator for ρ_j allows for significant flexibility in choosing the law \mathbb{K} . Following the discussion in Dai et al. (2023, Appendix B), there are two natural choices of unbiased estimators that could be used by making particular choices for the distribution of the discrete random variable used to simulate κ_c for $c \in \mathcal{C}$. We denote these $\tilde{\rho}_j^{(a)}$ and $\tilde{\rho}_j^{(b)}$ and are based, respectively, upon the GPE-1 and GPE-2 estimators of Fearnhead et al. (2008):

Definition 5 (GPE-1 for ρ_j (16)): Choosing the law of $\kappa_c \sim \text{Poi}((U_j^{(c)} - L_j^{(c)})\Delta_j)$ for $c \in \mathcal{C}$ leads to the following estimator:

$$\tilde{\rho}_j^{(a)}(\vec{\mathbf{x}}_{j-1}^{(C)}, \vec{\mathbf{x}}_j^{(C)}) := \prod_{c \in \mathcal{C}} \left(e^{-L_j^{(c)} \Delta_j} \cdot \prod_{\kappa_c=1}^{\kappa_c} \left[\frac{U_j^{(c)} - \phi_c(\mathbf{X}_{\xi_{c,\kappa_c}}^{(c)})}{U_j^{(c)} - L_j^{(c)}} \right] \right), \quad (23)$$

where $\exp\{\sum_{c=1}^C \Phi_c \Delta_j\} \cdot \tilde{\rho}_j^{(a)}$ is an unbiased estimator for ρ_j .

Definition 6 (GPE-2 for ρ_j (16)): Choosing the law of $\kappa_c \sim \text{NB}(\gamma_c, \beta_c)$ for $c \in \mathcal{C}$ with

$$\gamma_c := U_j^{(c)} \Delta_j - \int_{t_{j-1}}^{t_j} \phi_c \left(\mathbf{x}_{j-1}^{(c)} \cdot \frac{t_j - s}{\Delta_j} + \mathbf{x}_j^{(c)} \cdot \frac{s - t_{j-1}}{\Delta_j} \right) ds, \quad (24)$$

leads to the following estimator:

$$\tilde{\rho}_j^{(b)}(\vec{\mathbf{x}}_{j-1}^{(C)}, \vec{\mathbf{x}}_j^{(C)}) := \prod_{c \in \mathcal{C}} \left(e^{-U_j^{(c)} \Delta_j} \cdot \frac{\Delta_j^{\kappa_c} \cdot \Gamma(\beta_c) \cdot (\beta_c + \gamma_c)^{\beta_c + \kappa_c}}{\Gamma(\beta_c + \kappa_c) \beta_c^{\beta_c} \gamma_c^{\kappa_c}} \cdot \prod_{\kappa_c=1}^{\kappa_c} [U_j^{(c)} - \phi_c(\mathbf{X}_{\xi_{c,\kappa_c}}^{(c)})] \right), \quad (25)$$

where $\exp\{\sum_{c=1}^C \Phi_c \Delta_j\} \cdot \tilde{\rho}_j^{(b)}$ is an unbiased estimator for ρ_j .

The estimators $\tilde{\rho}_j^{(a)}$ and $\tilde{\rho}_j^{(b)}$ can be computed as detailed in Appendix F, and by means of Algorithm 5, by appealing to Dai et al. (2023, Algorithm 4 and Appendix. B). $\tilde{\rho}_j^{(a)}$ and $\tilde{\rho}_j^{(b)}$ have particularly desirable properties (by choosing $L_j^{(c)}$ and $U_j^{(c)}$ as in Proposition 3):

Proposition 7 *Let $a_j := \exp\{\sum_{c=1}^C \Phi_c \Delta_j\}$, then $a_j \tilde{\rho}_j^{(a)}(\vec{\mathbf{x}}^{(c)}, \mathbf{y}^{(c)})$ and $a_j \tilde{\rho}_j^{(b)}(\vec{\mathbf{x}}^{(c)}, \mathbf{y}^{(c)})$ are unbiased estimators of $\rho_j(\vec{\mathbf{x}}^{(c)}, \mathbf{y}^{(c)})$ which are positive with finite variance. In addition, $\tilde{\rho}_j^{(a)}(\vec{\mathbf{x}}^{(c)}, \mathbf{y}^{(c)}) \in [0, 1]$.*

Proof *See Fearnhead et al. (2008).* ■

As we discuss in Section 2.3, the critical consideration when choosing the law \mathbb{K} is to minimise the variance of the estimator. In our subsequent simulations, we will typically choose the GPE-2 estimator in Condition 6 as it has been empirically shown to have superior performance in Fearnhead et al. (2008, Section 5) and Dai et al. (2023, Section 3.5). Note that the mean run time for both the estimators $\tilde{\rho}_j^{(a)}$ and $\tilde{\rho}_j^{(b)}$ will be random, but will be finite and proportional to κ_c for a given layer $R_c \sim \mathcal{R}_c$.

2.3 Methodology

As we outlined earlier in this section, Theorem 1 suggests that we can simulate from the fusion target density $f^{(c)}$ by simulating $\mathfrak{X} \sim \mathbb{F}$ and retaining the T time marginal, $\mathbf{y}^{(c)}$. This can be achieved by simulating a number of proposals $\mathfrak{X} \sim \mathbb{P}$ and accepting (or importance weighting) the terminal time marginal $\mathbf{y}^{(c)}$ with probability proportional to the Radon-Nikodým derivative in (3). We are now able to implement each of these steps (as discussed in Sections 2.1 and 2.2 respectively), but we have considerable freedom over the details of the methodological approach.

The simplest approach is a *rejection sampler*: simulate a proposal from h_c (14) (by utilising Proposition 2), accept this proposal with probability $\prod_{j=0}^n \rho_j$, and conditional on acceptance return $\mathbf{y}^{(c)}$. As more fully discussed in Appendix A, this coincides methodologically with MCF if we set $\mathbf{\Lambda}_c = \mathbb{I}_d$ for $c \in \{1, \dots, C\}$ (although the formulation is different). The benefit of such a rejection sampler is it returns i.i.d. draws from $f^{(c)}$. However, it suffers from several inefficiencies. In particular, we would expect the acceptance probabilities ρ_j given in (4) to decay geometrically with increasing number of sub-posteriors, $|\mathcal{C}|$, as each term in this product is bounded by 1. Furthermore, the acceptance probability $\prod_{j=0}^n \rho_j$ will typically decay exponentially with increasing T . Consequently, a rejection sampling approach for this problem will ultimately be impractical in many practical settings as it will have very small acceptance probabilities. Similarly, the naive importance sampling adaptation of this approach (in which the proposal of the rejection sampler are all retained with a un-normalised importance weight of $\prod_{j=0}^n \tilde{\rho}_j$) will ultimately suffer from the same issues of robustness in practice.

Inspired by the importance sampling approach, the BF approach of Dai et al. (2023) introduced the auxiliary temporal partition \mathcal{P} in order to simulate from g_c using SMC: allowing for the *gradual* coalescence of the C stochastic processes. In particular, we can initialise an SMC algorithm by simulating N particles from the time 0 marginal in h_c (which consists of composing $|\mathcal{C}|$ samples from each of the sub-posterior densities to obtain $\vec{\mathbf{x}}_0^{(c)}$),

and assigning them an initial un-normalised importance weight given by $w'_{0,i} := \rho_0(\vec{\mathbf{x}}_{0,i}^{(c)})$ for $i \in \{1, \dots, N\}$. This initial particle set constitutes an approximation to the time 0 marginal of g_C , and can be sequentially propagated n times (i.e. $|\mathcal{P}| - 1$ times) through the temporal time mesh \mathcal{P} by simulating $\vec{\mathbf{x}}_{j,i}^{(c)} | \vec{\mathbf{x}}_{j-1,i}^{(c)} \sim \mathcal{N}_d(\vec{\mathbf{M}}_{j,i}^{(c)}, \mathbf{V}_j)$ as per (9) and (10) in Proposition 2. In our SMC formulation at each iteration ($j \in \{1, \dots, n\}$) the un-normalised importance weight of every particle is updated by a factor of $\tilde{\rho}_j(\vec{\mathbf{x}}_{j-1,i}^{(c)}, \vec{\mathbf{x}}_{j,i}^{(c)})$ as per (21). Upon normalisation, the resulting weighted particle set after the n th iteration is an approximation of both the time T marginal of g_C and our fusion target $f^{(c)}$. In particular,

$$f^{(c)}(\mathbf{y})d\mathbf{y} \approx \sum_{i=1}^N w_{n,i}^{(c)} \cdot \delta_{\mathbf{y}_i^{(c)}}(d\mathbf{y}). \quad (26)$$

As we remarked upon in Section 2.2, due to this normalisation of the particle set weights we can avoid the need to explicitly compute the constants Φ_c in (16), as they are simply constants which cancel.

As is common in SMC, to avoid *weight degeneracy* in which the variance of the importance weights degrades rapidly in n , we employ a *resampling* strategy (see for instance Gerber et al. (2019) for a recent investigation of the properties of many resampling schemes). In particular, we monitor the particle set for weight degeneracy by estimating its *effective sample size (ESS)* Kong et al. (1994). If the ESS falls below some user-specified threshold then at the beginning of the next iteration we resample the particle set to get N equally weighted particles. In all of our simulations in the subsequent sections, we used residual resampling (Higuchi, 1997; Liu and Chen, 1998; Whitley, 1994).

We term our resulting Fusion approach *Generalised Bayesian Fusion (GBF)* and summarise it in Algorithm 1.

2.4 Practical extensions of Generalised Bayesian Fusion

We now consider the practicalities of Algorithm 1. To generalise the algorithm further (and make it amenable to a recursive divide-and-conquer approach as in Divide-and-Conquer Fusion), we assume we have access to M importance weighted realisations of each sub-posterior, $\{\mathbf{x}_{0,k}^{(c)}, w_k^{(c)}\}_{k=1}^M$ for $c \in \mathcal{C}$. To initialise the algorithm, we start by composing M initial weighted particles by pairing the draws from each sub-posterior $\{\vec{\mathbf{x}}_{0,k}^{(c)}\}_{k=1}^M$, and compute the associated (un-normalised) partial weights $\{w_{0,k}^{(c)'}\}_{k=1}^M$ where $w_{0,k}^{(c)'} := (\prod_{c \in \mathcal{C}} w_k^{(c)}) \cdot \rho_0(\vec{\mathbf{x}}_{0,k}^{(c)})$ for $k = 1, \dots, M$. If we have $M \neq N$, we resample to obtain N samples from each sub-posterior, otherwise, we choose to only resample if the ESS is below some user-specified threshold. Note that in the **Input** step of Algorithm 1, we may have access to different numbers of samples from each sub-posterior: say M_c importance weighted samples for sub-posterior f_c (for $c \in \mathcal{C}$). In order to compose our M partial proposals in Step 1b, there are a number of approaches we could take. As presented above, if $M_c = M$

Algorithm 1 $\text{gbf}(\mathcal{C}, \{\{\mathbf{x}_{0,i}^{(c)}, w_i^{(c)}\}_{i=1}^M, \mathbf{\Lambda}_c\}_{c \in \mathcal{C}}, N, \mathcal{P})$: Generalised Bayesian Fusion (GBF).

1. **Initialisation** ($j = 0$):

- (a) **Input:** Importance weighted realisations $\{\mathbf{x}_{0,i}^{(c)}, w_i^{(c)}\}_{i=1}^M$ for $c \in \mathcal{C} := (c_1, \dots, c_{|\mathcal{C}|})$, the user-specified matrices, $\{\mathbf{\Lambda}_c : c \in \mathcal{C}\}$, the number of particles required, N , and temporal partition $\mathcal{P} := \{t_0, t_1, \dots, t_n : 0 =: t_0 < t_1 < \dots < t_n =: T\}$.
- (b) Compose the importance weighted realisations $\{\bar{\mathbf{x}}_{0,k}^{(C)}, w_{0,k}^{(C)'}\}_{k=1}^M$ where $w_{0,k}^{(C)'} := (\prod_{c \in \mathcal{C}} w_k^{(c)}) \cdot \rho_0(\bar{\mathbf{x}}_{0,k}^{(C)})$ for $k \in \{1, \dots, M\}$ as per (4).
- (c) $w_{0,k}^{(C)}$: For k in 1 to M , compute normalised weight $w_{0,k} = w_{0,k}^{(C)'} / \sum_{k'=1}^M w_{0,k'}^{(C)'}$.
- (d) g_0^M : Set $g_0^M(d\bar{\mathbf{x}}_0^{(C)}) := \sum_{k=1}^M w_{0,k}^{(C)} \cdot \delta_{\bar{\mathbf{x}}_{0,k}^{(C)}}(d\bar{\mathbf{x}}_0^{(C)})$
- (e) $\bar{\mathbf{x}}_{0,i}^{(C)}$: If $M \neq N$, for $i = 1, \dots, N$, resample $\bar{\mathbf{x}}_{0,i}^{(C)} \sim g_0^M$ and reset $w_{0,i}^{(C)} = \frac{1}{N}$.

2. **Iterative updates.** For $j \in \{1, \dots, n\}$:

- (a) **Resample:** If the ESS $:= \left(\sum_{i=1}^N w_{j-1,i}^{(C)}\right)^{-1}$ breaches the lower user-specified threshold, then for $i = 1, \dots, N$, resample $\bar{\mathbf{x}}_{j-1,i}^{(C)} \sim g_{j-1}^N$ and reset $w_{j-1,i}^{(C)} = \frac{1}{N}$.
- (b) For i in 1 to N ,
 - i. $\bar{\mathbf{x}}_{j,i}^{(C)}$: Simulate $\bar{\mathbf{x}}_{j,i}^{(C)} \sim \mathcal{N}_d(\bar{\mathbf{M}}_{j,i}^{(C)}, \mathbf{V}_j)$ as per Proposition 2.
 - ii. $w_{j,i}^{(C)'}$: Compute un-normalised weight $w_{j,i}^{(C)'} = w_{j-1,i}^{(C)} \cdot \tilde{\rho}_j(\bar{\mathbf{x}}_{j-1,i}^{(C)}, \bar{\mathbf{x}}_{j,i}^{(C)})$ as per (21) (using Algorithm 5).
- (c) $w_{j,i}^{(C)}$: For i in 1 to N , compute normalised weight $w_{j,i}^{(C)} = w_{j,i}^{(C)'} / \sum_{k'=1}^N w_{j,k'}^{(C)'}$.
- (d) g_j^N : Set $g_j^N(d\bar{\mathbf{x}}_j^{(C)}) := \sum_{i=1}^N w_{j,i}^{(C)} \cdot \delta_{\bar{\mathbf{x}}_{j,i}^{(C)}}(d\bar{\mathbf{x}}_j^{(C)})$.

3. **Output:** $\{\bar{\mathbf{x}}_{0,i}^{(C)}, \dots, \bar{\mathbf{x}}_{n-1,i}^{(C)}, \mathbf{y}_i^{(c)}, w_{n,i}^{(c)}\}_{i=1}^N$, where $\hat{f}^{(C)}(d\mathbf{y}) := g_n^N(d\mathbf{y}) \approx f^{(C)}(\mathbf{y})d\mathbf{y}$.

for $c \in \mathcal{C}$, we simply pair the sub-posterior draws index-wise. This is a basic merging strategy of the sub-posterior realisations and has the advantage that it can be implemented in $O(M)$ cost (and if $M_c \neq M$ for every $c \in \mathcal{C}$ one could simply sub-sample to obtain a common number of samples from each sub-posterior). However, as noted in Lindsten et al. (2017), while this approach has a low computational cost, it can lead to high variance when the product $\prod_{c \in \mathcal{C}} f_c(\mathbf{x}^{(c)})$ differs substantially from the corresponding marginal of $f^{(C)}$ — which one might expect to be the case in our setting when the sub-posteriors disagree.

We found this simple approach more than adequate in our simulations, but there are more sophisticated options available should they be required in still more challenging settings.

In particular, as described in Lindsten et al. (2017, Section 4.1), at the expense of a computational cost $O(\prod_{c \in \mathcal{C}} M_c)$, one could instead compose all possible permutations of the samples from each sub-posterior before weighting and then resampling to reduce the number of points in the approximation back to a pre-specified number, arriving at a better approximation at a greater cost. They termed this approach “mixture resampling” and also detailed a “lightweight mixture resampling” approach in which more than one permutation, but not all possible permutations, are used and found it to work well; as noted by Kuntz et al. (2022) such a strategy can be connected directly with the theory of incomplete U -statistics and consequently one might hope to realise much of the benefit of mixture resampling at a much reduced cost (see e.g. Kong and Zheng (2021)).

3. A divide-and-conquer approach to Fusion

A key drawback of the Monte Carlo Fusion and Bayesian Fusion approaches of Dai et al. (2019, 2023), and the Generalised Bayesian Fusion approach we introduced and outlined in Section 2, is that it lacks robustness with increasing number of sub-posteriors, $|\mathcal{C}|$. This is unsurprising as the extended target and proposal densities ($g_{\mathcal{C}}$ and $h_{\mathcal{C}}$) are $d(n|\mathcal{C}| + 1)$ -dimensional, and these become increasingly mismatched with increasing dimension. In particular, as a consequence of the definition of ρ_j in (16), the acceptance probability of any rejection-based scheme will decrease geometrically with increasing $|\mathcal{C}|$. Fundamentally, importance sampling variants of this will not address this bottleneck.

As presented both in Dai et al. (2019, 2023), Fusion is an example of a *fork-and-join* approach—*all* of the sub-posteriors are unified in a single step. In particular, within the GBF framework of Section 2 we set $\mathcal{C} := \{1, \dots, C\}$. This is illustrated in the *tree* diagram of Figure 1, where the *leaves* of the tree represent the available sub-posterior densities, the directed edges are used to illustrate the computational flow of MCF, and the *root* vertex of the tree is the desired fusion density, f (as given in (1)).




Figure 1: A tree representation of the fork-and-join approach of Monte Carlo Fusion.

As the goal of the methodology is to approximate f in (1), one could envision a recursive *divide-and-conquer* approach in which the sub-posteriors are combined in stages to recover f . There are a number of possible orderings in which we could combine sub-posteriors, and so we represent these orderings in *tree diagrams*, and term these *hierarchies* (see Figure 2). For instance as illustrated in Figure 2a, one approach would be to combine two sub-posteriors at a time (we term this a *balanced-binary tree* approach). In Figure 2a, the intermediate vertices represent intermediate (*auxiliary*) densities up to proportionality.

The approximation of the distribution associated with any non-leaf vertex is obtained by an application of Fusion methodology to the densities of the children of that vertex. A balanced-binary tree approach is perhaps the most natural way to combine sub-posteriors in a *truly distributed setting* (where the simulation of each sub-posterior has been conducted separately, and so the inferences we wish to combine are distributed). Another approach is given in Figure 2b, whereby sub-posteriors are fused one at a time (which we term a *progressive tree* approach). This is perhaps the most natural approach for an *online setting*. We focus on applying GBF to these two natural hierarchies for the remainder of this paper, although other hierarchies are certainly possible within our framework, and there is no limitation in unifying more than two vertices at any level of a tree (as suggested by both Section 2 and Figure 1).




Figure 2: Illustrative hierarchies for the Fusion problem of (1).

From this recursive perspective, sample approximations of auxiliary densities obtained at one level of any tree are themselves treated as sub-posteriors at the next level up. As such, one can iteratively apply the Fusion methodology of Section 2, working through the levels of the tree from the leaves to the root, using at each stage the output of one step as the input for the subsequent step. An advantage of our divide-and-conquer approach is that as fewer sub-posteriors are combined at each stage, we avoid (at each stage) the rapidly diminishing and variable importance weights.

A divide-and-conquer variant of Sequential Monte Carlo (D&C-SMC) was recently introduced in Lindsten et al. (2017). D&C-SMC generalises the classical SMC framework from sequences (or chains) to trees, such as those in Figures 1 and 2. The theoretical properties of D&C-SMC are increasingly-well characterized and include a strong law of large numbers, finite sample L_p errors bounds as well as a \sqrt{N} -central limit theorem under mild conditions (see Kuntz et al. (2023)). We thus embed our GBF approach within a D&C-SMC algorithm to address the robustness of Fusion with increasing $|C|$, albeit this being a trade-off with the cost of the repeated application of the methodology. In our recursive setting, we unify distributed sample approximations by operating on a tree of *auxiliary* Fusion densities. Let $\mathbb{T} = (\mathcal{V}, \mathcal{E})$ denote a tree with vertices \mathcal{V} and (directed) edge set \mathcal{E} . Let $\text{Leaf}(\mathbb{T})$ denote the leaves of the tree (which represent the sub-posteriors f_1, \dots, f_C), $\text{Root}(\mathbb{T})$ denote the root of the tree (which represents f) and $\text{Ch}(v)$ denote the children of vertex $v \in \mathcal{V}$ where

$\text{Ch}(t) = \emptyset$ if t is a leaf. Let $\mathcal{V} = \{v_0, v_1, \dots, v_C, \dots\}$ be the set of vertices, with $v_0 = \text{Root}(\mathbb{T})$, $\{v_1, \dots, v_C\} = \text{Leaf}(\mathbb{T})$ and as many intermediate vertices as are required to specify the tree.

For the purposes of utilising the methodology developed in Section 2, we define the following notation for non-leaf vertices (i.e. $v \notin \text{Leaf}(\mathbb{T})$): let $\mathcal{C}_v := \cup_{u \in \text{Ch}(v)} \mathcal{C}_u$ denote the index set representing the sub-posteriors that we want to unify for vertex $v \notin \text{Leaf}(\mathbb{T})$. In addition, to simplify the notation and avoid an unnecessary level of subscripts, we index densities and other quantities by v rather than \mathcal{C}_v when it is clear what is intended. In particular, let $\mathbf{\Lambda}_v := \mathbf{\Lambda}_{\mathcal{C}_v}$, $\vec{\mathbf{x}}_t^{(v)} := \vec{\mathbf{x}}_t^{(\mathcal{C}_v)}$, $\tilde{\mathbf{x}}_t^{(v)} := \tilde{\mathbf{x}}_t^{(\mathcal{C}_v)}$, $\mathbf{y}^{(v)} := \mathbf{y}^{(\mathcal{C}_v)}$ where $\mathbf{y}^{(v)} \sim f_v := f^{(\mathcal{C}_v)}$ for $v \notin \text{Leaf}(\mathbb{T})$. Let $\mathbb{W}_{\mathbf{\Lambda}_v, j}$ denote the law of a Brownian bridge $\{\mathbf{X}_t^{(v)}, t \in [t_{j-1}, t_j]\}$ with $\mathbf{X}_{t_{j-1}}^{(v)} := \mathbf{x}_{j-1}^{(v)}$ and $\mathbf{X}_{t_j}^{(v)} := \mathbf{x}_j^{(v)}$ with covariance $\mathbf{\Lambda}_v$ for $j \in \{1, \dots, n\}$. The extended target and proposal densities for vertex $v \notin \text{Leaf}(\mathbb{T})$ are denoted $g_v := g_{\mathcal{C}_v}$ and $h_v := h_{\mathcal{C}_v}$, respectively. Lastly, the importance sampling weights for $v \notin \text{Leaf}(\mathbb{T})$ are given by $\rho_0^{(v)}(\vec{\mathbf{x}}^{(v)}) := \rho_0(\vec{\mathbf{x}}^{(\mathcal{C}_v)})$ and $\rho_j^{(v)}(\vec{\mathbf{x}}^{(v)}, \mathbf{y}^{(v)}) := \rho_j(\vec{\mathbf{x}}^{(\mathcal{C}_v)}, \mathbf{y}^{(\mathcal{C}_v)})$ for all j .

To describe our *Divide-and-Conquer Fusion (D&C-Fusion)* approach, we specify an algorithm that is carried out at each vertex $v \in \mathcal{V}$ which leads to a recursive procedure; an initial call to $\text{D\&C-Fusion}(\text{Root}(\mathcal{V}), \dots)$ carries out the overall approach. For $v \in \mathcal{V}$, we define a procedure (as given in Algorithm 2), which returns a weighted particle set $\{\vec{\mathbf{x}}_{0,i}^{(v)}, \dots, \vec{\mathbf{x}}_{n-1,i}^{(v)}, \mathbf{y}_i^{(v)}, w_{n,i}^{(v)}\}_{i=1}^N$ where $w_{n,i}^{(v)}$ denotes the normalised importance weight of particle i for vertex $v \in \mathcal{V}$. From this particle set, we can take the marginal weighted samples for $\mathbf{y}^{(v)}$ to approximate the fusion density $f_v \propto \prod_{u \in \text{Ch}(v)} f_u$ for vertex $v \in \mathcal{V}$. Recall that the leaf vertices, v_c for $c \in \{1, \dots, C\}$, represent each of the sub-posteriors. It is possible to additionally incorporate importance sampling for the leaf vertices, but for simplicity we assume that we have access to unweighted samples for the sub-posteriors. Therefore, at these leaf vertices, we simply sample from the sub-posteriors. If independent sampling is not feasible, one could use MCMC to obtain unweighted sample approximations at the leaves. Formal arguments (under appropriate regularity conditions) could in principle follow an approach analogous to that in Finke et al. (2020). If v is a non-leaf vertex, we simply call Algorithm 1 by inputting the importance weighted samples $\{\mathbf{y}_i^{(u)}, w_i^{(u)}\}_{i=1}^N$ for $u \in \text{Ch}(v)$. As in standard SMC, although the auxiliary distributions are defined on larger spaces we do not need to retain sampled values which are not subsequently used; to obtain a more computationally manageable algorithm, we can choose to retain only the final parameter space marginal at each vertex (i.e. only returning $\{\mathbf{y}_i^{(v)}, w_i^{(v)}\}_{i=1}^N$) since we only require this to compute the importance weights in Algorithm 1 at each vertex $v \notin \text{Leaf}(\mathbb{T})$.

Note that in Algorithm 2, we allow the user to specify different temporal partitions at each node and level (i.e. $\{\mathcal{P}_u\}_{u \in \text{Ch}(v)}, \mathcal{P}_v$). As we explore fully in Section 4, when we develop guidance for user chosen tuning parameters, having this flexibility on the temporal partition can lead to a far more robust and efficient implementation of Algorithm 2.

Algorithm 2 D&C-Fusion(v, N, \mathcal{P}): Divide-and-Conquer Fusion (D&C-Fusion).

Given: Sub-posteriors, $\{f_u\}_{u \in \text{Leaf}(\mathbb{T})}$, and preconditioning matrices $\{\mathbf{\Lambda}_u\}_{u \in \mathbb{T}}$.

Input: Node in tree, v , the number of particles N , and (optionally) the temporal mesh partitions $\{\mathcal{P}_u\}_{u \in \text{Ch}(v)}$, \mathcal{P}_v .

1. For $u \in \text{Ch}(v)$,
 - (a) $\left\{ \vec{\mathbf{x}}_{0,i}^{(u)}, \dots, \vec{\mathbf{x}}_{n-1,i}^{(u)}, \mathbf{y}_i^{(u)}, w_{n,i}^{(u)} \right\}_{i=1}^N \leftarrow \text{D\&C-Fusion}(u, N, \mathcal{P}_u)$.
 2. If $v \in \text{Leaf}(\mathbb{T})$,
 - (a) For $i = 1, \dots, N$, sample $\mathbf{y}_i^{(v)} \sim f_v(\mathbf{y})$.
 - (b) **Output:** $\{\emptyset, \mathbf{y}_i^{(v)}, \frac{1}{N}\}_{i=1}^N$.
 3. If $v \notin \text{Leaf}(\mathbb{T})$,
 - (a) If \mathcal{P}_v is not inputted, apply guidance from Section 4.1 and Section 4.2.
 - (b) **Output:** Call $\text{gbf}(\text{Ch}(v), \{\{\mathbf{y}_i^{(u)}, w_i^{(u)}\}_{i=1}^N, \mathbf{\Lambda}_u\}_{u \in \text{Ch}(v)}, N, \mathcal{P}_v)$.
-

4. Implementational guidance for Generalised Bayesian Fusion

In this section we develop guidance for choosing the parameter T and the temporal partition \mathcal{P} (and so n implicitly) for our Generalised Bayesian Fusion (GBF) approach (Algorithm 1), the guidance for which can be used directly *at each node* within our Divide-and-Conquer Fusion approach (Algorithm 2). As GBF is fundamentally a sequential Monte Carlo (SMC) algorithm, we want to choose these hyperparameters in such a way to ensure that the discrepancy between subsequent proposal and target distributions are not degenerate. For this reason, and in common with Dai et al. (2023, Section 3), we look at the incremental weight changes and study the *current effective sample size (CESS)* associated with these weights:

$$\text{CESS}_j := \frac{\left(\sum_{i=1}^N \tilde{\rho}_{j,i}\right)^2}{\sum_{i=1}^N \tilde{\rho}_{j,i}^2} \text{ for } j = 1, \dots, n; \quad \text{CESS}_0 := \frac{\left(\sum_{i=1}^N \rho_{0,i}\right)^2}{\sum_{i=1}^N \rho_{0,i}^2}, \quad (27)$$

where $\rho_{0,i}$ and $\tilde{\rho}_{j,i}$ are given in (4) and (21) respectively.

In order to develop heuristics to choose hyper-parameters, we consider the idealised setting of combining multivariate Gaussian sub-posteriors with mean vector \mathbf{a}_c and covariance matrix $b|\mathcal{C}|\mathbf{\Lambda}_c/m$, for some $b > 0$, for $c \in \mathcal{C}$. The target is $f \sim \mathcal{N}_d(\tilde{\mathbf{a}}, b|\mathcal{C}|\mathbf{\Lambda}_c/m)$, where $\tilde{\mathbf{a}} := \left(\sum_{c \in \mathcal{C}} \mathbf{\Lambda}_c^{-1}\right)^{-1} \left(\sum_{c \in \mathcal{C}} \mathbf{\Lambda}_c^{-1} \mathbf{a}_c\right)$ and $\mathbf{\Lambda}_c := \left(\sum_{c \in \mathcal{C}} \mathbf{\Lambda}_c^{-1}\right)^{-1}$.

In BF, this idealised setting was used to help select T and n and, by imposing an additional assumption that the partition was a *regular* mesh, in turn \mathcal{P} . In this section we instead develop guidance for T (see Section 4.1) in the more sophisticated GBF setting, and then

in Section 4.2 investigate the more challenging selection of \mathcal{P} without assumption on its regularity (i.e. permitting an *irregular* choice of mesh)—and so we instead *implicitly find* n . These ideas can also be directly applied to improve BF itself, which we show later in our numerical results.

In our idealised setting, the key consideration is the degree to which the sub-posteriors disagree with one another. To measure how significant the *sub-posterior conflict* is we define

$$\sigma_{\mathbf{a}}^2 := \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\mathbf{a}_c - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\mathbf{a}_c - \tilde{\mathbf{a}}). \quad (28)$$

We further consider the two following conditions in order to explore how the algorithm hyperparameters should change according to sub-posterior heterogeneity:

Condition 8 $SH(\lambda)$. *The sub-posteriors obey the $SH(\lambda)$ condition (for some constant $\lambda > 0$) if*

$$\sigma_{\mathbf{a}}^2 = \frac{b(|\mathcal{C}| - 1)\lambda}{m}. \quad (29)$$

Remark 9 Interpretation of λ . *Of course $SH(\lambda)$ will always hold for some λ , and this condition can alternatively be interpreted as a definition of λ . We will be particularly interested in moderate values of λ close to 1 which will indicate only weak or no sub-posterior discrepancy. $SH(\lambda)$ is a natural condition, arising for instance if $\frac{m}{|\mathcal{C}|}$ of the data is randomly allocated to each sub-posteriors then $\sigma_{\mathbf{a}}^2 \sim \frac{b}{m} \chi_{|\mathcal{C}|-1}^2$ and have mean $\frac{b(|\mathcal{C}|-1)}{m}$. The λ for which $SH(\lambda)$ holds is therefore $\chi_{|\mathcal{C}|-1}^2 / (|\mathcal{C}| - 1)$ and therefore has mean 1 and variance $2 / (|\mathcal{C}| - 1)$. Consequently for large $|\mathcal{C}|$, we would expect λ to be close to 1. In this idealised i.i.d. case, these arguments duplicate classical ANOVA calculations.*

However the $SH(\lambda)$ condition for moderate $\lambda > 1$ is also of interest indicating weak discrepancy between sub-posteriors. This would occur (for instance) if the data consisted of disjoint segments of a long ergodic stationary sequence with no long-range dependence where, in this case, λ is an estimate of the integrated auto-correlation time of the sequence. For this reason, the scenario $\lambda < 1$ would not normally occur (particularly for large $|\mathcal{C}|$).

In the examples later on, we will set $\lambda = 1$ as default, since this is the natural iid scenario. However, as noted above, if we suspect that there is weak discrepancy between the sub-posteriors, or there is some dependency between the subsets of data, we may also choose λ to be slightly greater than 1 or alternatively estimate it from the data.

The defining characteristic of $SH(\lambda)$ is that λ is stable for large data sizes (large m). However for stronger sub-posterior discrepancy, just as the power of ANOVA tests become larger for larger data sets, λ will become much larger with m where there is a systematic difference in the data distributions between sub-posteriors. Now $SH(\lambda)$ will not adequately describe this dependence, and so we consider the following scenario instead:

Condition 10 $SSH(\gamma)$. *The sub-posteriors obey the super sub-posterior heterogeneity $SSH(\gamma)$ condition (for some constant $\gamma > 0$) if*

$$\sigma_{\mathbf{a}}^2 = b\gamma. \quad (30)$$

As with $SH(\lambda)$, this can alternatively be seen as a definition of γ . This setting can arise if the sub-posterior heterogeneity does not decay with data size m .

Remark 11 Choice of b : *In the case that the user-specified matrices $\{\mathbf{\Lambda}_c\}_{c \in \mathcal{C}}$ are chosen to be the estimated covariance matrices for each sub-posterior, then we would set $b = \frac{m}{|\mathcal{C}|}$. Therefore, the sub-posteriors $f_c \sim \mathcal{N}_d(\mathbf{a}_c, \frac{b|\mathcal{C}|}{m} \mathbf{\Lambda}_c)$ have variance which closely matches the sub-posterior variance. In general, we want to choose b such that $\frac{b|\mathcal{C}|}{m} \mathbf{\Lambda}_c$ is close to the variance of sub-posterior f_c for $c \in \mathcal{C}$.*

We study empirically our choices of tuning parameter (T , n and \mathcal{P}) in the idealised settings described by the $SH(\lambda)$ condition (of Condition 8) and $SSH(\gamma)$ condition (of Condition 10) in Sections I.1–I.2 respectively.

Note that the implementational guidance we provide in this section is for the general application and tuning of GBF methodology. In many practical settings there will be additional constraints which require further modification to GBF. This includes settings where latency between cores is problematic, or in scenarios where functional evaluations of the sub-posterior densities f_c are not available. In Appendix H we provide further direction on some of what we envisage to be the most common modifications.

4.1 Guidance for choosing T

In this section, we develop guidance on selecting T for the two idealised settings, $SH(\lambda)$ and $SSH(\gamma)$, defined in Conditions 8 and 10, respectively. In each setting, by first specifying the lower bound on the initial effective sample size that we desire, we can compute a minimum value of T which should be used in Algorithm 1. As choosing a larger value of T typically results in more iterations in GBF, we suggest using the minimum value of T which is suggested. The time horizon T only directly affects the initial weighting given to each of the N particles through ρ_0 in (4). Thus, to develop guidance for T we study $CESS_0$ in (27):

Theorem 12 *Let $f_c \sim \mathcal{N}_d(\mathbf{a}_c, \frac{b|\mathcal{C}|}{m} \mathbf{\Lambda}_c)$ for $c \in \mathcal{C}$, then considering the initial conditional effective sample size $CESS_0$ we have that as $N \rightarrow \infty$, the following convergence in probability holds*

$$N^{-1} CESS_0 \xrightarrow{p} \exp \left\{ - \frac{\sigma_{\mathbf{a}}^2 \left(\frac{b}{m} \right)}{\left(\frac{T}{|\mathcal{C}|} + \frac{b}{m} \right) \left(\frac{T}{|\mathcal{C}|} + \frac{2b}{m} \right)} \right\} \cdot \left[1 + \frac{\left(\frac{|\mathcal{C}|b}{Tm} \right)^2}{1 + \frac{2|\mathcal{C}|b}{Tm}} \right]^{-\frac{(|\mathcal{C}|-1)d}{2}}. \quad (31)$$

Proof See Appendix G. ■

The following corollary considers the effect of T on $CESS_0$ in the $SH(\lambda)$ and $SSH(\gamma)$ settings:

Corollary 13 *If for some constant $k_1 > 0$, T is chosen such that $T \geq \frac{b|\mathcal{C}|^{3/2}k_1}{m}$ for some constant k_1 , then the following lower bounds on $CESS_0$ hold:*

(a) *If $SH(\lambda)$ holds for some $\lambda > 0$, then*

$$\lim_{N \rightarrow \infty} N^{-1} CESS_0 \geq \exp \left\{ -\frac{\lambda}{k_1^2} - \frac{d}{2k_1^2} \right\}. \quad (32)$$

(b) *If $SSH(\gamma)$ holds for some $\gamma > 0$, and $T \geq k_2|\mathcal{C}|^{\frac{1}{2}}$ for some constant k_2 , then*

$$\lim_{N \rightarrow \infty} N^{-1} CESS_0 \geq \exp \left\{ -\frac{b\gamma}{k_1 k_2} - \frac{d}{2k_1^2} \right\}. \quad (33)$$

Proof See Appendix G. ■

We choose k_1 and k_2 by means of Remark 14, which in turn allows us to determine T . As required by Remark 14 we first set $\lambda = 1$ (see Remark 9), b (using Remark 11), and σ_a^2 as per (28).

Remark 14 *Choice of k_1, k_2 : To choose k_1 and k_2 , we first specify $\zeta \in (0, 1)$ to be a lower bound on the initial relative effective sample size that we would desire. We then can consider which situation that we are likely to be in, and then:*

1. *Under $SH(\lambda)$, suppose we want to ensure $N^{-1}CESS_0$ is above $\zeta \in (0, 1)$, from (32), we have $\exp \left\{ -\frac{\lambda}{k_1^2} - \frac{d}{2k_1^2} \right\} = \zeta$, which implies we choose $k_1 = \sqrt{-\frac{(\lambda + \frac{d}{2})}{\log(\zeta)}}$.*
2. *Under $SSH(\gamma)$, suppose we want to ensure $N^{-1}CESS_0$ is above $\zeta \in (0, 1)$, then from (33), we have*

$$\exp \left\{ -\frac{b\gamma}{k_1 k_2} - \frac{d}{2k_1^2} \right\} = \zeta. \quad (34)$$

Recall that for $SSH(\gamma)$, we must have $T \geq \max \left\{ \frac{b|\mathcal{C}|^{3/2}k_1}{m}, |\mathcal{C}|^{\frac{1}{2}}k_2 \right\}$. Since we wish T to be small, we would like k_1 and k_2 to be small, and thus we set these two terms equal to each other and find $k_2 = \frac{b|\mathcal{C}|k_1}{m}$. Substituting into (34), we then choose

$$k_1 = \sqrt{-\frac{(\frac{\gamma m}{\mathcal{C}} + \frac{d}{2})}{\log(\zeta)}}.$$

Given k_1 and k_2 , T can be chosen such that $T \geq \frac{b|\mathcal{C}|^{3/2}k_1}{m}$ if $SH(\lambda)$ holds, and $T \geq \max \left\{ \frac{b|\mathcal{C}|^{3/2}k_1}{m}, |\mathcal{C}|^{\frac{1}{2}}k_2 \right\}$ if $SSH(\gamma)$ holds. Typically we want to minimise iterations of Algorithm 1 Step 2, and so we choose the smallest T which satisfies the user-specified $\zeta \in (0, 1)$.

4.2 Guidance for choosing \mathcal{P}

In order to choose the temporal mesh \mathcal{P} we consider two approaches, each of which is considered and optimised by means of our CESS of (27): i) by first fixing n and assuming a *regular* mesh (as in Dai et al. (2023)), we then optimise for n by reference to the maximally tolerable degradation of CESS_j over any single iterate (see Section 4.2.1); (ii) by starting at $t_0 = 0$ we decide on the placement of t_1 such that we do not violate the maximally tolerable degradation of CESS_1 , and then iterate until we reach T , and so leading to a *irregular* (*adaptive*) mesh and implicitly choosing n (see Section 4.2.2). We summarise these two mesh constructions in Algorithms 3 and 4

To simplify the analysis of Algorithm 1, for which there is considerable flexibility in the choice of proposal distribution for our unbiased estimator of the importance weights (see Theorem 4 of Section 2.2), we assume that we have access to the *optimal* unbiased estimator. Fearnhead et al. (2008, Theorem 1) (and Dai et al. (2023, Appendix B)) show that the variance of the unbiased estimator $a_j \tilde{\rho}_j$ is minimised when $p(\kappa_c | R_c) \sim \text{Poi}(\lambda_c)$, where

$$\lambda_c := \left[\Delta_j \int_{t_{j-1}}^{t_j} \left(U_j^{(c)} - \phi(\mathbf{X}_t^{(c)}) \right)^2 dt \right]^{\frac{1}{2}}, \quad (35)$$

for $c \in \mathcal{C}$. With this choice the second moment is finite and $\mathbb{E} \left[(a_j \tilde{\rho}_j)^2 \right] \leq 1 < \infty$. In practice choosing this optimal distribution for \mathbb{K} is not possible since the integral in (35) cannot be evaluated directly. This is why in Section 2.2 we choose alternative simulatable distributions (as described in Conditions 5–6), which try to match this optimal distribution.

With this optimal choice, we establish the following theorem:

Theorem 15 *Let $p(\kappa_c | R_c)$ in (21) be a Poisson distribution with intensity given in (35), for $c \in \mathcal{C}$, and k_3, k_4 be positive constants. If $\lim_{\Delta_j \rightarrow 0}$ is taken over sequences of $\Delta_j := t_j - t_{j-1} \rightarrow 0$ with*

$$t_j - t_{j-1} \leq \tilde{\Delta}_j := \min \left\{ \frac{b^2 |\mathcal{C}| k_3}{\mathbb{E}[\nu_j] m^2}, \left(\frac{b^2 |\mathcal{C}| k_4}{2m^2 d} \right)^{\frac{1}{2}} \right\}, \quad (36)$$

where

$$\nu_j := \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \left(\mathbf{x}_{j-1}^{(c)} - \mathbf{a}_c \right)^\top \mathbf{\Lambda}_c^{-1} \left(\mathbf{x}_{j-1}^{(c)} - \mathbf{a}_c \right), \quad (37)$$

and the expectation $\mathbb{E}[\nu_j]$ is taken over $\tilde{\mathbf{x}}_{j-1}^{(c)}$, we have

$$\text{plim}_{\Delta_j \rightarrow 0} \text{plim}_{N \rightarrow \infty} N^{-1} \text{CESS}_j \geq e^{-k_3 - k_4}, \quad (38)$$

where plim denotes a limit in probability.

Proof See Appendix G. ■

Remark 16 In Theorem 15, ν_j (as defined in (37)) describes the scaled/weighted average variation of the $|\mathcal{C}|$ trajectories of the distribution of their proposed update locations with respect to their individual sub-posterior means (i.e. describing how far $\mathbf{x}_{j-1}^{(c)}$ is from \mathbf{a}_c). Since the GBF approach has $|\mathcal{C}|$ trajectories which are initialised from their respective sub-posterior distributions and coalesce to a common end point, this variation is mainly determined by a combination of: (i) how large the time horizon T is; (ii) how large the interval we are simulating over for this iteration (t_{j-1}, t_j]; and (iii) how much the sub-posteriors conflict which we determine by looking at the variation in their means as per (28). Given a weighted particle set from the $(j - 1)$ th iteration of the algorithm, $\{\tilde{\mathbf{x}}_{j-1,i}^{(c)}, w_{j-1,i}^{(c)}\}_{i=1}^N$, a natural estimator for $\mathbb{E}[\nu_j]$ is

$$\widehat{\mathbb{E}[\nu_j]} = \sum_{i=1}^N w_{j-1,i}^{(c)} \left(\frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \left(\mathbf{x}_{j-1,i}^{(c)} - \mathbf{a}_c \right)^\top \mathbf{\Lambda}_c^{-1} \left(\mathbf{x}_{j-1,i}^{(c)} - \mathbf{a}_c \right) \right). \quad (39)$$

Following Theorem 15 and Remark 16 we now have the additional problem of specifying k_3 and k_4 , and using the result to develop practical guidance. We do so by means of letting the user choose the meaning parameter $\zeta' \in (0, 1)$, which is we define to be a lower bound on the conditional effective sample size that they would tolerate. We can then select k_3 and k_4 such that $e^{-k_3 - k_4} = \zeta'$ and compute

$$t_j = \min \left\{ T, t_{j-1} + \tilde{\Delta}_j \right\} \quad (40)$$

recursively at each iteration until $j = n$ such that $t_n = T$.

Remark 17 Note that we expect to have very different performance with different choices of k_3 and k_4 . For instance, we can obtain a very high $CESS_j$ by simply choosing k_3 very small and set $k_4 = -\log(\zeta') - k_3$, which ultimately leads to having very small interval sizes $\tilde{\Delta}_j$. Choosing small interval sizes may help computationally simulating $\tilde{\rho}_j$, but this comes at the cost of having more iterations of the algorithm, leading to an increased communication between the cores. Natural choices for jointly specifying k_3 and k_4 are ones which lead to the largest interval size which still satisfies $N^{-1}CESS_j \geq \zeta' \in (0, 1)$, as this minimises the number of iterations of Algorithm 1 Step 2.

We now consider the previously outlined *regular* and *irregular (adaptive)* mesh selection of \mathcal{P} in Section 4.2.1 and Section 4.2.2 respectively.

4.2.1 A REGULAR MESH CONSTRUCTION

Imposing an additional assumption that the temporal partition \mathcal{P} is *regular* simplifies Algorithm 1 as it avoids us having to dynamically compute (36) at each iteration of Step 2. In particular, $\tilde{\Delta}_j = \Delta$ for each $j \in \{1, \dots, n\}$ where $n = \lceil T/\Delta \rceil$ (where $\lceil x \rceil$ denotes the smallest integer greater than or equal to x). This simplification of regularity was suggested in

Dai et al. (2023, Remark 6). They noted that for large datasets in which observations were randomly allocated to sub-posteriors, that one would expect sub-posterior heterogeneity to be small. Hence one would expect $\mathbb{E}[\nu_j]$ to be small (of $\mathcal{O}(m^{-1})$). In their simulations, Dai et al. (2023, Section 3 and 4) set $k_3 = k_4 = 1$ and $\Delta := t_{j-1} - t_j = \sqrt{(b^2|\mathcal{C}|k_4)/(2m^2d)}$ for all j . The rationale presented for these choices in Dai et al. (2023, Remark 6) does not hold in full generality so in this section, we instead develop a more systematic way to construct a regular mesh. In particular, setting $k_3 = k_4$ as they suggest is sub-optimal.

Given a user specified lower bound on CESS_j that they would tolerate (i.e. some $\zeta' \in (0, 1)$), we want to minimise the number of iterations of Algorithm 1 Step 2. This is achieved with reference to Theorem 15 (and in particular (36)). In particular, we choose a combination of k_3 and k_4 such that: (i) $\exp\{-k_3 - k_4\} \geq \zeta'$ (i.e. CESS_j for any j does not violate the chosen ζ'); and (ii), $\frac{b^2|\mathcal{C}|k_3}{\mathbb{E}[\nu_j]m^2} \geq \sqrt{\frac{b^2|\mathcal{C}|k_4}{2m^2d}}$ for each j .

The difficulty here is that at each iteration, we need the average variation of the trajectories $\mathbb{E}[\nu_j]$. Of course, this is not possible directly and so an estimate $\widehat{\mathbb{E}[\nu_j]}$ is computed as per the guidance of Remark 16. To ensure the chosen ζ' is not violated at *any* iteration we follow the guidance of (36) by using a supremum over all intervals of this estimator (i.e. $\sup_j \widehat{\mathbb{E}[\nu_j]}$). This choice allows us to specify k_3 and k_4 , and so in turn n and \mathcal{P} .

Remark 18 *For ease of practically implementing Algorithm 1, it is desirable to avoid any recursive definitions of n and \mathcal{P} (i.e. they are specified prior to calling Algorithm 1 Step 2 where they are required). In this setting we would need to estimate $\sup_j \widehat{\mathbb{E}[\nu_j]}$ based upon only the initial (weighted) sub-posterior realisations $\{\tilde{\mathbf{x}}_{0,i}^{(c)}, w_{0,i}^{(c)}\}_{i=1}^M$ obtained in Algorithm 1 Step 1b.*

Following Remark 16, we would expect $\mathbb{E}[\nu_j]$ to be maximised at $t = T$ (corresponding to (42)), but in some instance may also occur at $t = 0$ (corresponding to (43)). In most practical applications of GBF it will be at $t = T$ as the proposal for the coalescence of the $|\mathcal{C}|$ stochastic processes has a Gaussian distribution with mean $\tilde{\mathbf{x}}_0^{(c)}$ with variance $T\mathbf{\Lambda}_c$ (as a consequence of Proposition 2 and considering $s = 0$ and $t = T$). On the other hand, if the sub-posterior means are very close together, the largest variation in the trajectories from their respective means could occur at the start of the bridge. As such, we propose taking the larger value of those two scenarios to arrive at the following approximation:

$$\sup_j \widehat{\mathbb{E}[\nu_j]} \approx \max\{\Psi_1, \Psi_2\}, \quad (41)$$

where

$$\Psi_1 := \sum_{i=1}^M w_{0,i}^{(c)} \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \left(\tilde{\mathbf{x}}_{0,i}^{(c)} - \mathbf{a}_c \right)^\top \mathbf{\Lambda}_c^{-1} \left(\tilde{\mathbf{x}}_{0,i}^{(c)} - \mathbf{a}_c \right), \quad (42)$$

$$\Psi_2 := \sum_{i=1}^M w_{0,i}^{(c)} \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \left(\mathbf{x}_{0,i}^{(c)} - \mathbf{a}_c \right)^\top \mathbf{\Lambda}_c^{-1} \left(\mathbf{x}_{0,i}^{(c)} - \mathbf{a}_c \right), \quad (43)$$

and where $\tilde{x}_{0,i}^{(C)}$ is defined in (5) and $w_{0,i}^{(C)}$ are the initial particle weights given in Algorithm 1 Step 1b.

Our approximation of $\sup_j \widehat{\mathbb{E}}[\nu_j]$ has obvious limitations: it may not be conservative enough to ensure the user chosen ζ' is not breached; it may be too conservative and lead to choosing n too high. In practice we have found it to be a robust approximation.

Once we have a suitable estimate of $\sup_j \widehat{\mathbb{E}}[\nu_j]$, we need to find a suitable choice for k_3 and k_4 to ensure that we always choose the RHS side of (36) (as that leads to a regular mesh) and satisfies ζ' . As there are many combinations of k_3 and k_4 which can return a regular mesh, we aim to find the combination which returns the largest interval size. We can do this by means of the following proposition which considers the j th interval of the partition:

Proposition 19 *Considering the j th interval of \mathcal{P} (i.e. $[t_{j-1}, t_j]$), given a user-specified threshold $\zeta' \in (0, 1)$ and estimate $\widehat{\mathbb{E}}[\nu_j]$ of $\mathbb{E}[\nu_j]$, then the largest interval size which satisfies $N^{-1}CESS_j \geq \zeta'$ is given by*

$$\tilde{\Delta}_j = \sqrt{\frac{b^2|\mathcal{C}|k_{4,j}}{2m^2d}},$$

where,

$$k_{4,j} := \frac{\left(\frac{\widehat{\mathbb{E}}[\nu_j]^2 m^2}{2b^2|\mathcal{C}|d} - 2\log(\zeta')\right) - \sqrt{\left(2\log(\zeta') - \frac{\widehat{\mathbb{E}}[\nu_j]^2 m^2}{2b^2|\mathcal{C}|d}\right)^2 - 4\log(\zeta')^2}}{2}. \quad (44)$$

Proof See Appendix G. ■

Using Proposition 19, we can substitute our estimate of $\sup_j \widehat{\mathbb{E}}[\nu_j]$ into (44), and subsequently compute the regular interval size $\Delta := \sqrt{\frac{b^2|\mathcal{C}|k_4}{2m^2d}}$ and hence $n = \lceil T/\Delta \rceil$. In effect, here we are setting $k_4 = \sup_j k_{4,j}$. This process is summarised in Algorithm 3.

4.2.2 AN ADAPTIVE MESH CONSTRUCTION

Our presentation of Section 4.2.1 (as opposed to that of Dai et al. (2023)), naturally suggests an *adaptive* approach, leading to a partition \mathcal{P} with an *irregular* mesh. Since the construction of the regular mesh is based upon the *worst* case scenario of the trajectory variation, this leads to an excessive resolution of \mathcal{P} . In this section we will address this.

Instead, suppose we are at the beginning of the j th iteration of Algorithm 1 Step 2. At this point we have in effect simulated our $|\mathcal{C}|$ stochastic processes up to time $t_{j-1} < T$. We can now consider the placement of the next point in the partition (i.e. $\min(t_j, T)$) with reference to the user chosen $\zeta' \in (0, 1)$. In particular, we want the interval to be as large

Algorithm 3 Computing regular mesh \mathcal{P} .

Input: Time $T > 0$ and importance weighted particles $\{\vec{x}_{0,i}^{(C)}, w_{0,i}^{(C)}\}_{i=1}^M$.

1. Compute estimate of $\sup_j \widehat{\mathbb{E}[\nu_j]}$ as per (41).
 2. Compute k_4 using the estimate from Step 1 as per (44).
 3. Compute $\Delta := \sqrt{\frac{b^2|C|k_4}{2m^2d}}$ and let $n = \lceil T/\Delta \rceil$.
 4. For $j \in \{1, \dots, n\}$, let $t_j = \min\{T, t_{j-1} + \Delta\}$.
 5. **Output:** $\mathcal{P} := \{t_0, \dots, t_n\}$.
-

as possible while ensuring that the CESS $_j$ does not degrade by more than ζ' . To do this we can compute an estimate of $\mathbb{E}[\nu_j]$ as per (39) and appeal to Proposition 19 in order to choose $k_{4,j}$, and consequently the interval size Δ_j in order to set $t_j = \min(t_{j-1} + \Delta_j, T)$. Once we reach T we simply halt iterating Algorithm 1 Step 2.

In contrast to the regular mesh construction in Section 4.2.1, we cannot compute the temporal mesh prior to Algorithm 1 Step 2. Therefore, the computation of the interval size for iteration j must be done immediately after Step 2a and prior to Step 2b of Algorithm 1. In this setting, the number of steps in Algorithm 1, n , is not known in advance. Given the construction of the regular mesh assumes the *worst case* interval in selecting the mesh size, we would expect that n would be lower in our adaptive approach. Indeed, we show this empirically in our later simulation studies. We summarise this approach in Algorithm 4.

Algorithm 4 Computing adaptive mesh \mathcal{P} (computing Δ_j at iteration j immediately after Algorithm 1 Step 2a).

Input: Time $T > 0$ and importance weighted particles $\{\vec{x}_{j-1,i}^{(C)}, w_{j-1,i}^{(C)}\}_{i=1}^N$.

1. Compute $\widehat{\mathbb{E}[\nu_j]}$ as per (39).
 2. Compute k_4 with the estimate from Step 1 as per (44).
 3. Compute $t_j = \min\left\{T, t_{j-1} + \sqrt{\frac{b^2|C|k_4}{2m^2d}}\right\}$.
 4. **Output:** $\Delta_j := t_j - t_{j-1}$.
-

In general, it is known that adaptively specifying the sequence of target distributions in sequential Monte Carlo type algorithms can introduce a small bias. In practice, this bias is typically small enough that it can be neglected unless unbiasedness is required specifically, for example to allow the algorithm to be embedded within another. If the small bias is a problem it can be eliminated at the expense of a doubling of the computational cost by first

running an adaptive version of the algorithm and then running it a second time with the sequence of distributions fixed to those learned in the first run.

5. Examples

In this section we consider a number of models applied to a variety datasets, and suppose the dataset is randomly split into C (disjoint) subsets. We compare the performance of our Fusion methodologies (GBF and D&C-Fusion) with other established (approximate) methodologies. To compare performance, we consider their computational run-times and *Integrated Absolute Distance (IAD)*. To compute the IAD we average across each dimension the difference between the true target (*fusion*) density (f), and a kernel density estimate of the draws realised using a given methodology (\hat{f}). In particular,

$$\text{IAD} = \frac{1}{2d} \sum_{j=1}^d \int |\hat{f}(\theta_j) - f(\theta_j)| d\theta_j \in [0, 1]. \quad (45)$$

In the case where the true marginal density is not available analytically, we take as a proxy for the target f a kernel density estimate of f (for instance, obtained by sampling from f or using the output of an MCMC run). As a benchmark for the target f we use **Stan** (Carpenter et al., 2017) to implement an MCMC sampler for the target posterior distribution using the full dataset. In implementing Fusion methodologies we use the GPE-2 variants of Algorithm 2c and Algorithm 5 as before. Our implementation is as presented in Sections 2 and 3, following the guidance presented in Section 4 (but without the inclusion of any adaptations such as those presented in Appendix H). In implementing D&C-Fusion we use the balanced-binary tree hierarchy. For brevity of the main paper, all detailed derivations required for these specific examples have been put in Appendix J.

The established methodologies we consider are *Consensus Monte Carlo (CMC)* (Scott et al., 2016) (implemented using the `parallelMCMCcombine` package in **R** available from <https://github.com/cran/parallelMCMCcombine> (Miroshnikov and Conlon, 2014)), the kernel density averaging approach of Neiswanger et al. (2014) (which we term *KDEMC*, and also implemented using the `parallelMCMCcombine` **R** package), and the *Weierstrass Rejection Sampler (WRS)* (Wang and Dunson, 2013) (implemented using their **R** code available at <https://github.com/wwrechard/weierstrass>).

5.1 Simulation studies

In Appendix I, we study empirically the robustness of our Fusion algorithms in our two idealised settings—the $\text{SH}(\lambda)$ setting (Condition 8) and $\text{SSH}(\gamma)$ setting (Condition 10). We consider a range of different hyperparameter choices and illustrate in both settings that utilising both the guidance for T and the mesh \mathcal{P} , developed in Section 4, drastically improves the performance of both BF and GBF. By comparing the regular and adaptive meshes, we found that the adaptive mesh generally performed better as it provided similar

performance to the regular mesh but at a much reduced computational cost. The full details of these experiments can be found in Sections I.1 and I.2. In Section I.3 we compare the performance of Fusion methodologies with increasing dimensionality and found that our GBF and D&C-Fusion approaches offer the best performance with regards to dimension.

5.2 Robust regression

In this section we consider the ‘*Combined Cycle Power Plant*’ dataset available from the *UCI Machine Learning Repository* (Kaya et al., 2012; Tüfekci, 2014). The dataset comprises $m = 9568$ records of the **net hourly electrical output** of a combined cycle power plant over 6 years between 2006 and 2011, together with four (hourly averaged) *ambient* variables: **temperature**; **ambient-pressure**; **relative-humidity**; and, **exhaust-vacuum**.

To model electrical output using the ambient variables, we use a robust regression model:

$$y_i \sim t(\nu, X_i\boldsymbol{\beta}, \sigma), \quad i = 1, \dots, n,$$

$$\beta_{j'} \sim \mathcal{N}_1\left(\mu_{\beta_{j'}}, \sigma_{\beta_{j'}}^2\right), \quad j' = 0, \dots, p,$$

where $\mathbf{y} \in \mathbb{R}^n$ is the dependent variable (electrical output), $X \in \mathbb{R}^{n \times (p+1)}$ is the design matrix, $\boldsymbol{\beta} \in \mathbb{R}^{p+1}$ is the vector of predictor (ambient) variables which we want to perform inference on. For simplicity, we assume that ν , σ , $\mu_{j'}$ and $\sigma_{\beta_{j'}}$ for $j' = 0, \dots, p$ are known.

For our dataset $p = 4$, and so $d = 5$. We consider $C \in \{4, 8, 16, 32, 64, 128\}$ cores, each of which is assigned a random split of the data. We use **Stan** to sample the sub-posteriors (with $\mu_{j'} = 0$ and $\sigma_{\beta_{j'}}^2 = 10C$ for $j' = 0, \dots, p$), which we will attempt to unify as in (1). We use the approximate CMC, KDEMC and WRS approaches to do this, together with our D&C-Fusion approach. In implementing D&C-Fusion we set $N = 10000$, $\zeta = 0.5$, $\zeta' = 0.05$, and consider both the regular and adaptive mesh variants of the temporal partition, \mathcal{P} . We resample if the ESS falls below $0.5N$. The results are presented in Figure 3.

Figure 3a clearly shows that of all the approaches considered, D&C-Fusion provides the highest quality and most reliable sample approximation for f , and is the most robust to increasing C . Although more expensive computationally, D&C-Fusion has a cost which grows at the same rate as the approximate methodologies considered.

5.3 Negative Binomial regression

Here we consider the ‘*Bike Sharing*’ dataset available on the *UCI Machine Learning Repository* (Fanaee-T and Gama, 2014). The dataset contains $m = 17379$ records of the **total count of bikes on rental each hour**, together with seven variables: **seasonality** (a categorical variable with four levels: spring, summer, autumn, winter); **weekend** (binary, taking value 1 if a weekend, and 0 if not); **holiday**; (binary, taking value 1 if a holiday, and 0 if not); **rush-hour** (binary, taking value 1 if recorded on a weekday between 7AM-9AM




Figure 3: Comparison of competing methodologies to D&C-Fusion applied to a robust regression model using the power plant dataset (see Section 5.2).

or 4PM-7PM, and 0 if not); **weather** (binary, taking value 1 if ‘clear’, and 0 if not); **temperature** (continuous); and, **wind-speed** (continuous). We use treatment contrast coding to encode the **seasonality** via three binary variables.

To model the total count of bikes on rental, we use the following Negative binomial (NB) regression model:

$$y_i \sim \text{NB}(\mu_i, r), \quad \text{where } \log(\mu_i) = X_i \boldsymbol{\beta}, \quad i \in \{1, \dots, n\},$$

$$\beta_{j'} \sim \mathcal{N}_1(\mu_{\beta_{j'}}, \sigma_{\beta_{j'}}^2), \quad j' \in \{0, \dots, p\},$$

where $\mathbf{y} \in \mathbb{R}^n$ is our total count of bikes on rental, $X \in \mathbb{R}^{n \times (p+1)}$ is the design matrix, $\boldsymbol{\beta} \in \mathbb{R}^{p+1}$ is the vector of predictor variables. For simplicity, $r, \mu_{\beta_{j'}}, \sigma_{\beta_{j'}}^2$ for $j' = 0, \dots, p$ are assumed known.

For this data set $p = 9$, and so $d = 10$. As in Section 5.2, we split the dataset amongst $C \in \{4, 8, 16, 32, 64, 128\}$ cores, and use **Stan** with $\mu_{j'} = 0$ and $\sigma_{\beta_{j'}}^2 = 10C$ for $j' = 0, \dots, p$, to recover the respective sub-posteriors. To implement D&C-Fusion we set $N = 10000$, $\zeta = 0.2$, $\zeta' = 0.05$, and consider both regular and adaptive mesh variants of \mathcal{P} , and resample if the ESS drops below $0.5N$.

The results in Figure 4 again show that, when contrasted with existing (approximate) approaches, D&C-Fusion provides the most accurate sample approximation, and is robust and consistent with increasing C .




Figure 4: Comparison of competing methodologies to D&C-Fusion applied to a Negative Binomial regression model using the bike sharing dataset (see Section 5.3).

5.4 Logistic regression

In this section, we apply a logistic regression model to two different datasets (each of which highlight an aspect of Fusion):

$$y_i = \begin{cases} 1 & \text{with probability } \frac{\exp\{\mathbf{x}_i^\top \boldsymbol{\beta}\}}{1 + \exp\{\mathbf{x}_i^\top \boldsymbol{\beta}\}}, \\ 0 & \text{otherwise.} \end{cases} \quad (46)$$

5.4.1 SMALL DATA

Here we consider a *small data size* scenario ($m = 1000$), in which the data is simulated from a logistic regression model (46). This is a variant of Scott et al. (2016, Section 4.3), and is of interest as when the data is (randomly) split among the available cores both *exact* and *approximate* Fusion approaches struggle. This is due to the resulting sub-posteriors being naturally conflicting and lacking fully overlapping support with one another.

Each record of the simulated design matrix contained four covariates in addition to an intercept. The i th entry of the design matrix is given by $\mathbf{x}_i = [1, \zeta_{i,1}, \zeta_{i,2}, \zeta_{i,3}, \zeta_{i,4}]^\top$, where $\zeta_{i,1}, \zeta_{i,2}, \zeta_{i,3}, \zeta_{i,4}$ are random variables generated from a mixture density with a point-mass at zero (and so are either *activated* or not). In particular, we have for $j \in \{1, \dots, 4\}$ that $\zeta_{i,j} \sim p_j \mathcal{N}_1(1, 1) + (1 - p_j) \delta_0$. For this example we chose $p_1 = 0.2$, $p_2 = 0.3$, $p_3 = 0.5$ and $p_4 = 0.01$ (corresponding to a rarely activated covariate). Upon simulating the design matrix, binary observations were obtained by simulation using the parameters $\boldsymbol{\beta} = [-3, 1.2, -0.5, 0.8, 3]^\top$. In total there were a relatively small number of positive responses ($\sum_i y_i = 129$).

To conduct Fusion we first equally split the data between $C \in \{4, 8, 16, 32, 64\}$ cores. We again use `Stan` with Gaussian prior distributions with mean 0 and variance C on each parameter to find a sample approximation of each sub-posterior.

Together with the approximate methodologies, we implemented our D&C-Fusion approach with $N = 10000$, $\zeta = 0.2$, $\zeta' = 0.05$, and both regular and adaptive temporal partition meshes. Here, we also consider applying Generalised Bayesian Fusion (GBF) (i.e. directly applying Algorithm 1 with $C := \{1, \dots, C\}$ (which is equivalent to D&C-Fusion within a fork-and-join tree hierarchy, as per Figure 1)). We present the results in Figure 5.

Considering Figure 5a, we see again that D&C-Fusion achieves the best sample approximation, and the quality of the sample approximation is robust to increasing C . Note that our divide-and-conquer framework offers significant gains, with D&C-Fusion outperforming GBF in terms of robustness with C (even with the same tuning parameter guidance being followed). Note that CMC outperforms all other approximate methodologies, which leaves the practitioner with a clear decision: if a cheap but approximate methodology is needed use CMC, but if accuracy is the goal then D&C-Fusion should be used.





Figure 5: Comparison of competing methodologies to D&C-Fusion applied to a logistic regression model using a simulated small data set (see Section 5.4.1).

5.4.2 NYC FLIGHTS 2013 DATA

Finally, we study Fusion approaches for a logistic regression model (46) applied to the `nycflights13` dataset (obtained from the `nycflights13` **R** package available on CRAN (Wickham, 2021)). In this study we predict on-time arrival of airplanes, by creating binary observations for **arrival-delay** (taking the value 1 if the flight arrived 1 minute or more late, and 0 otherwise). We model this using $p = 20$ predictor variables (so $d = 21$). After removing any entries with NA values, in total the dataset was of size $m = 327346$. This dataset was split randomly across $C \in \{4, 8, 16, 32, 64, 128\}$ cores, and we used `Stan` to

find sample approximations of each sub-posterior (using Gaussian priors with mean 0 and variance C for each parameter). D&C-Fusion was implemented with $N = 30000$, $\zeta = 0.2$ and $\zeta' = 0.05$. The results are shown in Figure 6.

As before, D&C-Fusion provides the best sample approximation and is robust to increasing C , but comes at the expense of increased computational cost. Although approximate methodologies have been specifically developed to tackle Bayesian big-data problems, here we see that they struggle to recover f even in this idealised scenario. They additionally (and critically) lack robustness when scaled with C .



(a) Integrated Absolute Distance.

(b) Computational cost.

Figure 6: Comparison of competing methodologies to D&C-Fusion applied to a logistic regression model using the `nycflights13` dataset (see Section 5.4.2).

To further compare the methodologies, we consider fixing $C = 64$ and varying the computational budget for each method by varying the sample size N in order to study the effect of increased computation on IAD. We again compute the IAD against the same benchmark for the target f that was used above (based upon $N = 30000$ samples using `Stan`). Our results are shown in Figure 7.

The IAD of the approximate methodologies considered in Figure 7 had large variance, and so we run each of these methods 10 times and took an average of the IAD. To show the variability we also plot the minimum and maximum IAD achieved in the 10 runs. The longest run for each approximate methodology was one hour, or when it had become apparent that further computation was not improving IAD. As such, the CMC and KDEMC approaches were considered for a range of sample sizes from $N = 500$ to $N = 200000$, but KDEMC was only considered for $N = 500$ to $N = 50000$. For CMC and WRS, the average and variance of the IAD decreases with more computation, but both methods quickly reach a point where IAD no longer decreases. In Figure 7 we additionally plot a pink dashed line which is the minimum mean value IAD achieved for CMC, as this seems to be the point which the IAD of CMC converges to. For KDEMC increased computation does not improve the average IAD or its variability. As such CMC is clearly the best of the approximate




Figure 7: Integrated absolute distance against computational budget for competing methodologies to D&C-Fusion applied to a logistic regression model using the `nycflight` and fixing $C = 64$ (see Section 5.4.2).

methods, with it achieving the lowest computational cost and lowest IAD. For D&C-Fusion, we considered $N = 500$ to $N = 30000$. We did not perform replicate runs for D&C-Fusion due to the comparative lack of variability in results for this methodology. Of course, being an exact methodology Monte Carlo error can be further decreased by simply increasing N , but does achieve better IAD than CMC for its increased computational cost.

As there is a reasonably large number of data points on each core ($m/64 \approx 5000$) the sub-posteriors are approximately Gaussian, and hence CMC performs unsurprisingly well. Taking into account accuracy and computational budget, then CMC performs the best out of all approximate methodologies here. We are however left with the same conclusion: if the practitioner values accuracy, or they have a poor understanding of the biases induced by an approximate approach, then our D&C-Fusion methodology should be used.

6. Conclusion

The Fusion approach to unifying sub-posteriors into a coherent sample approximation of the posterior (as in (1)), offers fundamental advantages over approximation based approaches. In particular, Fusion avoids imposing any distributional approximation on the sub-posteriors, and so is more robust to a wider range of models, and circumvents needing to understand the impact of imposed approximations on the unified posterior. To date, Fusion approaches have had impractical computational cost in realistic settings, lacking robustness when considering: the number of sub-posteriors being unified; when unifying highly correlated sub-posteriors; the dimensionality of the sub-posteriors; and when considering conflicting sub-posteriors. In this paper, we have substantially addressed the practical issues of Fusion approaches by means of several theoretical and methodological extensions.

In Section 2 we introduced *Generalised Bayesian Fusion (GBF)*, which is a sequential Monte Carlo algorithm that incorporates available global information for each sub-posterior in order to construct informative proposals. As shown in Section 5.1, GBF addresses the lack of robustness when the sub-posteriors have strong correlation structure. By embedding GBF within the Divide-and-Conquer Sequential Monte Carlo (D&C-SMC) framework (Lindsten et al., 2017; Kuntz et al., 2023) in Section 3, we introduced *Divide-and-Conquer Fusion (D&C-Fusion)*, together with a number of *tree hierarchies*, which allow the sub-posteriors to be combined in stages to recover f . By using the provided guidance for selecting the hyperparameters required for the GBF approach (and developed in Section 4), we saw in Section I.3 that our D&C-Fusion approach was the most scalable Fusion approach to date with regards to dimension. In Section 5, we applied our D&C-Fusion methodology to a variety of models with realistic data sets and compared its performance with competing approximate methodologies. In all of these settings, our implementation of D&C-Fusion offered the best performance in terms of Integrated Absolute Distance to an appropriate benchmark, at a modest computational cost. Furthermore, the examples in Section 5 showed that D&C-Fusion is a robust approach to unifying large numbers of sub-posteriors.

There are a number of interesting avenues for extending the work of this paper. Perhaps most interesting is to adapt the D&C-Fusion approach to constraints in practical settings. As discussed in the introduction, one particularly promising direction is when considering (1) under *privacy constraints* of the individual sources (Yıldırım and Ermiş, 2019). In this setting, we may have a number of parties that wish to combine their distributional analysis on a common parameter space and model but cannot reveal their distribution due to confidentiality. This of course requires careful modifications to our approach and is an active area of research of the authors, and motivates variant *tree hierarchies* in D&C-Fusion.

Another application is when considering a truly distributed ‘*big data*’ setting where we have much larger datasets than ones considered in Section 5. In such settings, we may consider a large number of sub-posteriors C since the computational benefit of parallelisation for a divide-and-conquer method is typically proportional to the number of available processors (Nemeth and Sherlock, 2018). Although our divide-and-conquer approach is scalable with C , communication between different cores is expensive in a parallel setting (Scott et al., 2016). We discussed several practical implementation considerations in Appendix H.1, which aim to limit the amount of communication between cores for Algorithm 1, but a considered implementation of these techniques have yet to be explored. To make our Fusion methodology more applicable to large data settings, it would be particularly interesting to investigate embedding a sub-sampling approach within the Fusion algorithms (akin to the approaches of Pollock et al. (2020); Bouchard-Côté et al. (2018); Baker et al. (2019); Bierkens et al. (2019)). First steps in integrating sub-sampling into our D&C-Fusion are considered in Appendix H.2. We also note that there is a growing literature on implementing SMC approaches in parallel and distributed settings (see for instance Doucet and Lee (2018, Section 7.5.3)) which may also be interesting to integrate within Fusion.

From a theoretical perspective, current Fusion methodologies only consider sub-posteriors on a common parameter space. One direction of interest is extending Fusion methodology to

combine sub-posteriors with varying dimension. The *Markov Melding* framework of Goudie et al. (2019) where separate sub-models (potentially of differing dimension) are fitted to different data sources and then joined, is promising. In this setting, the *tree hierarchies* could be defined by the model itself. To mitigate computational robustness of Fusion with increasing dimension in this setting, it may be possible to further utilise the methodology in Lindsten et al. (2017).

Acknowledgments

We would like to thank Louis Aslett, Hector McKimm, Krzysztof Łatuszyński, Nicolas Chopin and Hongsheng Dai for helpful discussions on aspects of the paper. This work was supported by the Engineering and Physical Sciences Research Council under grant numbers EP/K034154/1, EP/K014463/1, EP/N510129/1, EP/R034710/1, EP/R018561/1 and EP/T004134/1 and by The Alan Turing Institute Doctoral Studentship and two Alan Turing Institute programmes; the Lloyd’s Register Foundation programme on ‘Data-centric engineering’ and the UK Government’s ‘Defence and security’ programme.

The data used in this paper are openly available from the following links: the code to reproduce the simulated data used for the simulation studies in Sections 5.1 and 5.4.1 can be found at <https://github.com/rchan26/DCFusion>; the ‘Combined Cycle Power Plant’ dataset (Section 5.2) can be accessed via the UCI Machine Learning Repository at <https://archive.ics.uci.edu/ml/datasets/Combined+Cycle+Power+Plant>; the ‘Bike Sharing’ dataset (Section 5.3) can be accessed via the UCI Machine Learning Repository at <https://archive.ics.uci.edu/ml/datasets/bike+sharing+dataset>; the `nycflights13` dataset (Section 5.4.2) is available through the `nycflights13` R package on CRAN available at <https://github.com/tidyverse/nycflights13>.

Appendix A. Connections with Monte Carlo Fusion and Bayesian Fusion

In this appendix, we more explicitly draw connections with the earlier *Monte Carlo Fusion (MCF)* approach of Dai et al. (2019), and *Bayesian Fusion (BF)* approach of Dai et al. (2023). In particular, we outline how our Generalised Bayesian Fusion (GBF) approach, which we develop in Section 2, improves upon these approaches. We do so by considering several toy examples to illustrate the benefits of the algorithmic developments we have presented in this paper.

Firstly, the theory and methodology developed in Section 2 admits the Monte Carlo Fusion Dai et al. (2019) and Bayesian Fusion (Dai et al., 2023) approaches as a special case and is established in the following corollaries:

Corollary 20 *Setting $\mathcal{P} := \{0, T\}$, $\mathbf{\Lambda}_c = \mathbb{I}_d$ for $c \in \mathcal{C} := \{1, \dots, C\}$, where \mathbb{I}_d is the identity matrix of dimension d and accepting a proposal $\mathbf{y}^{(c)}$ as a sample from (1) with probability $(\rho_0 \cdot \tilde{\rho}_1^{(a)}) (\vec{\mathbf{x}}^{(c)}, \mathbf{y}^{(c)}) \cdot \exp\{\sum_{c \in \mathcal{C}} \Phi_c T\}$, we recover the Monte Carlo Fusion approach of Dai et al. (2019, Algorithm 1).*

Corollary 21 *Setting $\mathbf{\Lambda}_c = \mathbb{I}_d$ for $c \in \mathcal{C} := \{1, \dots, C\}$, where \mathbb{I}_d is the identity matrix of dimension d , and applying the approach outlined in Algorithm 1 recovers the Bayesian Fusion approach of Dai et al. (2023, Algorithm 1).*

We note however that the MCF formulation to arrive at this algorithm is different and is based on the following proposition:

Proposition 22 *Suppose that p_c is the transition density of a Markov chain on \mathbb{R}^d with a stationary probability density proportional to f_c^2 . Then the $(|\mathcal{C}| + 1)d$ -dimensional probability density proportional to the integrable function*

$$g_{\mathcal{C}}^{MCF}(\vec{\mathbf{x}}^{(c)}, \mathbf{y}^{(c)}) := \prod_{c \in \mathcal{C}} \left[f_c^2(\mathbf{x}^{(c)}) \cdot p_c(\mathbf{y}^{(c)} | \mathbf{x}^{(c)}) \cdot \frac{1}{f_c(\mathbf{y}^{(c)})} \right], \quad (47)$$

admits marginal density $f^{(c)} \propto \prod_{c \in \mathcal{C}} f_c$ for $\mathbf{y}^{(c)} \in \mathbb{R}^d$.

Proof By integrating out $\vec{\mathbf{x}}^{(c)}$, we have

$$\begin{aligned} & \int_{\mathbb{R}_d} \dots \int_{\mathbb{R}_d} g_{\mathcal{C}}^{MCF}(\vec{\mathbf{x}}^{(c)}, \mathbf{y}^{(c)}) \, d\mathbf{x}^{(c_1)} \dots d\mathbf{x}^{(c_{|\mathcal{C}|})} \\ &= \prod_{c \in \mathcal{C}} \left[\int_{\mathbb{R}_d} f_c^2(\mathbf{x}^{(1)}) \cdot p_c(\mathbf{y}^{(c)} | \mathbf{x}^{(c)}) \cdot \frac{1}{f_c(\mathbf{y}^{(c)})} \, d\mathbf{x}^{(c)} \right] \\ &= \prod_{c \in \mathcal{C}} \left[\frac{f_c^2(\mathbf{y}^{(c)})}{f_c(\mathbf{y}^{(c)})} \right] \end{aligned}$$

$$= \prod_{c \in \mathcal{C}} f_c(\mathbf{y}^{(c)}) = f^{(c)}(\mathbf{y}^{(c)}). \quad (48)$$

Hence, $\mathbf{y}^{(c)}$ has marginal density $f^{(c)}$. ■

Dai et al. (2019) exploited Proposition 22 by noting that if the index set $\mathcal{C} := \{1, \dots, C\}$, then we recover the target *fusion* density f (as given in (1)). Since $g_{\mathcal{C}}^{MCF}$ will not typically be accessible directly, Dai et al. (2019) proposed sampling from $g_{\mathcal{C}}^{MCF}$ by constructing a suitable $(|\mathcal{C}|+1)d$ -dimensional proposal density (say, $h_{\mathcal{C}}$) for use within a rejection sampling algorithm (Dai et al., 2019, Algorithm 1), and then simply retaining the $\mathbf{y}^{(c)}$ marginal of any accepted draw as a realisation of $f^{(c)}$. Dai et al. (2019) showed that if p_c in Proposition 22 was chosen to be the transition density of a *constant volatility Langevin diffusion* at time T with invariant measure f_c^2 for each $c \in \mathcal{C}$ respectively, then for a (easily accessible) proposal $h_{\mathcal{C}}$ constructed by sampling a single draw from each sub-posterior ($\mathbf{x}^{(c)} \sim f_c$ for $c \in \mathcal{C}$), and then a single Gaussian random variable parameterised by the sub-posterior realisations (corresponding to the $\mathbf{y}^{(c)}$ -marginal), the acceptance probability was readily computable. Although the formulation of the approach was different, this corresponds algorithmically to a rejection sampling variant of GBF and setting $\Lambda_c = \mathbb{I}_d$ for all $c \in \mathcal{C}$ with $\mathcal{P} := \{0, T\}$.

The advantage of the BF and GBF approaches is that our formulation allows for a general temporal mesh \mathcal{P} . We have seen in Section 4 and Section 5.1 that the choice of T and \mathcal{P} can drastically alter the performance of the algorithm and so a clear advantage of BF and GBF over MCF is having a greater flexibility in hyperparameter selection. By being restricted to choosing $\mathcal{P} := \{0, T\}$, in many cases, it will not be possible to choose a T which is large enough for initialisation of the algorithm (i.e. to ensure CESS_0 is large) *and* to choose T small enough for CESS_1 to be sufficiently large. This trade-off in choice of T in MCF is ultimately why it can fail in many practical settings.

As discussed in the introduction, the existing MCF and BF approaches lack robustness in various key practical settings. For the remainder of this section, we re-visit these settings, and with the aid of illustrative examples, show that our new approach addresses these key bottlenecks. Since many of the key limitations are present with both approaches, we focus on comparing against *MCF* approach by setting $\mathcal{P} := \{0, T\}$ in Algorithm 1. We call this variant of GBF with $\mathcal{P} := \{0, T\}$ *Generalised Monte Carlo Fusion (GMCF)*. In particular, in Section A.1 we consider the effect of increasing sub-posterior correlation, in Section A.2 we consider the robustness with increasing numbers of sub-posteriors, and in Section A.3 we consider how to address conflicting sub-posteriors. Throughout this section we use the GPE-2 estimator of ρ_j as given in Definition 6, and use the Trapezoidal rule to estimate the mean γ_c in (24) and fix $\beta_c = 10$ for $c \in \mathcal{C}$. To compare the methodology, we compute both the computational run-times of each methodology and a metric which we term the *Integrated Absolute Distance (IAD)* (45).

A.1 Effect of correlation

One of our key contributions in this paper was the generalisation of BF which incorporated covariance information of the sub-posteriors within our algorithm. In this example, we focus on the illustrative case in which we wish to recover a bi-variate Gaussian target distribution, $f \propto f_1 f_2$, where $f_c \sim \mathcal{N}_2(\mathbf{0}, \Sigma)$ with,

$$\Sigma = \begin{pmatrix} 1.0 & \rho_{\text{corr}} \\ \rho_{\text{corr}} & 1.0 \end{pmatrix}.$$

As we are only considering combining two sub-posteriors in this section, we in effect consider only the GMCF approach. To study the impact of sub-posterior correlation on the robustness of MCF and GMCF (Algorithm 1 with $\mathcal{P} := \{0, T\}$) we can simply consider varying the single parameter ρ_{corr} , and compute the *Effective Sample Size* (ESS) per second averaged across 50 runs in order to compare the efficiency of each methodology. For simplicity, we assume we are able to sample directly from each sub-posterior, and for both methodologies we set $T = 1$. For the purposes of implementing GMCF, we simply set $\Lambda_c = \hat{\Sigma}_c$, where $\hat{\Sigma}_c$ is the *estimated* covariance matrix from the sub-posterior samples for $c = 1, 2$ (and so in effect we have incorporated global information into our proposals), and use a particle set size of $N = 10000$. The results are presented in Figure 8, which clearly show that GMCF is robust to increasing sub-posterior correlation, and offers a significant computational advantage over MCF (which in this case exhibits a strong degradation in efficiency and performance).




Figure 8: ESS per second (averaged over 50 runs) when contrasting Monte Carlo Fusion and Generalised Monte Carlo Fusion, along with increasing sub-posterior correlation, as per the example in Section A.1.

A.2 Effect of hierarchy

In our new formulation outlined in Section 2, we consider the more abstract setting of sampling from $f^{(\mathcal{C})} \propto \prod_{c \in \mathcal{C}} f_c$, where \mathcal{C} is the index set of sub-posteriors which we want

to unify. This abstraction of combining facilitates the recursive use of the algorithms to develop our Divide-and-Conquer Fusion (D&C-Fusion) approach in Section 3 - benefits of which are highlighted clearly in Section 5.4.1 (Here D&C-Fusion outperforms GBF as we increase C).

In this example, we consider the illustrative case of attempting to recover a univariate standard Gaussian target distribution. In particular, we have $f \propto \prod_{c=1}^C f_c$, where $f_c \sim \mathcal{N}(0, C)$ for $c = 1, \dots, C$. By simply varying C , we can study the robustness with increasing numbers of sub-posteriors of MCF (in effect the fork-and-join approach illustrated in Figure 1), and both our suggested versions of Divide-and-Conquer Generalised Monte Carlo Fusion (the balanced-binary tree approach illustrated in Figure 2a, and the progressive tree approach illustrated in Figure 2b). Note that in our chosen idealised setting, there is no advantage conferred with our embedded Generalised Monte Carlo Fusion methodology of Section 2, and so we are simply contrasting hierarchies. In all cases we use a particle set of size $N = 10000$ with resampling if $ESS < N/2$, set $T = 1$, use an appropriately scaled identity as the preconditioning (scalar) matrix, and average across 50 runs. The results are presented in Figure 9, which clearly show that, in contrast to the fork-and-join tree approach, both the balanced-binary tree and progressive tree approaches are robust in recovering the correct posterior distribution in the case of increasing C at the cost of modestly increased computational cost.




Figure 9: Illustrative comparison of the effect of using different hierarchies in Section A.2 (averaged over 50 runs).

A.3 Dealing with conflicting sub-posteriors

Directly unifying C conflicting sub-posteriors (sub-posteriors which have little common support and have high total-variation distance) using a fork-and-join approach as in MCF and Figure 1 is impractical. This can be understood with reference to (4) and (16), which indicates that importance weights will degrade rapidly in this setting.

An approach to deal with conflicting sub-posteriors is to temper the sub-posteriors (to an inverse temperature $\beta \in (0, 1]$ such that there is sufficient sub-posterior overlap), and then propose a suitable tree for which the recursive Divide-and-Conquer Generalised Monte Carlo Fusion approach we introduced in Section 3 could then be applied to recover (1). In particular,

$$f(\mathbf{x}) \propto \prod_{i=1}^{1/\beta} \left[\prod_{c=1}^C f_c^\beta(\mathbf{x}) \right], \quad \text{for } \frac{1}{\beta} \in \mathbb{N}. \quad (49)$$

One such generic tree is provided in Figure 10, in which the tempered sub-posteriors are first unified into $1/\beta \in \mathbb{N}$ tempered posteriors, which are then again unified into f .




Figure 10: Illustrative tree approach for the Fusion problem in the case of conflicting sub-posteriors as in Section A.3. $1/\beta$ copies of the C tempered (and over-lapping) sub-posteriors represent the leaves of the tree, which are unified into $1/\beta$ tempered versions of f (using a suitable tree and D&C-GMCF as in Section 3), and then unified again (using another tree, and D&C-GMCF) to recover f .

To illustrate the advantage of our D&C-GMCF and tempering approach in the case of conflicting sub-posteriors, we consider the scenario of unifying two Gaussian sub-posteriors with the same variance (1), but with different mean ($\pm\mu$). In particular, we have $f \propto f_1 f_2$ where $f_1 \sim \mathcal{N}(-\mu, 1)$ and $f_2 \sim \mathcal{N}(\mu, 1)$. By simply increasing μ we can emulate increasingly conflicting sub-posteriors and study how MCF (which is equivalent to the fork-and-join approach of Figure 1), behaves in terms of the IAD metric and computational time. We contrast this with our tempering approach, considering a range of temperatures $1/\beta \in \{2, 4, 8, 16\}$, and then following the guidance of Figure 10. In particular, we use our D&C-GMCF approach to unify the tempered sub-posteriors with the balanced-binary approach of Figure 2a for both the first and second stage in Figure 10. In all cases, we use a particle set size of $N = 10000$ with resampling if $\text{ESS} < N/2$, set $T = 1$, and average across 50 runs. The results are presented in Figure 11, and show clearly that our D&C-GMCF approach is significantly more robust to conflicting sub-posteriors than the MCF approach where no tempering is applied. A natural trade-off arises when applying the tempering approach suggested, in that decreasing β results in tempered sub-posteriors which are less conflicting and are easier to combine, but there is an increased computational cost in recovering f as an increased number of levels are added to the resulting tree.




Figure 11: Illustrative comparison of using no tempering (solid line), and tempering at 4 different levels together with D&C-GMCF, to combat conflicting sub-posteriors as per Section A.3 (averaged over 50 runs).

Appendix B. Proof of Theorem 1

Proof Following the approach of Dai et al. (2019, Appendix A), we begin by proving that the law of $|\mathcal{C}|$ independent Brownian motions initialised at $\mathbf{x}_0^{(c)} \sim f_c$ for $c \in \mathcal{C}$ and conditioned to coalesce at time T satisfies (2). Here, we use Doob h -transforms (Rogers and Williams, 2000, Chapter IV, Section 6.39) and define the following space-time harmonic function

$$h(t, \vec{\mathbf{x}}_t^{(c)}) = \int \prod_{c \in \mathcal{C}} \frac{1}{\sqrt{2\pi(T-t)|\Lambda_c|}} \exp \left\{ -\frac{(\mathbf{y} - \mathbf{x}_t^{(c)})^\top \Lambda_c^{-1} (\mathbf{y} - \mathbf{x}_t^{(c)})}{2(T-t)} \right\} d\mathbf{y}, \quad (50)$$

which represents the integrated density of coalescence at time T given the current state $\vec{\mathbf{x}}_t^{(c)}$. Then the $|\mathcal{C}|$ conditioned processes satisfy a SDE of the form,

$$d\vec{\mathbf{X}}_t^{(c)} = \vec{\Lambda}^{\frac{1}{2}} d\vec{\mathbf{W}}_t^{(c)} + \vec{\Lambda} \nabla \log \left(h(t, \vec{\mathbf{X}}_t^{(c)}) \right) dt, \quad (51)$$

where $\nabla \log \left(h(t, \vec{\mathbf{x}}_t^{(c)}) \right) =: \left(\mathbf{v}_t^{(c_1)}, \dots, \mathbf{v}_t^{(c_{|\mathcal{C}|})} \right)$ is a collection of $|\mathcal{C}|d$ -dimensional vectors and

$$\vec{\Lambda}^{\frac{1}{2}} = \begin{pmatrix} \Lambda_{c_1}^{\frac{1}{2}} & \mathbf{0}_{d \times d} & \cdots & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & \Lambda_{c_2}^{\frac{1}{2}} & \cdots & \mathbf{0}_{d \times d} \\ \vdots & \ddots & \ddots & \vdots \\ \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} & \cdots & \Lambda_{c_{|\mathcal{C}|}}^{\frac{1}{2}} \end{pmatrix}, \quad \vec{\Lambda} = \begin{pmatrix} \Lambda_{c_1} & \mathbf{0}_{d \times d} & \cdots & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & \Lambda_{c_2} & \cdots & \mathbf{0}_{d \times d} \\ \vdots & \ddots & \ddots & \vdots \\ \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} & \cdots & \Lambda_{c_{|\mathcal{C}|}} \end{pmatrix},$$

where $\Lambda_c^{\frac{1}{2}}$ is the (positive semi-definite) square root of Λ_c where $\Lambda_c^{\frac{1}{2}} \Lambda_c^{\frac{1}{2}} = \Lambda_c$ for $c \in \mathcal{C}$, and $\mathbf{0}_{d \times d}$ denotes the $d \times d$ matrix with all elements equal to 0.

Considering the c th term and letting $\Lambda_c^{-1} = \sum_{c \in \mathcal{C}} \Lambda_c^{-1}$, then

$$\begin{aligned}
 \mathbf{v}_t^{(c)} &= \frac{\int \left(\frac{\Lambda_c^{-1}(\mathbf{y} - \mathbf{x}_t^{(c)})}{T-t} \right) \prod_{c \in \mathcal{C}} \frac{1}{\sqrt{2\pi(T-t)|\Lambda_c|}} \exp \left\{ -\frac{(\mathbf{y} - \mathbf{x}_t^{(c)})^\top \Lambda_c^{-1}(\mathbf{y} - \mathbf{x}_t^{(c)})}{2(T-t)} \right\} d\mathbf{y}}{\int \prod_{c \in \mathcal{C}} \frac{1}{\sqrt{2\pi(T-t)|\Lambda_c|}} \exp \left\{ -\frac{(\mathbf{y} - \mathbf{x}_t^{(c)})^\top \Lambda_c^{-1}(\mathbf{y} - \mathbf{x}_t^{(c)})}{2(T-t)} \right\} d\mathbf{y}} \\
 &= \frac{\int \left(\frac{\Lambda_c^{-1} \mathbf{y}}{T-t} \right) \exp \left\{ -\frac{(\mathbf{y} - \tilde{\mathbf{x}}_t)^\top \Lambda_c^{-1}(\mathbf{y} - \tilde{\mathbf{x}}_t)}{2(T-t)} \right\} d\mathbf{y}}{\int \exp \left\{ -\frac{(\mathbf{y} - \tilde{\mathbf{x}}_t)^\top \Lambda_c^{-1}(\mathbf{y} - \tilde{\mathbf{x}}_t)}{2(T-t)} \right\} d\mathbf{y}} - \frac{\Lambda_c^{-1} \mathbf{x}_t^{(c)}}{T-t} \\
 &= \frac{\Lambda_c^{-1} (\tilde{\mathbf{x}}_t - \mathbf{x}_t^{(c)})}{T-t}.
 \end{aligned}$$

Consequently, we have

$$\nabla \log \left(h(t, \tilde{\mathbf{x}}_t^{(c)}) \right) = \left(\frac{\Lambda_{c_1}^{-1} (\tilde{\mathbf{x}}_t - \mathbf{x}_t^{(c_1)})}{T-t}, \dots, \frac{\Lambda_{c_{|\mathcal{C}|}}^{-1} (\tilde{\mathbf{x}}_t - \mathbf{x}_t^{(c_{|\mathcal{C}|})})}{T-t} \right), \quad (52)$$

and (2) holds.

Next, we show that under \mathbb{F} this common value has density f . Since \mathbb{P} is the measure for $|\mathcal{C}|$ coalesced Brownian motions (shown above), from (3), we can write \mathbb{F} as

$$\begin{aligned}
 d\mathbb{F}(\mathfrak{X}) &\propto d\mathbb{P}(\mathfrak{X}) \cdot \rho_0(\tilde{\mathbf{x}}_0^{(c)}) \cdot \prod_{c \in \mathcal{C}} \left[\exp \left\{ -\int_0^T \phi_c(\mathbf{X}_t^{(c)}) dt \right\} \right] \\
 &\propto \left[\prod_{c \in \mathcal{C}} f_c(\mathbf{x}_0^{(c)}) \right] \cdot \exp \left\{ -\frac{(\mathbf{y}^{(c)} - \tilde{\mathbf{x}}_0^{(c)})^\top \Lambda_c^{-1}(\mathbf{y}^{(c)} - \tilde{\mathbf{x}}_0^{(c)})}{2T} \right\} \cdot d\bar{\mathbb{W}}_\Lambda(\mathfrak{X}) \\
 &\quad \cdot \exp \left\{ -\sum_{c \in \mathcal{C}} \frac{(\tilde{\mathbf{x}}_0^{(c)} - \mathbf{x}_0^{(c)})^\top \Lambda_c^{-1}(\tilde{\mathbf{x}}_0^{(c)} - \mathbf{x}_0^{(c)})}{2T} \right\} \cdot \prod_{c \in \mathcal{C}} \left[\exp \left\{ -\int_0^T \phi_c(\mathbf{X}_t^{(c)}) dt \right\} \right] \\
 &= \left[\prod_{c \in \mathcal{C}} f_c(\mathbf{x}_0^{(c)}) \right] \cdot \exp \left\{ -\sum_{c \in \mathcal{C}} \frac{(\mathbf{y}^{(c)} - \mathbf{x}_0^{(c)})^\top \Lambda_c^{-1}(\mathbf{y}^{(c)} - \mathbf{x}_0^{(c)})}{2T} \right\} \cdot d\bar{\mathbb{W}}_\Lambda(\mathfrak{X}) \\
 &\quad \cdot \prod_{c \in \mathcal{C}} \left[\exp \left\{ -\int_0^T \phi_c(\mathbf{X}_t^{(c)}) dt \right\} \right], \quad (53)
 \end{aligned}$$

where $\bar{\mathbb{W}}_\Lambda$ denotes the law of $|\mathcal{C}|$ independent Brownian bridges $\{\mathbf{X}_t^{(c)}, t \in [0, T]\}_{c \in \mathcal{C}}$ starting at $\mathbf{X}_0^{(c)} := \mathbf{x}_0^{(c)}$ and ending at $\mathbf{X}_T^{(c)} := \mathbf{y}^{(c)}$ (with covariance Λ_c). Let $g_{\mathcal{C}}(\tilde{\mathbf{x}}_0^{(c)}, \mathbf{y}^{(c)})$ denote the marginal distribution of \mathbb{F} at $\tilde{\mathbf{x}}_0^{(c)}$ and $\tilde{\mathbf{x}}_T^{(c)} =: \mathbf{y}^{(c)}$, then we have

$$g_{\mathcal{C}}(\tilde{\mathbf{x}}_0^{(c)}, \mathbf{y}^{(c)}) \propto \prod_{c \in \mathcal{C}} \left[f_c(\mathbf{x}_0^{(c)}) \right] \cdot \exp \left\{ -\sum_{c \in \mathcal{C}} \frac{(\mathbf{y}^{(c)} - \mathbf{x}_0^{(c)})^\top \Lambda_c^{-1}(\mathbf{y}^{(c)} - \mathbf{x}_0^{(c)})}{2T} \right\}$$

$$\begin{aligned}
 & \cdot \prod_{c \in \mathcal{C}} \left[\exp \left\{ - \int_0^T \phi_c \left(\mathbf{X}_t^{(c)} \right) dt \right\} \right], \\
 & = \prod_{c \in \mathcal{C}} \left[f_c^2 \left(\mathbf{x}_0^{(c)} \right) \cdot p_c \left(\mathbf{y}^{(c)} \mid \mathbf{x}_0^{(c)} \right) \cdot \frac{1}{f_c \left(\mathbf{y}^{(c)} \right)} \right], \tag{54}
 \end{aligned}$$

where

$$\begin{aligned}
 p_c \left(\mathbf{y}^{(c)} \mid \mathbf{x}_0^{(c)} \right) & \propto \frac{f_c \left(\mathbf{y}^{(c)} \right)}{f_c \left(\mathbf{x}_0^{(c)} \right)} \cdot \exp \left\{ - \frac{\left(\mathbf{y}^{(c)} - \mathbf{x}_0^{(c)} \right)^\top \boldsymbol{\Lambda}_c^{-1} \left(\mathbf{y}^{(c)} - \mathbf{x}_0^{(c)} \right)}{2T} \right\} \\
 & \cdot \mathbb{E}_{\mathbb{W}_{\boldsymbol{\Lambda}_c}} \left[\exp \left\{ - \int_0^T \phi_c \left(\mathbf{X}_t^{(c)} \right) dt \right\} \right]. \tag{55}
 \end{aligned}$$

Using the Dacunha-Castelle representation (Dacunha-Castelle and Florens-Zmirou, 1986, Lemma 1), this is the transition density density of a Langevin diffusion with covariance matrix $\boldsymbol{\Lambda}_c$ over time $t \in [0, T]$. Critically, this diffusion process has invariant density proportional to f_c^2 , so

$$\int p \left(\mathbf{y}^{(c)} \mid \mathbf{x}_0^{(c)} \right) f_c^2 \left(\mathbf{x}_0^{(c)} \right) d\mathbf{x}_0^{(c)} = f_c^2 \left(\mathbf{y}^{(c)} \right).$$

By integrating out $\vec{\mathbf{x}}_0^{(c)}$ in (54), we can see that $g_{\mathcal{C}}(\vec{\mathbf{x}}_0^{(c)}, \mathbf{y}^{(c)})$ admits $f^{(c)}$ as a marginal. ■

Appendix C. Proof of Proposition 2

Proof Theorem a: We begin by deriving the joint density of $\vec{\mathbf{X}}_t^{(c)}$ conditional on the state at time s , $\vec{\mathbf{x}}_s^{(c)}$. Firstly, consider the $d(|\mathcal{C}| + 1)$ dimensional joint density of $\vec{\mathbf{X}}_t^{(c)}$ and end-point $\mathbf{y}^{(c)}$ conditional on $\vec{\mathbf{x}}_s^{(c)}$, which we denote as p_1 , then

$$-2 \log p_1 = D_1 + D_2,$$

where D_1 is the log-density of $\mathbf{y}^{(c)}$ conditional on $\vec{\mathbf{x}}_s^{(c)}$ and given by

$$D_1 = \sum_{c \in \mathcal{C}} \frac{\left(\mathbf{y}^{(c)} - \mathbf{x}_s^{(c)} \right)^\top \boldsymbol{\Lambda}_c^{-1} \left(\mathbf{y}^{(c)} - \mathbf{x}_s^{(c)} \right)}{T - s} + k_1$$

where k_1 is a constant; D_2 is the log-density of $\vec{\mathbf{X}}_t^{(c)}$ conditional on $\vec{\mathbf{x}}_s^{(c)}$ and $\mathbf{y}^{(c)}$ (which is simply the log-density of $|\mathcal{C}|$ Brownian bridges with respective covariance matrices $\boldsymbol{\Lambda}_c$ for $c \in \mathcal{C}$), given by

$$\begin{aligned}
 D_2 = & \\
 & \sum_{c \in \mathcal{C}} \frac{T - s}{(t - s)(T - t)} \left[\mathbf{x}_t^{(c)} - \frac{t - s}{T - s} \mathbf{y}^{(c)} - \frac{T - t}{T - s} \mathbf{x}_s^{(c)} \right]^\top \boldsymbol{\Lambda}_c^{-1} \left[\mathbf{x}_t^{(c)} - \frac{t - s}{T - s} \mathbf{y}^{(c)} - \frac{T - t}{T - s} \mathbf{x}_s^{(c)} \right]
 \end{aligned}$$

+ k_2 ,

where k_2 is a constant. We therefore have

$$\begin{aligned}
 -2 \log p_1 &= \frac{(\mathbf{y}^{(c)} - \tilde{\mathbf{x}}_s^{(c)})^\top \boldsymbol{\Lambda}_c^{-1} (\mathbf{y}^{(c)} - \tilde{\mathbf{x}}_s^{(c)})}{T-s} \\
 &+ \sum_{c \in \mathcal{C}} \left[\frac{t-s}{(T-t)(T-s)} \mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \mathbf{y}^{(c)} - \frac{2}{T-t} \mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \mathbf{x}_t^{(c)} + \frac{2}{T-s} \mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \mathbf{x}_s^{(c)} \right] \\
 &+ \sum_{c \in \mathcal{C}} \frac{T-s}{(t-s)(T-t)} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right]^\top \boldsymbol{\Lambda}_c^{-1} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right] + k_3 \\
 &= \frac{1}{T-s} \left[\mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \mathbf{y}^{(c)} - 2 \mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \tilde{\mathbf{x}}_s^{(c)} \right] \\
 &+ \left[\frac{t-s}{(T-t)(T-s)} \mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \mathbf{y}^{(c)} - \frac{2}{T-t} \mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \tilde{\mathbf{x}}_t^{(c)} + \frac{2}{T-s} \mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \tilde{\mathbf{x}}_s^{(c)} \right] \\
 &+ \sum_{c \in \mathcal{C}} \frac{T-s}{(t-s)(T-t)} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right]^\top \boldsymbol{\Lambda}_c^{-1} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right] + k_4 \\
 &= \left[\frac{1}{T-s} + \frac{t-s}{(T-t)(T-s)} \right] \mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \mathbf{y}^{(c)} - \frac{2}{T-t} \mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \tilde{\mathbf{x}}_t^{(c)} \\
 &+ \sum_{c \in \mathcal{C}} \frac{T-s}{(t-s)(T-t)} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right]^\top \boldsymbol{\Lambda}_c^{-1} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right] + k_4 \\
 &= \frac{1}{T-t} \left[\mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \mathbf{y}^{(c)} - 2 \mathbf{y}^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \tilde{\mathbf{x}}_t^{(c)} \right] \\
 &+ \sum_{c \in \mathcal{C}} \frac{T-s}{(t-s)(T-t)} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right]^\top \boldsymbol{\Lambda}_c^{-1} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right] + k_4 \\
 &= \frac{1}{T-t} \left[(\mathbf{y}^{(c)} - \tilde{\mathbf{x}}_t^{(c)})^\top \boldsymbol{\Lambda}_c^{-1} (\mathbf{y}^{(c)} - \tilde{\mathbf{x}}_t^{(c)}) \right] - \frac{1}{T-t} \tilde{\mathbf{x}}_t^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \tilde{\mathbf{x}}_t^{(c)} \\
 &+ \sum_{c \in \mathcal{C}} \frac{T-s}{(t-s)(T-t)} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right]^\top \boldsymbol{\Lambda}_c^{-1} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right] + k_4,
 \end{aligned}$$

where k_3 and k_4 are constants, and $\boldsymbol{\Lambda}_c := (\sum_{c \in \mathcal{C}} \boldsymbol{\Lambda}_c^{-1})^{-1}$.

Next, we integrate out $\mathbf{y}^{(c)}$ to obtain the $d|\mathcal{C}|$ -dimensional density of $\vec{\mathbf{X}}_t^{(c)}$ conditional on $\tilde{\mathbf{x}}_s^{(c)}$, which we denote p_2 :

$$\begin{aligned}
 -2 \log p_2 &= -\frac{1}{T-t} \tilde{\mathbf{x}}_t^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \tilde{\mathbf{x}}_t^{(c)} + \sum_{c \in \mathcal{C}} \frac{T-s}{(t-s)(T-t)} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right]^\top \boldsymbol{\Lambda}_c^{-1} \left[\mathbf{x}_t^{(c)} - \frac{T-t}{T-s} \mathbf{x}_s^{(c)} \right] \\
 &+ k_5 \\
 &= -\frac{1}{T-t} \tilde{\mathbf{x}}_t^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \tilde{\mathbf{x}}_t^{(c)} + \sum_{c \in \mathcal{C}} \frac{T-s}{(t-s)(T-t)} \left[\mathbf{x}_t^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \mathbf{x}_t^{(c)} - 2 \left(\frac{T-t}{T-s} \right) \mathbf{x}_t^{(c)\top} \boldsymbol{\Lambda}_c^{-1} \mathbf{x}_s^{(c)} \right] \\
 &+ k_6,
 \end{aligned}$$

where k_5 and k_6 are constants. Noting that

$$\begin{aligned}\tilde{\mathbf{x}}_t^{(c)\top} \Lambda_c^{-1} \tilde{\mathbf{x}}_t^{(c)} &= \left[\left(\sum_{c \in \mathcal{C}} \Lambda_c^{-1} \right)^{-1} \left(\sum_{c \in \mathcal{C}} \Lambda_c^{-1} \mathbf{x}_t^{(c)} \right) \right]^\top \left(\sum_{c \in \mathcal{C}} \Lambda_c^{-1} \right) \left[\left(\sum_{c \in \mathcal{C}} \Lambda_c^{-1} \right)^{-1} \left(\sum_{c \in \mathcal{C}} \Lambda_c^{-1} \mathbf{x}_t^{(c)} \right) \right] \\ &= \left(\sum_{c \in \mathcal{C}} \Lambda_c^{-1} \mathbf{x}_t^{(c)} \right)^\top \Lambda_c \left(\sum_{c \in \mathcal{C}} \Lambda_c^{-1} \mathbf{x}_t^{(c)} \right) \\ &= \sum_{i,j \in \mathcal{C}} \mathbf{x}_t^{(i)\top} \left(\Lambda_i^{-1} \Lambda_c \Lambda_j^{-1} \right) \mathbf{x}_t^{(j)}.\end{aligned}$$

So we have,

$$\begin{aligned}-2 \log p_2 &= -\frac{1}{T-t} \sum_{i,j \in \mathcal{C}} \mathbf{x}_t^{(i)\top} \left(\Lambda_i^{-1} \Lambda_c \Lambda_j^{-1} \right) \mathbf{x}_t^{(j)} + \frac{T-s}{(t-s)(T-t)} \sum_{c \in \mathcal{C}} \mathbf{x}_t^{(c)\top} \Lambda_c^{-1} \mathbf{x}_t^{(c)} \\ &\quad - \frac{2}{t-s} \sum_{c \in \mathcal{C}} \mathbf{x}_t^{(c)\top} \Lambda_c^{-1} \mathbf{x}_s^{(c)} + k_6 \\ &= \frac{T-s}{(t-s)(T-t)} \sum_{c \in \mathcal{C}} \mathbf{x}_t^{(c)\top} \Lambda_c^{-1} \mathbf{x}_t^{(c)} - \frac{1}{T-t} \sum_{c \in \mathcal{C}} \mathbf{x}_t^{(c)\top} \left(\Lambda_c^{-1} \Lambda_c \Lambda_c^{-1} \right) \mathbf{x}_t^{(c)} \\ &\quad - \frac{1}{T-t} \sum_{\substack{i,j \in \mathcal{C} \\ i \neq j}} \mathbf{x}_t^{(i)\top} \left(\Lambda_i^{-1} \Lambda_c \Lambda_j^{-1} \right) \mathbf{x}_t^{(j)} - \frac{2}{t-s} \sum_{c \in \mathcal{C}} \mathbf{x}_t^{(c)\top} \Lambda_c^{-1} \mathbf{x}_s^{(c)} + k_6 \\ &= \bar{\mathbf{x}}_t^\top \mathbf{V}_{s,t}^{-1} \bar{\mathbf{x}}_t^{(c)} - \frac{2}{t-s} \bar{\mathbf{x}}_t^\top \mathbf{L}^{-1} \bar{\mathbf{x}}_s^{(c)} + k_6\end{aligned}$$

where

$$\mathbf{V}_{s,t}^{-1} = \begin{pmatrix} \Sigma_{11}^{-1} & \Sigma_{12}^{-1} & \cdots & \Sigma_{1|\mathcal{C}|}^{-1} \\ \Sigma_{21}^{-1} & \Sigma_{22}^{-1} & \cdots & \Sigma_{2|\mathcal{C}|}^{-1} \\ \vdots & \vdots & \ddots & \vdots \\ \Sigma_{|\mathcal{C}|1}^{-1} & \Sigma_{|\mathcal{C}|2}^{-1} & \cdots & \Sigma_{|\mathcal{C}||\mathcal{C}|}^{-1} \end{pmatrix} \in \mathbb{R}^{|\mathcal{C}|d \times |\mathcal{C}|d}, \quad (56)$$

with

$$\begin{aligned}\Sigma_{ii}^{-1} &= \frac{T-s}{(t-s)(T-t)} \Lambda_{c_i}^{-1} - \frac{1}{T-t} \left(\Lambda_{c_i}^{-1} \Lambda_c \Lambda_{c_i}^{-1} \right) \in \mathbb{R}^{d \times d}, \\ \Sigma_{ij}^{-1} &= -\frac{1}{T-t} \left(\Lambda_{c_i}^{-1} \Lambda_c \Lambda_{c_j}^{-1} \right) \in \mathbb{R}^{d \times d},\end{aligned}$$

for $i, j = 1, \dots, |\mathcal{C}|$, and

$$\mathbf{L}^{-1} = \begin{pmatrix} \Lambda_{c_1}^{-1} & \mathbf{0}_{d \times d} & \cdots & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & \Lambda_{c_2}^{-1} & \cdots & \mathbf{0}_{d \times d} \\ \vdots & \ddots & \vdots & \vdots \\ \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} & \cdots & \Lambda_{c_{|\mathcal{C}|}}^{-1} \end{pmatrix} \in \mathbb{R}^{|\mathcal{C}|d \times |\mathcal{C}|d},$$

where $\mathbf{0}_{d \times d}$ is the $d \times d$ with all elements zero. We finally complete the square to get

$$-2 \log p_2 = \bar{\mathbf{x}}_t^{(c)\top} \mathbf{V}_{s,t}^{-1} \bar{\mathbf{x}}_t^{(c)} - 2 \bar{\mathbf{x}}_t^{(c)\top} \mathbf{V}_{s,t}^{-1} \bar{\mathbf{M}}_{s,t}^{(c)} + k_6,$$

where

$$\vec{M}_{s,t}^{(c)} = \frac{\mathbf{V}_{s,t} \mathbf{L}^{-1} \vec{\mathbf{x}}_s^{(c)}}{t-s}.$$

Inverting $\mathbf{V}_{s,t}^{-1}$ in (56), we obtain (10) and subsequently we can get the expression for $\mathbf{M}_{s,t}^{(c)}$ in (9) to prove the statement in part a of Theorem 1.

For part b, for $c \in \mathcal{C}$, the law of $\{\mathbf{X}_t^{(c)}, t \in (0, T)\}$ conditional on endpoints $\mathbf{x}_0^{(c)}$ and $\mathbf{y}^{(c)}$ is that of a Brownian bridge. This statement in the theorem holds from the standard properties of Brownian bridges (with covariance matrix $\mathbf{\Lambda}_c$). In particular, considering the distribution of $\mathbf{X}_q^{(c)}$ at an intermediate point $q \in (s, t)$ given the positions $\mathbf{X}_s^{(c)} = \mathbf{x}_s^{(c)}$ and $\mathbf{X}_t^{(c)} = \mathbf{x}_t^{(c)}$ at times s and t respectively, then we have

$$\begin{aligned} & \mathbb{P}\left(\mathbf{X}_q = \mathbf{w} \mid \mathbf{X}_s^{(c)} = \mathbf{x}_s^{(c)}, \mathbf{X}_t^{(c)} = \mathbf{x}_t^{(c)}\right) \\ & \propto \mathbb{P}\left(\mathbf{X}_t^{(c)} = \mathbf{x}_t^{(c)} \mid \mathbf{X}_s^{(c)} = \mathbf{x}_s^{(c)}, \mathbf{X}_q = \mathbf{w}\right) \cdot \mathbb{P}\left(\mathbf{X}_q = \mathbf{w} \mid \mathbf{X}_s^{(c)} = \mathbf{x}_s^{(c)}\right) \\ & \propto \mathbb{P}\left(\mathbf{X}_t^{(c)} = \mathbf{x}_t^{(c)} \mid \mathbf{X}_q = \mathbf{w}\right) \cdot \mathbb{P}\left(\mathbf{X}_q = \mathbf{w} \mid \mathbf{X}_s^{(c)} = \mathbf{x}_s^{(c)}\right) \\ & \propto \exp\left(-\frac{(\mathbf{x}_t^{(c)} - \mathbf{w})^\top \mathbf{\Lambda}_c^{-1} (\mathbf{x}_t^{(c)} - \mathbf{w})}{2(t-q)}\right) \cdot \exp\left(-\frac{(\mathbf{w} - \mathbf{x}_s^{(c)})^\top \mathbf{\Lambda}_c^{-1} (\mathbf{w} - \mathbf{x}_s^{(c)})}{2(q-s)}\right), \end{aligned}$$

and hence we arrive at the result in the statement. \blacksquare

Appendix D. Proof of Proposition 3

Proof First note that we can rewrite (6) as follows,

$$\phi_c(\mathbf{x}) = \frac{1}{2} \left(\left\| \mathbf{\Lambda}_c^{\frac{1}{2}} \nabla \log f_c(\mathbf{x}) \right\|^2 + \text{Tr}(\mathbf{\Lambda}_c \nabla^2 \log f_c(\mathbf{x})) \right). \quad (57)$$

Let $R_c := R_c(\mathbf{X}_{[0,T]}^{(c)})$ denote a compact subset of \mathbb{R}^d for which $\mathbf{X}_t^{(c)}$ is constrained in time $[0, T]$ for $c \in \mathcal{C}$, then to bound the first term in (57), we first use the triangle inequality by noting

$$\begin{aligned} \max_{\mathbf{x} \in R_c} \left\| \mathbf{\Lambda}_c^{\frac{1}{2}} \nabla \log f_c(\mathbf{x}) \right\| &= \max_{\mathbf{x} \in R_c} \left\| \mathbf{\Lambda}_c^{\frac{1}{2}} \nabla \log f_c(\hat{\mathbf{x}}^{(c)}) + \mathbf{\Lambda}_c^{\frac{1}{2}} \left(\nabla \log f_c(\mathbf{x}) - \nabla \log f_c(\hat{\mathbf{x}}^{(c)}) \right) \right\| \\ &\leq \left\| \mathbf{\Lambda}_c^{\frac{1}{2}} \nabla \log f_c(\hat{\mathbf{x}}^{(c)}) \right\| + \max_{\mathbf{x} \in R_c} \left\| \mathbf{\Lambda}_c^{\frac{1}{2}} \left(\nabla \log f_c(\mathbf{x}) - \nabla \log f_c(\hat{\mathbf{x}}^{(c)}) \right) \right\|, \end{aligned} \quad (58)$$

where $\hat{\mathbf{x}}^{(c)}$ is a user-specified point in \mathbb{R}^d . Focusing now on bounding the second term of (58), then we express this as a line integral between \mathbf{x} and $\hat{\mathbf{x}}^{(c)}$ so

$$\max_{\mathbf{x} \in R_c} \left\| \mathbf{\Lambda}_c^{\frac{1}{2}} \left(\nabla \log f_c(\mathbf{x}) - \nabla \log f_c(\hat{\mathbf{x}}^{(c)}) \right) \right\| = \max_{\mathbf{x} \in R_c} \left\| \mathbf{\Lambda}_c^{-\frac{1}{2}} \int_0^{\|\mathbf{x} - \hat{\mathbf{x}}^{(c)}\|} \mathbf{\Lambda}_c \nabla^2 \log f(\mathbf{x} + u\mathbf{n}) \mathbf{n} \, du \right\|,$$

where $\mathbf{u} = \mathbf{x} + u\mathbf{n}$, where \mathbf{n} is a unit-vector. We have

$$\begin{aligned} \max_{\mathbf{x} \in R_c} \left\| \Lambda_c^{-\frac{1}{2}} \int_0^{\|\mathbf{x} - \hat{\mathbf{x}}^{(c)}\|} \Lambda_c \nabla \log f(\mathbf{x} + u\mathbf{n}) \mathbf{n} \, du \right\| &\leq \max_{\mathbf{x} \in R_c} \left\| \Lambda_c^{-\frac{1}{2}} (\mathbf{x} - \hat{\mathbf{x}}^{(c)}) \right\| \\ &\quad \cdot \sup_{\mathbf{n}; \mathbf{x} \in R_c} \left\| \Lambda_c \nabla^2 \log f(\mathbf{x} + u\mathbf{n}) \mathbf{n} \right\| \\ &\leq \max_{\mathbf{x} \in R_c} \left\| \Lambda_c^{-\frac{1}{2}} (\mathbf{x} - \hat{\mathbf{x}}^{(c)}) \right\| \cdot P^{\Lambda_c}, \end{aligned} \quad (59)$$

where P^{Λ_c} is defined in (19). Putting together (58) and (59), we have

$$\max_{\mathbf{x} \in R_c} \left\| \Lambda_c^{\frac{1}{2}} \nabla \log f_c(\mathbf{x}) \right\| \leq \left\| \Lambda_c^{\frac{1}{2}} \nabla \log f_c(\hat{\mathbf{x}}^{(c)}) \right\| + \max_{\mathbf{x} \in R_c} \left\| \Lambda_c^{-\frac{1}{2}} (\mathbf{x} - \hat{\mathbf{x}}^{(c)}) \right\| \cdot P^{\Lambda_c}.$$

Since for a matrix $A \in \mathbb{R}^d$, $\text{Tr}(A) \leq d \cdot \gamma(A)$, we can bound the second term in (57) as follows:

$$\max_{\mathbf{x} \in R_c} |\text{Tr}(\Lambda_c \nabla^2 \log f_c(\mathbf{x}))| \leq d \cdot P^{\Lambda_c},$$

and hence we can bound ϕ_c as follows:

$$\max_{\mathbf{x} \in R_c} |\phi_c(\mathbf{x})| \leq \frac{1}{2} \left[\left(\left\| \Lambda_c^{\frac{1}{2}} \nabla \log f_c(\hat{\mathbf{x}}^{(c)}) \right\| + \max_{\mathbf{x} \in R_c} \left\| \Lambda_c^{-\frac{1}{2}} (\mathbf{x} - \hat{\mathbf{x}}^{(c)}) \right\| \cdot P^{\Lambda_c} \right)^2 + d \cdot P^{\Lambda_c} \right].$$

Noting that in (57) that the first term is squared, then the lower and upper bounds of $\phi_c(\mathbf{x})$ for $\mathbf{x} \in R_c$ are given by (17) and (18) respectively. \blacksquare

Appendix E. Proof of Theorem 4

Proof Following in the style of Beskos et al. (2006, 2008); Fearnhead et al. (2008) and Dai et al. (2023, Appendix B), for $j = 1, \dots, n$, we have

$$\begin{aligned} &\mathbb{E}_{\bar{\mathcal{R}}} \mathbb{E}_{\bar{\mathcal{W}}|\bar{\mathcal{R}}} \mathbb{E}_{\bar{\mathcal{K}}} \mathbb{E}_{\bar{\mathcal{U}}} [a_j \tilde{\rho}_j] \\ &= \mathbb{E}_{\bar{\mathcal{R}}} \mathbb{E}_{\bar{\mathcal{W}}|\bar{\mathcal{R}}} \mathbb{E}_{\bar{\mathcal{K}}} \mathbb{E}_{\bar{\mathcal{U}}} \left[\prod_{c \in \mathcal{C}} \left(\frac{\Delta_j^{\kappa_c} \cdot e^{-(U_j^{(c)} - \Phi_c) \Delta_j}}{\kappa_c! \cdot p(\kappa_c | R_c)} \prod_{k_c=1}^{\kappa_c} (U_j^{(c)} - \phi_c(\mathbf{X}_{\xi_c, k_c}^{(c)})) \right) \right] \\ &= \mathbb{E}_{\bar{\mathcal{R}}} \mathbb{E}_{\bar{\mathcal{W}}|\bar{\mathcal{R}}} \mathbb{E}_{\bar{\mathcal{K}}} \left[\prod_{c \in \mathcal{C}} \left(\frac{\Delta_j^{\kappa_c} \cdot e^{-(U_j^{(c)} - \Phi_c) \Delta_j}}{\kappa_c! \cdot p(\kappa_c | R_c)} \cdot \left[\int_{t_{j-1}}^{t_j} \frac{U_j^{(c)} - \phi_c(\mathbf{X}_t^{(c)})}{\Delta_j} dt \right]^{\kappa_c} \right) \right] \\ &= \mathbb{E}_{\bar{\mathcal{R}}} \mathbb{E}_{\bar{\mathcal{W}}|\bar{\mathcal{R}}} \left[\prod_{c \in \mathcal{C}} \left(\sum_{k_c=0}^{\infty} \frac{\Delta_j^{k_c} \cdot e^{-(U_j^{(c)} - \Phi_c) \Delta_j}}{k_c! \cdot p(k_c | R_c)} \cdot \left[\int_{t_{j-1}}^{t_j} \frac{U_j^{(c)} - \phi_c(\mathbf{X}_t^{(c)})}{\Delta_j} dt \right]^{k_c} \right) \right] \\ &= \mathbb{E}_{\bar{\mathcal{R}}} \mathbb{E}_{\bar{\mathcal{W}}|\bar{\mathcal{R}}} \left[\prod_{c \in \mathcal{C}} e^{-(U_j^{(c)} - \Phi_c) \Delta_j} \cdot \left(\sum_{k_c=0}^{\infty} \frac{\Delta_j^{k_c}}{k_c! \cdot p(k_c | R_c)} \cdot \left[\int_{t_{j-1}}^{t_j} \frac{U_j^{(c)} - \phi_c(\mathbf{X}_t^{(c)})}{\Delta_j} dt \right]^{k_c} \right) \right] \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E}_{\tilde{\mathcal{R}}|\mathcal{R}} \mathbb{E}_{\mathbb{W}|\tilde{\mathcal{R}}} \left[\prod_{c \in \mathcal{C}} e^{-(U_j^{(c)} - \Phi_c) \Delta_j} \cdot \exp \left\{ \int_{t_{j-1}}^{t_j} \left(U_j^{(c)} - \phi_c \left(\mathbf{X}_t^{(c)} \right) \right) dt \right\} \right] \\
 &= \prod_{c \in \mathcal{C}} \mathbb{E}_{\mathbb{W}_{\Lambda_c, j}} \left[\exp \left\{ - \int_{t_{j-1}}^{t_j} \left(\phi_c \left(\mathbf{X}_t^{(c)} \right) - \Phi_c \right) dt \right\} \right] =: \rho_j,
 \end{aligned}$$

and hence $a_j \tilde{\rho}_j$ is an unbiased estimator for ρ_j . \blacksquare

Appendix F. Unbiased Estimation of ρ_j

Computing $\tilde{\rho}_1^{(a)}$ and $\tilde{\rho}_1^{(b)}$ by means of *layer information* in the case where $\Lambda_c = \mathbb{I}_d$ is detailed explicitly in Dai et al. (2023, Algorithm 4, Appendix B). In the case where $\Lambda_c \neq \mathbb{I}_d$, we simulate layers by appealing to a suitable transformation. In particular, we transform the start and end points of the Brownian bridge with transformation matrix $\Lambda_c^{-\frac{1}{2}}$, letting $\mathbf{z}_{j-1}^{(c)} := \Lambda_c^{-\frac{1}{2}} \mathbf{x}_{j-1}^{(c)}$ and $\mathbf{z}_j^{(c)} := \Lambda_c^{-\frac{1}{2}} \mathbf{x}_j^{(c)}$. The resulting Brownian bridge sample path, $\mathbf{z}_t^{(c)} := \Lambda_c^{-\frac{1}{2}} \mathbf{X}_t^{(c)}$, has identity covariance structure and thus we can use existing methods for simulating *layered* Brownian bridge sample paths $\mathbf{z}_t^{(c)}$ with law $\mathbb{W}_{\mathbb{I}_d}$ from $\mathbf{z}_{j-1}^{(c)}$ to $\mathbf{z}_j^{(c)}$. By finding a bounding hyper cube for the reverse transformed bounds, we are able to find appropriate layer information for the case $\Lambda_c \neq \mathbb{I}_d$. We are now able with minimal modification to apply the approach of Dai et al. (2023), as given in Algorithm 5.

Appendix G. Proof of Theorem 12, Corollary 13, Theorem 15, Proposition 19

Proof (Theorem 12) Considering the initial conditional effective sample size, CESS_0 , we have

$$\begin{aligned}
 N^{-1} \text{CESS}_0 &:= N^{-1} \left[\frac{\left(\sum_{i=1}^N \rho_{0,i} \right)^2}{\sum_{i=1}^N \rho_{0,i}^2} \right] \rightarrow \frac{(\mathbb{E}[\rho_{0,i}])^2}{\mathbb{E}[\rho_{0,i}^2]} \\
 &= \frac{\mathbb{E} \left[\exp \left\{ - \sum_{c \in \mathcal{C}} \frac{(\tilde{\mathbf{x}}_0^{(c)} - \mathbf{x}_0^{(c)})^\top \Lambda_c^{-1} (\tilde{\mathbf{x}}_0^{(c)} - \mathbf{x}_0^{(c)})}{2T} \right\} \right]^2}{\mathbb{E} \left[\exp \left\{ - \sum_{c \in \mathcal{C}} \frac{(\tilde{\mathbf{x}}_0^{(c)} - \mathbf{x}_0^{(c)})^\top \Lambda_c^{-1} (\tilde{\mathbf{x}}_0^{(c)} - \mathbf{x}_0^{(c)})}{T} \right\} \right]} \\
 &= \frac{\mathbb{E} \left[e^{-\frac{|\mathcal{C}| \sigma^2}{2T}} \right]^2}{\mathbb{E} \left[e^{-\frac{|\mathcal{C}| \sigma^2}{T}} \right]}, \tag{60}
 \end{aligned}$$

where $\sigma^2 := \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\tilde{\mathbf{x}}_0^{(c)} - \mathbf{x}_0^{(c)})^\top \Lambda_c^{-1} (\tilde{\mathbf{x}}_0^{(c)} - \mathbf{x}_0^{(c)})$ where $\mathbf{x}_0^{(c)} \sim \mathcal{N}_d(\mathbf{a}_c, \frac{b|\mathcal{C}|}{m} \Lambda_c)$. To get an expression for $N^{-1} \text{CESS}_0$, we begin by obtaining the moment generating function (mgf)

Algorithm 5 Simulating $\tilde{\rho}_j$.

 1. For $c \in \mathcal{C}$

- (a) $\mathbf{z}_{j-1}^{(c)}, \mathbf{z}_j^{(c)}$: Transform the path, setting $\mathbf{z}_{j-1}^{(c)} := \mathbf{\Lambda}_c^{-\frac{1}{2}} \mathbf{x}_{j-1}^{(c)}$, and $\mathbf{z}_j^{(c)} := \mathbf{\Lambda}_c^{-\frac{1}{2}} \mathbf{x}_j^{(c)}$.
- (b) R_c : Set $R_c := \mathbf{\Lambda}_c^{\frac{1}{2}} R_c^{(z)}$, where $R_c^{(z)} \sim \mathcal{R}_c^{(z)}$ as per Pollock et al. (2016, Algorithm 14).
- (c) $L_X^{(c)}, U_X^{(c)}$: Compute lower and upper bounds, $L_X^{(c)}$ and $U_X^{(c)}$, of $\phi_c(\mathbf{x})$ for $\mathbf{x} \in R_c$ (as per (17) and (18), or otherwise).
- (d) p_c : Choose $p(\cdot|R_c)$ using either GPE-1 (Condition 5) or GPE-2 (Condition 6).
- (e) κ_c, ξ : Simulate $\kappa_c \sim p(\cdot|R_c)$, and simulate $\xi_{c,1}, \dots, \xi_{c,\kappa_c} \sim \mathcal{U}[t_{j-1}, t_j]$.
- (f) $\mathbf{z}^{(c)}$: Simulate $\mathbf{z}_{\xi_{c,1}}^{(c)}, \dots, \mathbf{z}_{\xi_{c,\kappa_c}}^{(c)} \sim \mathbb{W}_{\mathbb{I}_d} | R_c^{(z)}$ as per Pollock et al. (2016, Algorithm 15).
- (g) $\mathbf{X}^{(c)}$: Reverse transform the path, setting $\mathbf{X}_{\xi_{c,k_c}}^{(c)} = \mathbf{\Lambda}_c^{\frac{1}{2}} \mathbf{z}_{\xi_{c,k_c}}^{(c)}$ for $k_c \in \{1, \dots, \kappa_c\}$.

2. Output:

$$\tilde{\rho}_j = \prod_{c \in \mathcal{C}} \left[\frac{\Delta_j^{\kappa_c} \cdot e^{-U_j^{(c)} \Delta_j}}{\kappa_c! \cdot p(\kappa_c | R_c)} \prod_{k_c=1}^{\kappa_c} \left(U_j^{(c)} - \phi_c \left(\mathbf{X}_{\xi_{c,k_c}}^{(c)} \right) \right) \right].$$

 for σ^2 . First note

$$\frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\mathbf{x}_0^{(c)} - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\mathbf{x}_0^{(c)} - \tilde{\mathbf{a}}) = \sigma^2 + \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}}). \quad (61)$$

Considering the term $\frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\mathbf{x}_0^{(c)} - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\mathbf{x}_0^{(c)} - \tilde{\mathbf{a}})$ and letting $\mathbf{Y}_c := \mathbf{\Lambda}_c^{-\frac{1}{2}} (\mathbf{x}_0^{(c)} - \tilde{\mathbf{a}})$, then \mathbf{Y}_c has mean $\mathbf{\Lambda}_c^{-\frac{1}{2}} (\mathbf{a}_c - \tilde{\mathbf{a}})$ and variance $\frac{b|\mathcal{C}|}{m} \mathbb{I}_d$. Hence $\sqrt{\frac{m}{b|\mathcal{C}|}} \mathbf{Y}_c$ has mean $\sqrt{\frac{m}{b|\mathcal{C}|}} \mathbf{\Lambda}_c^{-\frac{1}{2}} (\mathbf{a}_c - \tilde{\mathbf{a}})$ and variance \mathbb{I}_d , and so let

$$\lambda = \sum_{c \in \mathcal{C}} \left\| \sqrt{\frac{m}{b|\mathcal{C}|}} \mathbf{\Lambda}_c^{-\frac{1}{2}} (\mathbf{a}_c - \tilde{\mathbf{a}}) \right\|^2 = \frac{m}{b|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\mathbf{a}_c - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\mathbf{a}_c - \tilde{\mathbf{a}}) = \frac{m}{b} \sigma_{\mathbf{a}}^2,$$

then $\frac{m}{b|\mathcal{C}|} \sum_{c \in \mathcal{C}} \|\mathbf{Y}_c\|^2 \sim \chi^2(|\mathcal{C}|d, \lambda)$ distribution (i.e. $\frac{m}{b|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\mathbf{x}_0^{(c)} - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\mathbf{x}_0^{(c)} - \tilde{\mathbf{a}})$ has a non-central $\chi^2(|\mathcal{C}|d, \lambda)$ distribution) with mgf

$$M_1(s) := \frac{\exp \left\{ \frac{\lambda s}{1-2s} \right\}}{(1-2s)^{\frac{|\mathcal{C}|d}{2}}}. \quad (62)$$

Secondly, consider $\frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}}) = \frac{1}{|\mathcal{C}|} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}})$, where $\mathbf{\Lambda}_c^{-1} := \sum_{c \in \mathcal{C}} \mathbf{\Lambda}_c^{-1}$. Then since $\tilde{\mathbf{x}}_0^{(c)} \sim \mathcal{N}_d(\tilde{\mathbf{a}}, \frac{b|\mathcal{C}|}{m} \mathbf{\Lambda}_c)$, then $\mathbf{Z} := \sqrt{\frac{m}{b|\mathcal{C}|}} \mathbf{\Lambda}_c^{-\frac{1}{2}} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}}) \sim$

$\mathcal{N}_d(\mathbf{0}, \mathbb{I}_d)$ and so $\|\mathbf{Z}\|^2 \sim \chi_d^2$ (i.e. $\frac{m}{b|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}})$ has χ_d^2 distribution) with mgf

$$M_2(s) := (1 - 2s)^{-\frac{d}{2}}. \quad (63)$$

From (61), we have

$$\sigma^2 = \frac{b}{m} \underbrace{\left[\frac{m}{b|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\mathbf{x}_0^{(c)} - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\mathbf{x}_0^{(c)} - \tilde{\mathbf{a}}) \right]}_{\sim \chi^2(|\mathcal{C}|d, \lambda)} - \frac{b}{m} \underbrace{\left[\frac{m}{b|\mathcal{C}|} \sum_{c \in \mathcal{C}} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}})^\top \mathbf{\Lambda}_c^{-1} (\tilde{\mathbf{x}}_0^{(c)} - \tilde{\mathbf{a}}) \right]}_{\sim \chi_d^2},$$

where $\lambda = \frac{m\sigma_{\mathbf{a}}^2}{b}$. Therefore, using (62) and (63), the mgf for σ^2 is given by

$$M_{\sigma^2}(s) = \frac{M_1\left(\frac{sb}{m}\right)}{M_2\left(\frac{sb}{m}\right)} = \exp\left\{\frac{m\sigma_{\mathbf{a}}^2 s}{m - 2sb}\right\} \cdot \left(1 - 2\frac{sb}{m}\right)^{-\frac{(|\mathcal{C}|-1)d}{2}}, \quad \text{where } \frac{sb}{m} < \frac{1}{2}. \quad (64)$$

Given the mgf of σ^2 , then

$$\begin{aligned} N^{-1}\text{CESS}_0 &\rightarrow \frac{\mathbb{E}\left[e^{-\frac{|\mathcal{C}|\sigma^2}{2T}}\right]^2}{\mathbb{E}\left[e^{-\frac{|\mathcal{C}|\sigma^2}{T}}\right]} = \frac{M_{\sigma^2}\left(-\frac{|\mathcal{C}|}{2T}\right)^2}{M_{\sigma^2}\left(-\frac{|\mathcal{C}|}{T}\right)} \\ &= \frac{\left[\exp\left\{\frac{m\sigma_{\mathbf{a}}^2\left(-\frac{|\mathcal{C}|}{2T}\right)}{m-2\left(-\frac{|\mathcal{C}|b}{2T}\right)}\right\} \cdot \left(1 - 2\left(-\frac{|\mathcal{C}|}{2T}\right)\frac{b}{m}\right)^{-\frac{(|\mathcal{C}|-1)d}{2}}\right]^2}{\exp\left\{\frac{m\sigma_{\mathbf{a}}^2\left(-\frac{|\mathcal{C}|}{T}\right)}{m-2\left(-\frac{|\mathcal{C}|b}{T}\right)}\right\} \cdot \left(1 - 2\left(-\frac{|\mathcal{C}|}{T}\right)\frac{b}{m}\right)^{-\frac{(|\mathcal{C}|-1)d}{2}}}\right)} \\ &= \frac{\exp\left\{-\frac{m\sigma_{\mathbf{a}}^2\left(\frac{|\mathcal{C}|}{T}\right)}{m+\frac{|\mathcal{C}|b}{T}}\right\} \cdot \left(1 + \frac{|\mathcal{C}|b}{Tm}\right)^{-(|\mathcal{C}|-1)d}}{\exp\left\{-\frac{m\sigma_{\mathbf{a}}^2\left(\frac{|\mathcal{C}|}{T}\right)}{m+2\left(\frac{|\mathcal{C}|b}{T}\right)}\right\} \cdot \left(1 + 2\left(\frac{|\mathcal{C}|b}{Tm}\right)\right)^{-\frac{(|\mathcal{C}|-1)d}{2}}}\right)} \\ &= \exp\left\{-\frac{\sigma_{\mathbf{a}}^2}{\frac{T}{|\mathcal{C}|} + \frac{b}{m}}\right\} \cdot \exp\left\{\frac{\sigma_{\mathbf{a}}^2}{\frac{T}{|\mathcal{C}|} + \frac{2b}{m}}\right\} \cdot \left[\frac{\left(1 + \frac{|\mathcal{C}|b}{Tm}\right)^2}{1 + 2\left(\frac{|\mathcal{C}|b}{Tm}\right)}\right]^{-\frac{(|\mathcal{C}|-1)d}{2}} \\ &= \exp\left\{-\frac{\sigma_{\mathbf{a}}^2\left(\frac{b}{m}\right)}{\left(\frac{T}{|\mathcal{C}|} + \frac{b}{m}\right)\left(\frac{T}{|\mathcal{C}|} + \frac{2b}{m}\right)}\right\} \cdot \left[1 + \frac{\left(\frac{|\mathcal{C}|b}{Tm}\right)^2}{1 + \frac{2|\mathcal{C}|b}{Tm}}\right]^{-\frac{(|\mathcal{C}|-1)d}{2}}, \end{aligned}$$

and so Theorem 12 immediately follows. \blacksquare

Proof (Corollary 13) Under Condition 8, $\sigma_{\mathbf{a}}^2 = \frac{(|\mathcal{C}|-1)\lambda}{m} < \frac{|\mathcal{C}|\lambda}{m}$, so for the first term in (31),

$$\exp\left\{-\frac{\sigma_{\mathbf{a}}^2\left(\frac{b}{m}\right)}{\left(\frac{T}{|\mathcal{C}|} + \frac{b}{m}\right)\left(\frac{T}{|\mathcal{C}|} + \frac{2b}{m}\right)}\right\} \geq \exp\left\{-\frac{\sigma_{\mathbf{a}}^2 b |\mathcal{C}|^2}{T^2 m}\right\}$$

$$\begin{aligned}
 &\geq \exp \left\{ -\frac{b^2 |\mathcal{C}|^3 \lambda}{T^2 m^2} \right\} \\
 &\geq \exp \left\{ -\frac{\lambda}{k_1^2} \right\}, \tag{65}
 \end{aligned}$$

where $T \geq \frac{b|\mathcal{C}|^{3/2}k_1}{m}$ for some constant $k_1 > 0$, and for the second term in (31), then

$$\begin{aligned}
 \left[1 + \frac{\left(\frac{|\mathcal{C}|b}{Tm}\right)^2}{1 + \frac{2|\mathcal{C}|b}{Tm}} \right]^{-\frac{(|\mathcal{C}|-1)d}{2}} &\geq \left[\exp \left\{ \frac{\left(\frac{|\mathcal{C}|b}{Tm}\right)^2}{1 + \frac{2|\mathcal{C}|b}{Tm}} \right\} \right]^{-\frac{(|\mathcal{C}|-1)d}{2}} \\
 &= \exp \left\{ -\frac{\left(\frac{|\mathcal{C}|b}{Tm}\right)^2 (|\mathcal{C}|-1)d}{2\left(1 + \frac{2|\mathcal{C}|b}{Tm}\right)} \right\} \\
 &\geq \exp \left\{ -\frac{\left(\frac{|\mathcal{C}|^3 b^2}{T^2 m^2}\right) d}{2} \right\} \\
 &\geq \exp \left\{ -\frac{d}{2k_1^2} \right\}, \tag{66}
 \end{aligned}$$

with $T \geq \frac{b|\mathcal{C}|^{3/2}k_1}{m}$. Hence, under Condition 8 and choosing $T \geq \frac{b|\mathcal{C}|^{3/2}k_1}{m}$, combining the bounds from (65) and (66) gives (32). Under Condition 10, $\sigma_{\mathbf{a}}^2 = b\gamma$, if we assume $T \geq \frac{b|\mathcal{C}|^{3/2}k_1}{m}$ for some constant $k_1 > 0$, and $T \geq |\mathcal{C}|^{\frac{1}{2}}k_2$ for some constant $k_2 > 0$, then

$$\left(\frac{T}{|\mathcal{C}|} + \frac{b}{m} \right) \left(\frac{T}{|\mathcal{C}|} + \frac{2b}{m} \right) \geq \frac{T^2}{|\mathcal{C}|^2} \geq \frac{bk_1 k_2}{m},$$

and so we have

$$\exp \left\{ -\frac{\sigma_{\mathbf{a}}^2 \left(\frac{b}{m}\right)}{\left(\frac{T}{|\mathcal{C}|} + \frac{b}{m}\right) \left(\frac{T}{|\mathcal{C}|} + \frac{2b}{m}\right)} \right\} \geq \exp \left\{ -\frac{\frac{b^2 \gamma}{m}}{\frac{bk_1 k_2}{m}} \right\} = \exp \left\{ -\frac{b\gamma}{k_1 k_2} \right\}. \tag{67}$$

Hence, under Condition 10 and choosing T such that $T \geq \frac{b|\mathcal{C}|^{3/2}k_1}{m}$ and $T \geq |\mathcal{C}|^{\frac{1}{2}}k_2$, we can combine the bounds from (67) and (66) to obtain the bound in (33). \blacksquare

Proof (Theorem 15) As $N \rightarrow \infty$, we have

$$N^{-1} \text{CESS}_j := N^{-1} \left[\frac{\left(\sum_{i=1}^N \tilde{\rho}_{j,i}\right)^2}{\sum_{i=1}^N \tilde{\rho}_{j,i}^2} \right] = \left[\frac{\left(N^{-1} \sum_{i=1}^N a_j \tilde{\rho}_{j,i}\right)^2}{N^{-1} \sum_{i=1}^N (a_j \tilde{\rho}_{j,i})^2} \right] \rightarrow \frac{\mathbb{E}[a_j \tilde{\rho}_j]^2}{\mathbb{E}[(a_j \tilde{\rho}_j)^2]},$$

where $a_j := \exp\{\sum_{c \in \mathcal{C}} \Phi_c \Delta_j\}$. Since $a_j \tilde{\rho}_j$ is an unbiased estimate of ρ_j (see Theorem 4), then

$$\mathbb{E}[a_j \tilde{\rho}_j] = \prod_{c \in \mathcal{C}} \mathbb{E}_{\mathbb{W}_{\Lambda_c, j}} \left(\exp \left\{ -\int_{t_{j-1}}^{t_j} \left(\phi_c(\mathbf{X}_t^{(c)}) - \Phi_c \right) \right\} \right)$$

$$= \mathbb{E}_{\bar{\mathbb{W}}_{\Lambda}} \left(\exp \left\{ - \sum_{c \in \mathcal{C}} \int_{t_{j-1}}^{t_j} \phi_c \left(\mathbf{X}_t^{(c)} \right) dt \right\} \right) \cdot a_j$$

where $\bar{\mathbb{W}}_{\Lambda}$ denotes the law of the collection of Brownian bridges $\{\mathbb{W}_{\Lambda_c, j} : c \in \mathcal{C}\}$ for each j . Note that under the optimal distribution for $p(\kappa_c | R_c)$ (a Poisson distribution with intensity given in (35)), then $\mathbb{E} \left[(a_j \tilde{\rho}_j)^2 \right] \leq 1$ (Fearnhead et al., 2008; Dai et al., 2023), so

$$\lim_{N \rightarrow \infty} N^{-1} \text{CESS}_j \geq \mathbb{E} [a_j \tilde{\rho}_j]^2 = \left[\mathbb{E}_{\bar{\mathbb{W}}_{\Lambda}} \left(\exp \left\{ - \sum_{c \in \mathcal{C}} \int_{t_{j-1}}^{t_j} \phi_c \left(\mathbf{X}_t^{(c)} \right) dt \right\} \right) \right]^2 \cdot a_j^2.$$

If $f_c \sim \mathcal{N}_d(\mathbf{a}_c, \frac{b|\mathcal{C}|}{m} \mathbf{\Lambda}_c)$, then $\phi_c(\mathbf{x}) = \frac{1}{2} \left(\left(\frac{m}{b|\mathcal{C}|} \right)^2 (\mathbf{x} - \mathbf{a}_c)^\top \mathbf{\Lambda}_c^{-1} (\mathbf{x} - \mathbf{a}_c) - \frac{md}{b|\mathcal{C}|} \right)$ which has global lower bound $\Phi_c = -\frac{1}{2} \left(\frac{md}{b|\mathcal{C}|} \right)$ (since the minimum of ϕ_c occurs at the mean, \mathbf{a}_c). Then by considering small intervals (t_{j-1}, t_j) and taking the limit of $\Delta_j := t_j - t_{j-1} \rightarrow 0$, then

$$\begin{aligned} & \lim_{\Delta_j \rightarrow 0} \lim_{N \rightarrow \infty} N^{-1} \text{CESS}_j \\ & \geq \lim_{\Delta_j \rightarrow 0} \left[\mathbb{E} \left(\mathbb{E} \left\{ \mathbb{E} \left(\exp \left\{ - \sum_{c \in \mathcal{C}} \int_{t_{j-1}}^{t_j} \phi_c \left(\mathbf{X}_t^{(c)} \right) dt \right\} \middle| \boldsymbol{\xi}_j, \bar{\mathbf{x}}_{j-1}^{(c)} \right\} \middle| \bar{\mathbf{x}}_{j-1}^{(c)} \right) \right) \right]^2 \cdot a_j^2 \\ & \geq \lim_{\Delta_j \rightarrow 0} \left[\mathbb{E} \left(\mathbb{E} \left\{ \mathbb{E} \left(\exp \left\{ - \frac{\Delta_j}{2} \sum_{c \in \mathcal{C}} \left(\frac{m}{b|\mathcal{C}|} \right)^2 (\mathbf{x}_j^{(c)} - \mathbf{a}_c)^\top \mathbf{\Lambda}_c^{-1} (\mathbf{x}_j^{(c)} - \mathbf{a}_c) \right\} \middle| \boldsymbol{\xi}_j, \bar{\mathbf{x}}_{j-1}^{(c)} \right\} \middle| \bar{\mathbf{x}}_{j-1}^{(c)} \right) \right) \right]^2 \\ & \geq \left[\mathbb{E} \left(\mathbb{E} \left\{ \lim_{\Delta_j \rightarrow 0} \mathbb{E} \left(\exp \left\{ - \frac{\Delta_j}{2} \sum_{c \in \mathcal{C}} \left(\frac{m}{b|\mathcal{C}|} \right)^2 (\mathbf{x}_j^{(c)} - \mathbf{a}_c)^\top \mathbf{\Lambda}_c^{-1} (\mathbf{x}_j^{(c)} - \mathbf{a}_c) \right\} \middle| \boldsymbol{\xi}_j, \bar{\mathbf{x}}_{j-1}^{(c)} \right\} \middle| \bar{\mathbf{x}}_{j-1}^{(c)} \right) \right) \right]^2, \end{aligned}$$

(by using a trapezoidal rule approximation of the integral and exploiting the use of small intervals) where $\lim_{\Delta_j \rightarrow 0}$ and expectations are exchanged using the dominated convergence theorem (as the exponential term is bounded above by 1 and its expectation exists (Dai et al., 2023, Appendix C)).

From (75) in Corollary 23, we note that $\mathbf{x}_j^{(c)}$ only depends $\mathbf{x}_{j-1}^{(c)}$ through $\boldsymbol{\xi}_j$ and $\boldsymbol{\zeta}_j^{(c)}$ for all $c \in \mathcal{C}$, and we have

$$\mathbf{x}_j^{(c)} \middle| \boldsymbol{\xi}_j, \bar{\mathbf{x}}_{j-1}^{(c)} \sim \mathcal{N}_d \left(\mathbb{E} \left[\mathbf{x}_j^{(c)} \middle| \boldsymbol{\xi}_j, \bar{\mathbf{x}}_{j-1}^{(c)} \right], \frac{T - t_j}{T - t_{j-1}} \Delta_j \mathbf{\Lambda}_c \right),$$

and consequently,

$$\left(\frac{T - t_j}{T - t_{j-1}} \Delta_j \right)^{-1} \sum_{c \in \mathcal{C}} (\mathbf{x}_j^{(c)} - \mathbf{a}_c)^\top \mathbf{\Lambda}_c^{-1} (\mathbf{x}_j^{(c)} - \mathbf{a}_c) \sim \chi^2(|\mathcal{C}|d, \lambda'_j),$$

with moment generating function $M_j(s) := \exp \left\{ \frac{\lambda'_j s}{1 - 2s} \right\} \cdot (1 - 2s)^{-\frac{|\mathcal{C}|d}{2}}$, where

$$\lambda'_j = \left(\frac{T - t_j}{T - t_{j-1}} \Delta_j \right)^{-1} \sum_{c \in \mathcal{C}} \left(\mathbb{E} \left[\mathbf{x}_j^{(c)} \middle| \boldsymbol{\xi}_j, \bar{\mathbf{x}}_{j-1}^{(c)} \right] - \mathbf{a}_c \right)^\top \mathbf{\Lambda}_c^{-1} \left(\mathbb{E} \left[\mathbf{x}_j^{(c)} \middle| \boldsymbol{\xi}_j, \bar{\mathbf{x}}_{j-1}^{(c)} \right] - \mathbf{a}_c \right)$$

$$= \left(\frac{T - t_j}{T - t_{j-1}} \Delta_j \right)^{-1} |\mathcal{C}| \sigma_{t_j}^2,$$

with

$$\sigma_{t_j}^2 := \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \left(\mathbb{E} \left[\mathbf{x}_j^{(c)} \mid \boldsymbol{\xi}_j, \bar{\mathbf{x}}_{j-1}^{(c)} \right] - \mathbf{a}_c \right)^\top \boldsymbol{\Lambda}_c^{-1} \left(\mathbb{E} \left[\mathbf{x}_j^{(c)} \mid \boldsymbol{\xi}_j, \bar{\mathbf{x}}_{j-1}^{(c)} \right] - \mathbf{a}_c \right).$$

Letting $s = -\frac{1}{2} \left(\frac{m}{b|\mathcal{C}|} \right)^2 \left(\frac{T - t_j}{T - t_{j-1}} \right) \Delta_j^2$, then

$$\begin{aligned} \lim_{\Delta_j \rightarrow 0} \lim_{N \rightarrow \infty} N^{-1} \text{CESS}_j &\geq \left[\mathbb{E} \left(\mathbb{E} \left\{ \lim_{\Delta_j \rightarrow 0} \exp \left\{ \frac{\lambda_j s}{1 - 2s} \right\} \mid \bar{\mathbf{x}}_{j-1}^{(c)} \right\} \right) \right]^2 \cdot (1 - 2s)^{-|\mathcal{C}|d} \\ &\geq \left[\mathbb{E} \left(\mathbb{E} \left\{ \lim_{\Delta_j \rightarrow 0} \exp \left\{ \frac{-\frac{1}{2} \left(\frac{m^2}{b^2|\mathcal{C}|} \right) \sigma_{t_j}^2 \Delta_j}{1 - 2s} \right\} \mid \bar{\mathbf{x}}_{j-1}^{(c)} \right\} \right) \right]^2 \cdot (1 - 2s)^{-|\mathcal{C}|d}. \end{aligned}$$

From (75), we have

$$\mathbb{E} \left[\mathbf{x}_j^{(c)} \mid \boldsymbol{\xi}_j, \bar{\mathbf{x}}_{j-1}^{(c)} \right] = \left[\frac{\Delta_j^2}{T - t_{j-1}} \right]^{\frac{1}{2}} \boldsymbol{\xi}_j + \frac{T - t_j}{T - t_{j-1}} \mathbf{x}_{j-1}^{(c)} + \frac{t_j - t_{j-1}}{T - t_{j-1}} \tilde{\mathbf{x}}_{j-1},$$

and so we have $\lim_{\Delta_j \rightarrow 0} \sigma_{t_j}^2 =: \nu_j$ where ν_j is given in (37). Using Jensen's inequality, we can get

$$\begin{aligned} \lim_{\Delta_j \rightarrow 0} \lim_{N \rightarrow \infty} N^{-1} \text{CESS}_j &\geq \lim_{\Delta_j \rightarrow 0} \left[\exp \left\{ \frac{-\frac{1}{2} \mathbb{E}[\nu_j] \left(\frac{m^2}{b^2|\mathcal{C}|} \right) \Delta_j}{1 - 2s} \right\} \right]^2 \cdot (1 - 2s)^{-|\mathcal{C}|d} \\ &\geq \lim_{\Delta_j \rightarrow 0} \exp \left\{ \frac{-\mathbb{E}[\nu_j] \left(\frac{m^2}{b^2|\mathcal{C}|} \right) \Delta_j}{1 - 2s} \right\} \cdot (1 - 2s)^{-|\mathcal{C}|d}. \end{aligned} \quad (68)$$

Consider the first term in (68), then taking the limit $\Delta_j \rightarrow 0$ implies that $s \rightarrow 0$, and if $\Delta_j \leq \frac{b^2|\mathcal{C}|k_3}{\mathbb{E}[\nu_j]m^2}$ for some $k_3 > 0$, then

$$\exp \left\{ \frac{-\mathbb{E}[\nu_j] \left(\frac{m^2}{b^2|\mathcal{C}|} \right) \Delta_j}{1 - 2s} \right\} \geq \exp \{-k_3\}. \quad (69)$$

Similarly for the second term in (68), if $\Delta_j \leq \left(\frac{b^2|\mathcal{C}|k_4}{2m^2d} \right)^{\frac{1}{2}}$, we have

$$\begin{aligned} (1 - 2s)^{-|\mathcal{C}|d} &\geq \exp \{4s|\mathcal{C}|d\} \\ &= \exp \left\{ 4|\mathcal{C}|d \left(-\frac{1}{2} \left(\frac{m}{b|\mathcal{C}|} \right)^2 \left(\frac{T - t_j}{T - t_{j-1}} \right) \Delta_j^2 \right) \right\} \\ &= \exp \left\{ -2 \left(\frac{m^2}{b^2|\mathcal{C}|} \right) d \Delta_j^2 \right\} \geq \exp \{-k_4\}. \end{aligned} \quad (70)$$

Combining the bounds in (69) and (70), and taking the limit $\Delta_j \rightarrow 0$ over sequences of $t_j - t_{j-1} \rightarrow 0$, with (36), we arrive at the result given in the theorem. \blacksquare

Proof (Proposition 19 Using Theorem 15, then for iteration j , we want to choose $\exp\{-k_{3,j} - k_{4,j}\} = \zeta' \in (0, 1)$, and so $k_{3,j} = -\log(\zeta') - k_{4,j}$. By substituting this into (36), we can choose the mesh size as

$$\tilde{\Delta}_j = \min \left\{ \frac{b^2|\mathcal{C}|[-\log(\zeta') - k_{4,j}]}{\mathbb{E}[\nu_j]m^2}, \left(\frac{b^2|\mathcal{C}|k_{4,j}}{2m^2d} \right)^{\frac{1}{2}} \right\}, \quad (71)$$

where $k_{4,j} < -\log(\zeta')$ (in order to ensure that $k_{3,j} > 0$). Here, we want the largest interval which satisfies $N^{-1}\text{CESS}_j \geq \zeta'$. This corresponds to choosing $k_{4,j}$ with

$$\begin{aligned} & \frac{b^2|\mathcal{C}|[-\log(\zeta') - k_{4,j}]}{\mathbb{E}[\nu_j]m^2} = \left(\frac{b^2|\mathcal{C}|k_{4,j}}{2m^2d} \right)^{\frac{1}{2}} \\ \implies & \frac{b^4|\mathcal{C}|^2[-\log(\zeta') - k_{4,j}]^2}{\mathbb{E}[\nu_j]^2m^4} = \frac{b^2|\mathcal{C}|k_{4,j}}{2m^2d} \\ \implies & [-\log(\zeta') - k_{4,j}]^2 = \frac{\mathbb{E}[\nu_j]^2m^2}{2b^2|\mathcal{C}|d}k_{4,j} \\ \implies & \log(\zeta')^2 + 2k_{4,j}\log(\zeta') + k_{4,j}^2 = \frac{\mathbb{E}[\nu_j]^2m^2}{2b^2|\mathcal{C}|d}k_{4,j} \\ \implies & k_{4,j}^2 + \left(2\log(\zeta') - \frac{\mathbb{E}[\nu_j]^2m^2}{2b^2|\mathcal{C}|d} \right) k_{4,j} + \log(\zeta')^2 = 0. \end{aligned} \quad (72)$$

Applying the quadratic formula to solve (72) gives

$$k_{4,j} = \frac{\left(\frac{\mathbb{E}[\nu_j]^2m^2}{2b^2|\mathcal{C}|d} - 2\log(\zeta') \right) \pm \sqrt{\left(2\log(\zeta') - \frac{\mathbb{E}[\nu_j]^2m^2}{2b^2|\mathcal{C}|d} \right)^2 - 4\log(\zeta')^2}}{2}.$$

Note that we have the constraints that $0 < k_{4,j} < -\log(\zeta')$, and since from (72), we have

$$k_{4,j}^2 + \left(2\log(\zeta') - \frac{\mathbb{E}[\nu_j]^2m^2}{2b^2|\mathcal{C}|d} \right) k_{4,j} = -\log(\zeta')^2,$$

then we will always choose the smaller root and arrive at the statement of the theorem. \blacksquare

Appendix H. Practical implementation considerations

In many practical settings there will be additional constraints which require us to modify Algorithm 1 appropriately. Examples include settings where latency between cores is problematic, or in scenarios where functional evaluations of the sub-posterior densities f_c are not available. In this section, consider several modifications to Algorithm 1 to make it more amenable to certain application areas. To clarify, the implementation of our methodology in examples presented in Section 5 do not exploit these modifications that we present below.

H.1 Reducing communication between the cores

For our GBF approach, we highlight two steps where communication between cores could be reduced. In particular, it is possible to limit the amount of communication necessary when initialising the particle set, and also when we propagate the particles in the iterative steps of the algorithm. In a distributed/parallel setting, it is desirable to reduce the number of communication between cores since there is a latency penalty for each communication leading to a more computationally expensive algorithm.

In Algorithm 1 Step 1b, the particles are composed by pairing the sub-posterior draws index-wise to obtain $\{\tilde{\mathbf{x}}_{0,i}^{(c)}\}_{i=1}^M$ which requires a communication between the cores. To fully initialise the algorithm, we must assign importance weights to the particles which requires an additional two communications between the cores; namely a communication back to the individual cores to provide the weighted mean of the particles $\tilde{\mathbf{x}}_{0,i}$, and a communication between the cores to compute $\rho_{0,i}(\tilde{\mathbf{x}}_{0,i}^{(c)})$ (since (4) can be decomposed into a product of $|\mathcal{C}|$ terms corresponding to the individual contributions from each sub-posterior. Following the approach of Dai et al. (2023, Section 3.7.1), let $\tilde{\boldsymbol{\theta}} \in \mathbb{R}^d$ be a weighted average of approximate modes (or means) of each sub-posterior. Noting that this can be computed in a single pre-processing step prior to initialisation, then we can modify the proposal mechanism for the initial draw to be from the density

$$\tilde{f}_c(\mathbf{x}_0^{(c)}) \propto \exp\left\{-\frac{(\mathbf{x}_0^{(c)} - \tilde{\boldsymbol{\theta}})^\top \boldsymbol{\Lambda}_c^{-1}(\mathbf{x}_0^{(c)} - \tilde{\boldsymbol{\theta}})}{2T}\right\} \cdot f_c(\mathbf{x}_0^{(c)}), \quad (73)$$

then by modifying the algorithm by replacing ρ_0 with

$$\tilde{\varrho}_0 := \exp\left\{\frac{(\tilde{\mathbf{x}}_0^{(c)} - \tilde{\boldsymbol{\theta}})^\top \boldsymbol{\Lambda}_c^{-1}(\tilde{\mathbf{x}}_0^{(c)} - \tilde{\boldsymbol{\theta}})}{2T}\right\}, \quad (74)$$

where $\boldsymbol{\Lambda}_c^{-1} := (\sum_{c \in \mathcal{C}} \boldsymbol{\Lambda}_c^{-1})$, we can see that

$$\tilde{\varrho}_0(\tilde{\mathbf{x}}_0^{(c)}) \cdot \prod_{c \in \mathcal{C}} \tilde{f}_c(\mathbf{x}_0^{(c)}) \propto \rho_0(\tilde{\mathbf{x}}_0^{(c)}) \cdot \prod_{c \in \mathcal{C}} f_c(\mathbf{x}_0^{(c)}).$$

Since we subsequently re-normalise the importance weights, we do not need to compute any constant of proportionality for $\tilde{\varrho}_0$. Adopting this approach means that we can sample from \tilde{f}_c on each core independently and evaluate the modified importance weight without any further communication between the cores. This therefore reduces the number of communications required to initialise the particle set from three (in the original formulation) to two (since this approach does require one communication in order to compute $\tilde{\boldsymbol{\theta}}$). The modified initialisation is summarised in Algorithm 6.

There is also scope to reduce the number of communications required to propagate the particle set in Algorithm 1 Step 2(b)i. To propagate the particles, there is a communication between the cores in order to compute $\vec{\mathbf{M}}_j^{(c)} := \vec{\mathbf{M}}_{t_{j-1}, t_j}^{(c)}$ as per (9) since this requires the

Algorithm 6 Particle set initialisation modification (to replace Algorithm 1 Step 1b).

1(b) For k in 1 to M ,

- (i) $\vec{\mathbf{x}}_{0,k}^{(c)}$: For $c \in \mathcal{C}$, simulate $\mathbf{x}_{0,k}^{(c)} \sim \tilde{f}_c$ (73). Set $\vec{\mathbf{x}}_{0,k}^{(c)} := (\mathbf{x}_{0,k}^{(c_1)}, \dots, \mathbf{x}_{0,k}^{(c_{|c|})})$.
 - (ii) Compute un-normalised weight $w_{0,k}^{(c)'} := (\prod_{c \in \mathcal{C}} w_k^{(c)}) \cdot \tilde{\varrho}_0(\vec{\mathbf{x}}_{0,k}^{(c)})$ as per (74).
-

current position of each of the $|\mathcal{C}|$ trajectories. Once we have computed this and propagated the samples, a further communication back to the cores would be necessary so that each core can compute their contribution to the $\tilde{\rho}_j$ importance weight. Alternatively, we can utilise Corollary 23 so that each of the $|\mathcal{C}|$ processes can propagate their own individual particles to compose $\vec{\mathbf{x}}_j^{(c)}$.

Corollary 23 *Simulating $\vec{\mathbf{x}}_j^{(c)} \sim \mathcal{N}_d(\vec{\mathbf{M}}_j^{(c)}, \mathbf{V}_j)$, the required transition from $\vec{\mathbf{x}}_{j-1}^{(c)}$ to $\vec{\mathbf{x}}_j^{(c)}$ in Algorithm 1 Step 2(b)i, can be expressed as*

$$\mathbf{x}_j^{(c)} = \left[\frac{\Delta_j^2}{T - t_{j-1}} \right]^{\frac{1}{2}} \boldsymbol{\xi}_j + \left[\frac{T - t_j}{T - t_{j-1}} \Delta_j \right]^{\frac{1}{2}} \boldsymbol{\eta}_j^{(c)} + \mathbf{M}_j^{(c)}, \quad (75)$$

where $\boldsymbol{\xi}_j \sim \mathcal{N}_d(\mathbf{0}, \boldsymbol{\Lambda}_c)$, $\boldsymbol{\eta}_j^{(c)} \sim \mathcal{N}_d(\mathbf{0}, \boldsymbol{\Lambda}_c)$ and $\mathbf{M}_j^{(c)}$ is the sub-vector of $\vec{\mathbf{M}}_j^{(c)}$ corresponding to the c th component given by (9).

Proof From Proposition a, we have $\vec{\mathbf{x}}_j^{(c)} \sim \mathcal{N}_d(\vec{\mathbf{M}}_j^{(c)}, \mathbf{V}_j)$ where $\vec{\mathbf{M}}_j^{(c)} := \vec{\mathbf{M}}_{t_{j-1}, t_j}^{(c)}$ is given by (9) and $\mathbf{V}_j := \mathbf{V}_{t_{j-1}, t_j}$ is given by (10). From (75), the mean and covariance matrix of $\vec{\mathbf{x}}_j^{(c)}$ given $\vec{\mathbf{x}}_{j-1}^{(c)}$ are also given by $\vec{\mathbf{M}}_j^{(c)}$ and \mathbf{V}_j as required. ■

By using Corollary 23, we can see that the interaction between the $|\mathcal{C}|$ trajectories occurs through their weighted mean $\tilde{\mathbf{x}}_{j-1}$ at the previous iteration. This can be computed at the previous iteration, and we can communicate this along with the common Gaussian vector $\boldsymbol{\xi}_j$ at the same time. This therefore removes an unnecessary additional communication between the cores at every iteration, resulting in a much more efficient algorithm if latency is a concern. This approach is presented in Algorithm 7.

Algorithm 7 Particle set propagation modification (to replace Algorithm 1 Step 2(b)i).

2(b)i.

- (A) For $c \in \mathcal{C}$, simulate $\mathbf{x}_{j,i}^{(c)} | (\tilde{\mathbf{x}}_{j-1,i}, \mathbf{x}_{j-1,i}^{(c)})$ in (75).
 - (B) Set $\vec{\mathbf{x}}_{j,i}^{(c)} := (\mathbf{x}_{j,i}^{(c_1)}, \dots, \mathbf{x}_{j,i}^{(c_{|c|})})$ and compute $\tilde{\mathbf{x}}_{j,i} := (\sum_{c \in \mathcal{C}} \boldsymbol{\Lambda}_c^{-1})^{-1} (\sum_{c \in \mathcal{C}} \boldsymbol{\Lambda}_c^{-1} \mathbf{x}_{j,i}^{(c)})$.
-

H.2 Alternative methods for updating the particle set weights

In this paper, we have assumed that we have been able to compute functionals of each sub-posterior f_c for $c \in \mathcal{C}$, however there are many settings where it may be impractical or infeasible to do so. This may be case if there is some form of intractability of the sub-posteriors (see for instance Andrieu and Roberts (2009)), or maybe the evaluation of such quantities may be simply too computationally expensive (for instance in large data settings (Pollock et al., 2020; Bouchard-Côté et al., 2018; Bierkens et al., 2019; Dai et al., 2023)). In these settings, we no longer are able to evaluate ϕ_c in (6) which is necessary to update the particle weights in the iterative steps of the BF algorithm. However, it is possible to consider alternative unbiased estimators for $\tilde{\rho}_j$ in Step 2c.

Corollary 24 (Dai et al., 2023, Corollary 3) *The estimator*

$$\tilde{\varrho}_j(\vec{\mathbf{x}}_{j-1}^{(c)}, \vec{\mathbf{x}}_j^{(c)}) := \prod_{c \in \mathcal{C}} \left(\frac{\Delta_j^{\kappa_c} \cdot e^{-\bar{U}_X^{(c)} \Delta_j}}{\kappa_c! \cdot p(\kappa_c | R_c)} \prod_{k_c=1}^{\kappa_c} \left(\bar{U}_X^{(c)} - \tilde{\phi}_c(\mathbf{X}_{\xi_c, k_c}^{(c)}) \right) \right), \quad (76)$$

where $\tilde{\phi}_c$ is an unbiased estimator of ϕ_c and $\bar{U}_j^{(c)}$ is a constant such that $\tilde{\phi}_c(\mathbf{x}) \leq \bar{U}_j^{(c)}$ for $\mathbf{x} \in R_c$.

Proof This follows directly from Theorem 4. ■

The estimator $\tilde{\varrho}_j$ in Corollary 24 can therefore be used as a substitute for $\tilde{\rho}_j$ in Algorithm 1 Step 2c. However, we must be careful in constructing $\tilde{\varrho}_j$ since its introduction typically increases the variance of the estimator which ultimately causes higher variance in the particle set weights in the BF algorithm. In particular, by using Corollary 24, the number of expected functional evaluations will change from K to K' and so we must consider the growth in the ratio K'/K as $m_c \rightarrow \infty$ (Pollock et al., 2020; Dai et al., 2023). However, as noted, introducing an alternative unbiased estimator may be necessary to apply the BF approach to some settings.

For instance, consider the example setting provided in Dai et al. (2023, Appendix E), where we have a large number of data points associated to each sub-posterior (i.e we have $m_c \gg 1$ data points for core $c \in \mathcal{C}$) then computing ϕ_c in (6) is an expensive $\mathcal{O}(m_c)$ operation. However, since ϕ_c is linear in terms of $\nabla \log f_c(\mathbf{x})$ and $\nabla^2 \log f_c(\mathbf{x})$, it is simple to construct an unbiased estimator $\tilde{\phi}_c$ for ϕ_c . In the setting, we also assume the sub-posteriors admit a structure with conditional independence and can be factorised as follows,

$$f_c(\mathbf{x}) \propto \prod_{i=1}^{m_c} l_{i,c}(\mathbf{x}). \quad (77)$$

Then since ϕ_c is linear in terms of $\nabla \log l_{i,c}(\mathbf{x})$ and $\nabla^2 \log l_{i,c}(\mathbf{x})$, then we could use the following naive unbiased estimator for ϕ_c^{dl} :

$$\tilde{\phi}_c(\mathbf{x}) = \frac{m_c}{2} \left(\nabla \log l_{I,c}(\mathbf{x}^*)^\top \mathbf{\Lambda}_c \nabla \log l_{J,c}(\mathbf{x}^*) + \text{Tr}(\mathbf{\Lambda}_c \nabla^2 \log l_{I,c}(\mathbf{x}^*)) \right), \quad (78)$$

where $I, J \stackrel{\text{iid}}{\sim} \mathcal{U}\{1, \dots, m_c\}$. Although using such an estimator has the advantage of having $\mathcal{O}(1)$ cost when evaluating, this comes at the cost of an $\mathcal{O}(m_c)$ inflation in the expected number of evaluations when evaluating $\tilde{\rho}_j$ over $\tilde{\rho}_j$. However, following the approach of Pollock et al. (2020, Section 4) and Dai et al. (2023, Appendix E), we first want to suitably choose some *control variates* to construct our estimator, and compute $\nabla \log f_c$ and $\nabla^2 \log f_c$ at points *close* to either the mode of the sub-posterior, $\hat{\mathbf{x}}_c$, or the mode of the target posterior $\hat{\mathbf{x}}$ (where close means within $\mathcal{O}(m_c^{-\frac{1}{2}})$ of the true respective modes). Computing these control variates will typically be one-time $\mathcal{O}(m_c)$ computations.

Let

$$\tilde{\alpha}_{I,c}(\mathbf{x}) := n \cdot [\nabla \log l_{I,c}(\mathbf{x}) - \nabla \log l_{I,c}(\mathbf{x}^*)], \quad (79)$$

$$\tilde{H}_{I,c}(\mathbf{x}) := n \cdot [\nabla^2 \log l_{I,c}(\mathbf{x}) - \nabla^2 \log l_{I,c}(\mathbf{x}^*)], \quad (80)$$

then since $\log f_c(\mathbf{x}) = \sum_{i=1}^{m_c} \log l_{i,c}(\mathbf{x})$, we have

$$\mathbb{E}_{\mathcal{A}} [\tilde{\alpha}_{I,c}(\mathbf{x})] = \alpha_c(\mathbf{x}), \quad \mathbb{E}_{\mathcal{A}} [\tilde{H}_{I,c}(\mathbf{x})] = H_c(\mathbf{x}). \quad (81)$$

where $\alpha_c(\mathbf{x}) := \nabla \log f_c(\mathbf{x}) - \nabla \log f_c(\mathbf{x}^*)$ and $H_c(\mathbf{x}) := \nabla^2 \log f_c(\mathbf{x}) - \nabla^2 \log f_c(\mathbf{x}^*)$ and \mathcal{A} is the law of $I \sim \mathcal{U}\{1, \dots, n\}$.

Noting that $\phi_c(\mathbf{x})$ in (6) can be re-expressed as

$$\phi_c(\mathbf{x}) = \frac{1}{2} [\alpha_c(\mathbf{x})^\top \mathbf{\Lambda}_c (2\nabla \log f_c(\mathbf{x}^*) + \alpha_c(\mathbf{x})) + \text{Tr}(\mathbf{\Lambda}_c H_c(\mathbf{x}))] + C^*, \quad (82)$$

where $C^* := \frac{1}{2} (\nabla \log f_c(\mathbf{x}^*)^\top \mathbf{\Lambda}_c \nabla \log f_c(\mathbf{x}^*) + \text{Tr}(\mathbf{\Lambda}_c \nabla^2 \log f_c(\mathbf{x}^*)))$, then this leads to the following unbiased estimator for ϕ_c :

$$\tilde{\phi}_c(\mathbf{x}) := \frac{1}{2} [\alpha_{I,c}(\mathbf{x})^\top (2\nabla \log f_c(\mathbf{x}^*) + \alpha_{J,c}(\mathbf{x})) + \text{Tr}(\mathbf{\Lambda}_c \tilde{H}_{I,c}(\mathbf{x}))] + C^*, \quad (83)$$

where $I, J \stackrel{\text{iid}}{\sim} \mathcal{U}\{1, \dots, m_c\}$, i.e. if now we let \mathcal{A} be the law of $I, J \stackrel{\text{iid}}{\sim} \mathcal{U}\{1, \dots, m_c\}$, we have $\mathbb{E}_{\mathcal{A}} [\tilde{\phi}_c(\mathbf{x})] = \phi_c(\mathbf{x})$.

Here the evaluations of the constants $\|\nabla \log f_c(\mathbf{x}^*)\|^2$, $\text{Tr}(\nabla^2 \log f_c(\mathbf{x}^*))$ are of $\mathcal{O}(m_c)$ cost, but they only need to be computed once prior to calling Algorithm 1. The unbiased estimator $\tilde{\phi}_c(\mathbf{x})$ uses only double draws from $\{1, \dots, m_c\}$, although Pollock et al. (2020) notes that it would be possible to replace this by averaging over multiple draws (sampling from $\{1, \dots, m_c\}$ with replacement) which could have advantages of reducing the variance of the estimator at the cost of increasing the number of data points to evaluate at.

Appendix I. Simulation studies

In this section we study empirically the performance of our Fusion algorithms (Sections 2 and 3), and selection of tuning parameters (T , n and \mathcal{P} as discussed in Section 4) in our

two idealised key settings—the SH(λ) setting (Condition 8) and SSH(γ) setting (Condition 10) described in Section 4. We do this in Sections I.1 and I.2 respectively. For simplicity, here we focus on BF and GBF, noting that GBF is simply D&C-Fusion with a *fork-and-join tree* hierarchy (as in Figure 1). Finally, in Section I.3 we compare the performance of Fusion methodologies (including D&C-Fusion with a *balanced-binary* tree hierarchy) with increasing dimensionality. In Section 5, we consider more substantive examples using real data. Note that the earlier Bayesian Fusion approach is simply a special case of our GBF approach with $\mathbf{\Lambda}_c = \mathbb{I}_d$ for $c \in \mathcal{C}$, and so comparison with this work is straight-forward.

To compare the performance of different approaches we consider their computational cost (both the total run time, and n which represents the number of iterations of Algorithm 1 Step 2 and so is a proxy for the amount of communication between cores), and *Integrated Absolute Distance (IAD)* defined in (45).

Throughout this section we use the GPE-2 estimator of ρ_j as given in Definition 6, and use the Trapezoidal rule to estimate the mean γ_c in (24) and set $\beta_c = 10$ for $c \in \mathcal{C}$. Code to run these simulation studies can be found at <https://github.com/rchan26/DCFusion>.

I.1 Sub-posterior Homogeneity

We first study the guidance developed for T and \mathcal{P} in Section 4 for GBF (Algorithm 1) in the SH(λ) setting of Condition 8. Recall, this is the setting in which we are combining homogeneous sub-posteriors, and would naturally arise if a dataset was split randomly across several cores. To study this setting, we consider the idealised scenario of combining $C = 10$ bi-variate Gaussian sub-posteriors, with a range of data sizes from $m = 1000$ to $m = 40000$, which have been randomly split across the $C = 10$ cores. In particular, each sub-posterior has mean $\mathbf{0} = (0, 0)$ and variance $\frac{C}{m}\mathbf{\Sigma}$, where $\mathbf{\Sigma} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$ with $\rho = 0.9$. For this example, we apply both BF and GBF with a fixed particle set size of $N = 10000$.

To verify the guidance for T and \mathcal{P} , we consider varying T and \mathcal{P} with increasing data size m , and the impact this has on CESS_0 and CESS_j (for $j \in \{1, \dots, n\}$). We consider the four following choices of T and \mathcal{P} :

1. a fixed choice of T and n to obtain \mathcal{P} (for GBF, $T = 1$ and $n = 5$, and for BF, $T = 0.005$ and $n = 5$),
2. using the recommended T from Section 4.1 and fixed $n = 5$ to obtain \mathcal{P} ,
3. using the recommended T and \mathcal{P} using a regular mesh (as outlined in Algorithm 3 and Section 4.2.1),
4. using the recommended T and \mathcal{P} using an adaptive mesh (as outlined in Algorithm 4 and Section 4.2.2).

In implementing the BF and GBF, we set our lower tolerable bounds for the initial (CESS_0) and the iterative (CESS_j) conditional effective sample sizes to be $0.5N$ (i.e. we set $\zeta = \zeta' = 0.5$), resampling if ESS falls below $0.5N$. To summarise how one might practically use our guidance to choose T and \mathcal{P} , we present our approach in Remark 25.

Remark 25 *We set the tuning parameters for BF and GBF (for the $\text{SH}(\lambda)$ setting of Section I.1) as follows:*

1. *Following the guidance outlined in Remark 14, and with $\zeta = 0.5, \lambda = 1$ and $d = 2$, we have $k_1 = \sqrt{-\frac{(\lambda + \frac{d}{2})}{\log(\zeta)}} \approx 1.7$. For GBF, $\mathbf{\Lambda}_c$ is the estimated covariance matrices for sub-posterior $c \in \{1, \dots, C\}$, so $b = \frac{m}{C}$ (see Remark 11), and we choose $T = C^{\frac{1}{2}}k_1$. For BF, $\mathbf{\Lambda}_c = \mathbb{I}_d$ for $c \in \{1, \dots, C\}$, so we have $b = 1$ and so we choose $T = C^{3/2}k_1/m$.*
2. *When using the regular mesh, we use Algorithm 3 to obtain \mathcal{P} . First let $\zeta' = 0.5$ then for GBF we have $b = \frac{m}{C}$, and so $\Delta_j = \Delta = \sqrt{\frac{k_4}{2Cd}}$ for each j , where k_4 is computed as per (44) and computing an estimate of the supremum of $\widehat{\mathbb{E}[\nu_j]}$ as per (41). For BF, $b = 1$, so $\Delta_j = \Delta = \sqrt{\frac{Ck_4}{2m^2d}}$ for each j .*
3. *When using the adaptive mesh, we use Algorithm 4 to obtain Δ_j recursively at each iteration to construct \mathcal{P} . We let $\zeta' = 0.5$ and for GBF (where $b = \frac{m}{C}$) we compute $t_j = \min\{T, t_{j-1} + \Delta_j\}$ where $\Delta_j = \sqrt{\frac{k_4}{2Cd}}$ at each iteration of Algorithm 1, until we have $t_j = T$. For the standard BF approach, note that $b = 1$ so we must compute $\Delta_j = \sqrt{\frac{Ck_4}{2m^2d}}$ instead at each iteration.*

The conditional effective sample size of the GBF and (standard) BF approaches with increasing data size in this $\text{SH}(\lambda)$ setting are shown in Figure 12.

First considering the results from fixing T and n in Figure 12a, we can see that BF lacks robustness with increasing data size. Here CESS_0 improves with increasing data size (m), which is due to the sub-posteriors becoming increasingly similar with m in this idealised scenario. However, as we increase m the fixed choice for T (and hence the size of the intervals) becomes increasingly inappropriate for the sub-posteriors, which leads to a degradation in average CESS_j . In contrast, GBF incorporates global information about the sub-posteriors (i.e. the variance of the sub-posteriors), so there is no change in performance with m . Here there is a trade-off with the choice of T : a small T leads to poor behaviour on initialisation (i.e. low CESS_0), but good behaviour at each iteration (i.e. high average CESS_j).

Considering Figure 12b, we see that scaling T following the guidance developed in Section 4.1 immediately stabilises CESS_0 , although CESS_j performance is still poor (n is too small). In Figure 12c and Figure 12d, we see that utilising both the guidance for T and the mesh \mathcal{P} drastically improves the performance of both BF and GBF. In both cases GBF outperforms

BF: it achieves higher average CESS_j , and the variance of CESS_j is lower. Given BF is a special case of GBF, this improvement can be ascribed to the use of estimated covariance matrices for Λ_c . In particular, this choice leads to a lower variance unbiased estimator for ρ_j , and an improved proposal h^{bf} (14) for g^{bf} (15).

From Figure 12e we see that with BF that without our guidance on T and \mathcal{P} , average IAD is poor, and the variance of the IAD is very large. In contrast, GBF with our guidance is robust across the different scenarios. Comparing the regular and adaptive meshes simply using CESS_0 and CESS_j would imply that the regular mesh is performing better (since it has slightly better CESS_j), however the adaptive mesh is slightly more computationally efficient as shown by having a smaller mesh size, n , (illustrated in Figure 12f) and having a faster algorithm run-time (illustrated in Figure 12g). By looking at the IAD obtained for these approaches, we can see that we are able to obtain similar performance at a lower cost with the adaptive mesh construction.


I.2 Sub-posterior Heterogeneity

Now we study the guidance for T and \mathcal{P} for GBF (Algorithm 1) developed in Section 4 in the $\text{SSH}(\gamma)$ setting of Condition 10. This represents the setting where sub-posterior heterogeneity does not decay with data size. Here, we consider the scenario of combining $C = 2$ bi-variate Gaussian sub-posteriors, $f_c \sim \mathcal{N}(\mu_c, \frac{2}{m}\Sigma)$, where $\mu_1 = -(0.25, 0.25)$ and $\mu_2 = (0.25, 0.25)$ and $\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$, with $\rho = 0.9$. We again consider a range of data sizes, which ranges from $m = 250$ to $m = 2500$ and are randomly split between $C = 2$ cores. We apply BF and GBF with a fixed particle set size of $N = 10000$.


In this setting as m increases the sub-posterior heterogeneity increases, which is a consequence of the sub-posteriors having diminishing overlapping support. In BF (where $\Lambda_1 = \Lambda_2 = \mathbb{I}_d$) this heterogeneity is not captured, and $\sigma_a^2 = 0.125$ irrespective of m . By contrast, the generalised approach is able to capture the heterogeneity with m with simply the inclusion of the estimated covariance matrices $\{\Lambda_c\}_{c=1,2}$.

As with the previous example in Section I.1, we will investigate the effect of varying T and \mathcal{P} with m , and its impact upon CESS_0 and CESS_j . We consider the following choices for T and \mathcal{P} :

1. a fixed choice of T and n to obtain \mathcal{P} (for GBF, $T = 2$ and $n = 5$, and for BF, $T = 0.01$ and $n = 5$),
2. using the recommended T from Section 4.1 and fixed $n = 5$ to obtain \mathcal{P} ,
3. using the recommended T and \mathcal{P} using a regular mesh (as outlined in Algorithm 3 and Section 4.2.1),
4. using the recommended T and \mathcal{P} using a, adaptive mesh (as outlined in Algorithm 4 and Section 4.2.2).




(a) Fixed user-specified T and n .




(b) Recommended T and fixed n .

Figure 12: Bivariate Gaussian example in $SH(\lambda)$ setting with increasing data size. In Figures 12a, 12b, 12c, 12d solid lines denote initial CESS ($CESS_0$), and dotted lines denote averaged CESS in subsequent iterations ($\frac{1}{n} \sum_{j=1}^n CESS_j$), and crosses denote $CESS_j$ for each $j \in \{1, \dots, n\}$.




(i) BF




(ii) GBF

(c) Recommended T and recommended regular mesh \mathcal{P} .




(i) BF




(ii) GBF

(d) Recommended T and recommended adaptive mesh \mathcal{P} .




(i) BF




(ii) GBF

(e) Integrated absolute distance: lines connect the mean IAD (averaged over ten runs) while the points denote the individual IAD achieved on each run.

Figure 12: Bivariate Gaussian example in $SH(\lambda)$ setting (continued).




(i) BF




(ii) GBF

(f) Comparison of mesh sizes between regular and adaptive schemes.



(i) BF



(ii) GBF

(g) Mean average computational run-times (based on ten runs).

Figure 12: Bivariate Gaussian example in $SH(\lambda)$ setting (continued).

When applying the guidance, we set the lower tolerable bounds on the initial (CESS_0) and iterative (CESS_j) conditional effective sample sizes to be $0.5N$ (i.e. we set $\zeta = \zeta' = 0.5$), and re-sample if ESS drops below $0.5N$. Again, for helping with the practical interpretation of our extensive guidance for selecting T and \mathcal{P} , we summarise our approach in Remark 26:

Remark 26 *We set the tuning parameters for BF and GBF (for the $\text{SSH}(\gamma)$ setting of Section I.2) as follows:*

1. *We follow the guidance outlined in Remark 14, noting that $\zeta = 0.5$ and $d = 2$. For GBF, $\mathbf{\Lambda}_{c=1,2}$ are the estimated covariance matrices for each of the sub-posteriors, so $b = \frac{m}{C}$ (see Remark 11), and $\gamma = m\sigma_{\mathbf{a}}^2/C$ (where $\sigma_{\mathbf{a}}^2$ is estimated from the sub-posterior samples). Consequently, we can compute $k_1 = k_2 = \sqrt{-\frac{(\frac{\gamma m}{C} + \frac{d}{2})}{\log(\zeta)}}$, and choose $T = C^{\frac{1}{2}}k_1$. For BF, $\mathbf{\Lambda}_{c=1,2} = \mathbb{I}_d$, so $b = 1$ and $\gamma = \sigma_{\mathbf{a}}^2$, and so we can compute $k_1 = \sqrt{-\frac{(\frac{\gamma m}{C} + \frac{d}{2})}{\log(\zeta)}}$ and $k_2 = \frac{Ck_1}{m}$, and choose $T = \frac{C^{3/2}k_1}{m} = C^{\frac{1}{2}}k_2$.*
2. *When using the regular mesh, we use Algorithm 3 to obtain \mathcal{P} . As $\zeta' = 0.5$, we have for GBF $b = \frac{m}{C}$, and so $\Delta_j = \Delta = \sqrt{\frac{k_4}{2Cd}}$ for each j where k_4 is computed as per (44) (with $\sup_j \widehat{\mathbb{E}[\nu_j]}$ computed as per (41)). For BF we have $b = 1$, so $\Delta_j = \Delta = \sqrt{\frac{Ck_4}{2m^2d}}$ for each j .*
3. *When using the adaptive mesh, we use Algorithm 4 to obtain Δ_j recursively at each iteration to construct \mathcal{P} . With $\zeta' = 0.5$ for the GBF (where $b = \frac{m}{C}$) we compute $t_j = \min\{T, t_{j-1} + \Delta_j\}$ where $\Delta_j = \sqrt{\frac{k_4}{2Cd}}$ at each iteration of Algorithm 1 until we have $t_j = T$. For BF with $b = 1$ we have instead $\Delta_j = \sqrt{\frac{Ck_4}{2m^2d}}$ at each iteration.*

CESS for BF and GBF with increasing m in this $\text{SSH}(\gamma)$ setting are shown in Figure 13. We can immediately see that the $\text{SSH}(\gamma)$ setting is much more challenging than the idealised $\text{SH}(\lambda)$ setting of Section I.1 and Figure 12, which is unsurprising as in this case the sub-posteriors are becoming increasingly mismatched as data size increases.

In Figure 13a, we see that fixing T and n is not ideal for either BF or GBF. As shown in Figure 13b, there is a positive effect for both BF and GBF in using our recommended scaling of T in the quality of the initialisation. In Figure 13c and Figure 13d, where both the guidance for T and \mathcal{P} are implemented, we see a substantial improvement in the performance of both approaches with respect to CESS , with again our new GBF approach outperforming BF.

In Figure 13c we see that the use of a regular mesh in choosing \mathcal{P} , following our guidance, provides robust CESS_j with low variance. Indeed, it appears to outperform the adaptive mesh approach for \mathcal{P} (see Figure 13d). However, as discussed in Section 4.2.2, the regular mesh is overly conservative, and when we factor in the reduced number of iterations required

in the adaptive case (Figure 13f), along with the overall reduction in computational cost (Figure 13g) for comparable IAD (Figure 13e), we see that the use of an adaptive mesh is preferable.






Figure 13: Bivariate Gaussian example in $SSH(\gamma)$ setting with increasing data size. In Figures 13a, 13b, 13c, 13d solid lines denote initial $CESS$ ($CESS_0$), and dotted lines denote averaged $CESS$ in subsequent iterations ($\frac{1}{n} \sum_{j=1}^n CESS_j$), and crosses denote $CESS_j$ for each $j \in \{1, \dots, n\}$.




(i) BF




(ii) GBF

(c) Recommended T and recommended regular mesh \mathcal{P} .




(i) BF




(ii) GBF

(d) Recommended T and recommended adaptive mesh \mathcal{P} .




(i) BF




(ii) GBF

(e) Integrated absolute distance: lines connect the mean IAD (averaged over ten runs) while the points denote the individual IAD achieved on each run.

Figure 13: Bivariate Gaussian Example in $\text{SSH}(\gamma)$ setting (continued).




(i) BF




(ii) GBF

(f) Comparison of mesh sizes between regular and adaptive schemes.



(i) BF



(ii) GBF

(g) Mean average computational run-times (based on ten runs).


Figure 13: Bivariate Gaussian Example in SSH(γ) setting (continued).

I.3 Dimension Scaling

In this section we empirically study the performance of Fusion approaches (BF, GBF and D&C-Fusion) with increasing dimensionality. To do so we consider a d -dimensional multi-variate Gaussian $f \propto \prod_{c=1}^C f_c$, where we let $C = 8$ and $f_c \sim \mathcal{N}_d(\mathbf{0}, C\Sigma)$, and where

$$\begin{aligned} \Sigma_{ii} &= 1, & \text{for all } i \in \{1, \dots, d\}, \\ \Sigma_{ij} &= 0.9, & \text{for all } i \neq j, (i, j) \in \{1, \dots, d\}, \end{aligned}$$

and simply vary d (in steps from $d = 1$ to $d = 100$). For BF and GBF we use an adaptive mesh for \mathcal{P} , and for D&C-Fusion we consider both a regular and adaptive mesh for \mathcal{P} with a balanced-binary tree hierarchy. In all cases we use the guidance developed in Section 4. As we are in the SH(λ) setting (the true sub-posterior means are the same), we set $\lambda = 1$. The lower bounds of the tolerable initial and iterative CESS are set to $0.05N$ (i.e. $\zeta = \zeta' = 0.05$) and we resample if the ESS drops below $0.5N$, where here we have $N = 10000$. The results are presented in Figure 14.



(a) Integrated Absolute Distance.

(b) Computational cost.

Figure 14: Comparison of Fusion methodologies with increasing dimensionality (in the setting of Section I.3). In Figure 14a, lines connect the mean IAD (averaged over ten runs) while the points denote the individual IAD achieved on each run.

As shown in Figure 14a, the performance of all Fusion methods degrades with increasing dimensionality: both in terms of the average IAD and also the variance. As our target exhibits high correlation between components, BF struggles here even in low dimensions, whereas the GBF and D&C-Fusion approaches we have developed in this paper offer much better scaling with dimension. D&C-Fusion comfortably outperforms existing Fusion approaches for even moderate dimensionality in terms of IAD and computational cost.

Appendix J. Calculations for examples

In this section, we provide the calculations necessary to implement the Fusion algorithms discussed in this paper. In particular, to implement Generalised Bayesian Fusion (Section 2) and Divide-and-Conquer Fusion (Section 3), we must be able to compute ϕ_c given in (6). This requires the computation of the first and second order derivatives of the log sub-posterior densities.

Furthermore, it is necessary to compute bounds of ϕ_c . As noted in Section 2.2, if it is not possible to (or simply difficult to) find tight bounds for ϕ_c , we can use the general bounds given in Proposition 3. To use these general bounds, we must find an upper bound on the matrix norm of $\mathbf{\Lambda}_c \nabla^2 \log f_c(\mathbf{x})$ for $\mathbf{x} \in R_c$ (i.e. find P^{Λ_c} given in (19)), which can be done by computing the matrix norm of the matrix which bounds the matrix $\mathbf{\Lambda}_c \nabla^2 \log f_c(\mathbf{x})$ element-wise.

We note that in some cases, it may be easier to find the bound on the matrix norm of the Hessian of the *transformed* sub-posterior, $f_c^{(z)}(\mathbf{z})$ where $\mathbf{z} := \mathbf{\Lambda}_c^{-\frac{1}{2}} \mathbf{x}$. In particular, rather than finding a bound in (19), we can focus on finding the bound

$$P^{\Lambda_c} \geq \max_{\mathbf{z} \in R_c^{(z)}} \gamma \left(\nabla^2 \log f_c^{(z)}(\mathbf{z}) \right), \quad (84)$$

which is equivalent to finding the bound in (19).

In Appendix F, we detailed how we can simulate $\tilde{\rho}_j$. In particular, in Algorithm 5, we perform a transformation on the space and in Step 1b, we compute the layer information $R_c^{(z)}$ and so we can directly use this to find local element-wise bounds $\nabla^2 \log f_c^{(z)}(\mathbf{z})$ for $\mathbf{z} \in R_c^{(z)}$. Therefore, to find P^{Λ_c} , we just need to find bounds on the second order derivatives of the log-sub-posterior in the transformed space $\mathbf{z} := \mathbf{\Lambda}_c^{-\frac{1}{2}} \mathbf{x}$ so that we can compute the matrix norm of the matrix which bounds $\nabla^2 \log f_c^{(z)}(\mathbf{z})$ element-wise.

J.1 Logistic Regression

In Section 5.4, we considered applying our Fusion methodologies to a logistic regression example with Gaussian prior distributions for the parameters. In particular, our sub-posterior densities were given by the posterior for Bayesian logistic regression with $\mathcal{N}_d(\mu_j, C\sigma_{\beta_j}^2)$ prior for β_j for $j = 0, \dots, p$ is given by

$$f_c(\boldsymbol{\beta}) := \pi(\boldsymbol{\beta}|\mathbf{y}) = \left[\prod_{i=1}^n \frac{e^{X_i \boldsymbol{\beta} \cdot \mathbf{y}_i}}{1 + e^{X_i \boldsymbol{\beta}}} \right] \cdot \left[\prod_{j=0}^p \frac{1}{\sqrt{2\pi C \sigma_{\beta_j}^2}} \exp \left(-\frac{(\beta_j - \mu_j)^2}{2C \sigma_{\beta_j}^2} \right) \right] \quad (85)$$

where $X \in \mathbb{R}^{n \times (p+1)}$ is the design matrix so $X_i \boldsymbol{\beta} = \beta_0 + \beta_1 X_{i1} + \dots + \beta_p X_{ip}$. The log-posterior is given by

$$\log f_c(\boldsymbol{\beta}) = \sum_{i=1}^n \left[X_i \boldsymbol{\beta} \cdot \mathbf{y}_i - \log(1 + e^{X_i \boldsymbol{\beta}}) \right] - \sum_{j=0}^p \frac{(\beta_j - \mu_j)^2}{2C\sigma_{\beta_j}^2} + \text{constant}. \quad (86)$$

The first derivative of the log-posterior with respect to β_k for $k = 0, \dots, p$, is given by

$$\begin{aligned} \frac{\partial \log f_c(\boldsymbol{\beta})}{\partial \beta_k} &= \sum_{i=1}^n \left[X_{ik} \cdot \mathbf{y}_i - \frac{X_{ik} e^{X_i \boldsymbol{\beta}}}{1 + e^{X_i \boldsymbol{\beta}}} \right] - \frac{(\beta_k - \mu_k)}{C\sigma_{\beta_k}^2} \\ &= \sum_{i=1}^n \left[X_{ik} \cdot \left(\mathbf{y}_i - \frac{1}{1 + e^{-X_i \boldsymbol{\beta}}} \right) \right] - \frac{(\beta_k - \mu_k)}{C\sigma_{\beta_k}^2} \end{aligned} \quad (87)$$

and the second order derivatives of the log-posterior are given by

$$\frac{\partial^2 \log f_c(\boldsymbol{\beta})}{\partial \beta_k^2} = - \sum_{i=1}^n \frac{X_{ik}^2 e^{X_i \boldsymbol{\beta}}}{(1 + e^{X_i \boldsymbol{\beta}})^2} - \frac{1}{C\sigma_{\beta_k}^2}, \quad (88)$$

$$\frac{\partial^2 \log f_c(\boldsymbol{\beta})}{\partial \beta_k \partial \beta_l} = - \sum_{i=1}^n \frac{X_{ik} X_{il} e^{X_i \boldsymbol{\beta}}}{(1 + e^{X_i \boldsymbol{\beta}})^2} \text{ for } k \neq l, \quad (89)$$

for $k, l = 0, \dots, p$. We can use these directly to compute ϕ_c given in (6).

To compute the bounds of ϕ_c , we can utilise the bounds provided in Proposition 3 (or in (17) and (18)). To do so, we must be able to compute an upper bound of the matrix norm $\boldsymbol{\Lambda}_c \nabla^2 \log f_c(\mathbf{x})$ for $\mathbf{x} \in R_c$ where R_c denotes the simulated layer information, i.e. to compute (19). While this can be done by computing the matrix norm of the matrix which bounds the matrix $\boldsymbol{\Lambda}_c \nabla^2 \log f_c(\mathbf{x})$ element-wise, we noted above that it is typically easier to find bounds on the matrix norm of $\nabla^2 \log f_c^{(z)}(\mathbf{z})$ where $\mathbf{z} := \boldsymbol{\Lambda}_c^{-\frac{1}{2}} \boldsymbol{\beta}$, and instead we can focus on finding a bound in the transformed space, i.e. compute (84).

In this logistic regression setting, let $\mathbf{z} = \boldsymbol{\Lambda}_c^{-\frac{1}{2}} \boldsymbol{\beta}$ then the transformed posterior density is given by

$$f_c^{(z)}(\mathbf{z}) = \pi(\boldsymbol{\beta} | X, \mathbf{y}) |J|, \quad (90)$$

where $J = \boldsymbol{\Lambda}_c^{-\frac{1}{2}}$ is the Jacobian matrix with elements $J_{ij} = \frac{\partial z_i}{\partial \beta_j} = \boldsymbol{\Lambda}_{c,ij}^{-\frac{1}{2}}$. We have

$$\log f_c^{(z)}(\mathbf{z}) = \log \pi(\boldsymbol{\beta} | X, \mathbf{y}) + \log |J|. \quad (91)$$

Since $\boldsymbol{\beta} = \boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}$, we have

$$\begin{aligned} f_c^{(z)}(\mathbf{z}) &:= \pi(\boldsymbol{\beta} | X, \mathbf{y}) \cdot |\boldsymbol{\Lambda}_c^{-\frac{1}{2}}| \\ &= \left[\prod_{i=1}^n \frac{e^{X_i \boldsymbol{\beta} \cdot \mathbf{y}_i}}{1 + e^{X_i \boldsymbol{\beta}}} \right] \cdot \left[\prod_{j=0}^p \frac{1}{\sqrt{2\pi C \sigma_{\beta_j}^2}} \exp \left(-\frac{(\beta_j - \mu_j)^2}{2C\sigma_{\beta_j}^2} \right) \right] \cdot |\boldsymbol{\Lambda}_c^{-\frac{1}{2}}| \end{aligned}$$

$$= \left[\prod_{i=1}^n \frac{e^{X_i(\Lambda_c^{\frac{1}{2}} \mathbf{z}) \cdot y_i}}{1 + e^{X_i(\Lambda_c^{\frac{1}{2}} \mathbf{z})}} \right] \cdot \left[\prod_{j=0}^p \frac{1}{\sqrt{2\pi C \sigma_{\beta_j}^2}} \exp \left(-\frac{\left((\Lambda_c^{\frac{1}{2}} \mathbf{z})_j - \mu_j \right)^2}{2C \sigma_{\beta_j}^2} \right) \right] \cdot |\Lambda_c^{-\frac{1}{2}}|, \quad (92)$$

so

$$\log f_c^{(z)}(\mathbf{z}) = \sum_{i=1}^n \left[X_i(\Lambda_c^{\frac{1}{2}} \mathbf{z}) \cdot y_i - \log \left(1 + e^{X_i(\Lambda_c^{\frac{1}{2}} \mathbf{z})} \right) \right] - \sum_{j=0}^p \frac{\left((\Lambda_c^{\frac{1}{2}} \mathbf{z})_j - \mu_j \right)^2}{2C \sigma_{\beta_j}^2} + \text{constant}. \quad (93)$$

We first note that since $\beta = \Lambda_c^{\frac{1}{2}} \mathbf{z}$, then $\beta_i = (\Lambda_c^{\frac{1}{2}} \mathbf{z})_i = \sum_k \Lambda_{ik}^{\frac{1}{2}} z_k$. So we have

$$\begin{aligned} \frac{\partial (X_i \Lambda_c^{\frac{1}{2}} \mathbf{z})}{\partial z_k} &= \frac{\partial}{\partial z_k} \sum_j X_{ij} \beta_j \\ &= \frac{\partial}{\partial z_k} \sum_j X_{ij} \left(\sum_k \Lambda_{jk}^{\frac{1}{2}} z_k \right) \\ &= \sum_j X_{ij} \Lambda_{jk}^{\frac{1}{2}} \\ &= (X \Lambda_c^{\frac{1}{2}})_{ik} \end{aligned} \quad (94)$$

and also we have

$$\frac{\partial (\Lambda_c^{\frac{1}{2}} \mathbf{z})_i}{\partial z_k} = \frac{\partial}{\partial z_k} \sum_j \Lambda_{ij}^{\frac{1}{2}} z_j = \Lambda_{ik}^{\frac{1}{2}} \quad (95)$$

Using (94) and (95), then the first derivative of the log transformed posterior with respect to β_k for $k = 0, \dots, p$, is given by

$$\frac{\partial \log f_c^{(z)}(\mathbf{z})}{\partial z_k} = \sum_{i=1}^n \left[(X \Lambda_c^{\frac{1}{2}})_{ik} \cdot \left(y_i - \frac{1}{1 + e^{-(X_i \Lambda_c^{\frac{1}{2}} \mathbf{z})}} \right) \right] - \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \left((\Lambda_c^{\frac{1}{2}} \mathbf{z})_j - \mu_j \right)}{C \sigma_{\beta_j}^2}. \quad (96)$$

Then the second order derivatives are given by

$$\frac{\partial^2 \log f_c^{(z)}(\mathbf{z})}{\partial z_k \partial z_l} = - \sum_{i=1}^n \frac{(X \Lambda_c^{\frac{1}{2}})_{ik} (X \Lambda_c^{\frac{1}{2}})_{il} e^{(X_i \Lambda_c^{\frac{1}{2}} \mathbf{z})}}{\left(1 + e^{(X_i \Lambda_c^{\frac{1}{2}} \mathbf{z})} \right)^2} - \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \Lambda_{jl}^{\frac{1}{2}}}{C \sigma_{\beta_j}^2}, \quad (97)$$

for $k, l = 0, \dots, p$.

To find bounds for ϕ_c , we must now try to find bounds on the second derivatives given above and compute the matrix norm of the matrix made up of these bounds (which ultimately bounds $\nabla^2 \log f_c^{(z)}(\mathbf{z})$ element-wise). For this example, we can find global and lower bounds of the second derivatives. Note however, we typically will expect better performance with the local bounds on P^{Λ_c} (84) (as this will typically lead to the expected number of points we need to evaluate while performing Poisson thinning, κ_c , to be lower) despite these bounds being slightly more expensive to compute in practice.

J.1.1 GLOBAL BOUNDS OF P^{Λ_c}

We first note that $\frac{e^x}{(1+e^x)^2} \leq \frac{1}{4}$ for all x (and this maximum occurs at $x = 0$). We can utilise this to obtain a global bound:

$$\sup \left[\left| \frac{\partial^2 \log f_c^{(z)}(\mathbf{z})}{\partial z_k \partial z_l} \right| \right] = \sum_{i=1}^n \frac{|X \Lambda_c^{\frac{1}{2}}|_{ik} \cdot |X \Lambda_c^{\frac{1}{2}}|_{il}}{4} + \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \Lambda_{jl}^{\frac{1}{2}}}{C \sigma_{\beta_j}^2}. \quad (98)$$

J.1.2 LOCAL BOUNDS OF P^{Λ_c}

Local bounds can be obtained if we can find local bounds for

$$G_1(\mathbf{z}) := \frac{e^{(X_i \Lambda_c^{\frac{1}{2}}) \mathbf{z}}}{\left(1 + e^{(X_i \Lambda_c^{\frac{1}{2}}) \mathbf{z}}\right)^2}, \quad (99)$$

for $i = 1, \dots, n$. In that case, we have

$$\sup_{\mathbf{z} \in R^{(z)}} \left[\left| \frac{\partial^2 \log f_c^{(z)}(\mathbf{z})}{\partial z_k \partial z_l} \right| \right] = \sum_{i=1}^n \left[|X \Lambda_c^{\frac{1}{2}}|_{ik} \cdot |X \Lambda_c^{\frac{1}{2}}|_{il} \cdot \max_{\mathbf{z} \in R^{(z)}} \{G_1(\mathbf{z})\} \right] + \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \Lambda_{jl}^{\frac{1}{2}}}{C \sigma_{\beta_j}^2}. \quad (100)$$

To compute $\max_{\mathbf{z} \in R^{(z)}} \{G_1(\mathbf{z})\}$, see Section J.3.2 and Algorithm 8 and set $r = 1$.

J.2 Robust Regression

In Section 5.2, we considered a robust regression example (using a student- t distribution) with Gaussian prior distributions for the parameters. In particular, our sub-posterior densities were given by the posterior for Bayesian robust regression with $\mathcal{N}_d(\mu_j, C \sigma_{\beta_j}^2)$ prior for β_j for $j = 0, \dots, p$ is given by

$$f_c(\boldsymbol{\beta}) = \pi(\boldsymbol{\beta} | X, \mathbf{y}) := \left[\prod_{i=1}^n \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2}) \sqrt{\pi \nu} \sigma} \left(1 + \frac{1}{\nu} \left(\frac{y_i - X_i \boldsymbol{\beta}}{\sigma} \right)^2 \right)^{-\left(\frac{\nu+1}{2}\right)} \right]$$

$$\cdot \left[\prod_{j=0}^p \frac{1}{\sqrt{2\pi C\sigma_{\beta_j}^2}} \exp\left(-\frac{(\beta_j - \mu_j)^2}{2C\sigma_{\beta_j}^2}\right) \right]. \quad (101)$$

The log-posterior is given by

$$\log f_c(\boldsymbol{\beta}) = -\left(\frac{\nu+1}{2}\right) \sum_{i=1}^n \log\left(1 + \frac{1}{\nu\sigma^2} (y_i - X_i\boldsymbol{\beta})^2\right) - \sum_{j=0}^p \frac{(\beta_j - \mu_j)^2}{2C\sigma_{\beta_j}^2} + \text{constant}. \quad (102)$$

The first derivative of the log-posterior with respect to β_k for $k = 0, \dots, p$ is given by

$$\begin{aligned} \frac{\partial \log \pi(\boldsymbol{\beta}|X, \mathbf{y})}{\partial \beta_k} &= -\left(\frac{\nu+1}{2}\right) \sum_{i=1}^n \frac{-\frac{2X_{ik}}{\nu\sigma^2} (y_i - X_i\boldsymbol{\beta})}{1 + \frac{1}{\nu\sigma^2} (y_i - X_i\boldsymbol{\beta})^2} - \frac{(\beta_k - \mu_k)}{C\sigma_{\beta_k}^2} \\ &= (\nu+1) \sum_{i=1}^n \frac{X_{ik}(y_i - X_i\boldsymbol{\beta})}{\nu\sigma^2 + (y_i - X_i\boldsymbol{\beta})^2} - \frac{(\beta_k - \mu_k)}{C\sigma_{\beta_k}^2}, \end{aligned} \quad (103)$$

and the second order derivatives of the log-posterior are given by

$$\frac{\partial^2 \log \pi(\boldsymbol{\beta}|X, \mathbf{y})}{\partial \beta_k^2} = (\nu+1) \sum_{i=1}^n \frac{X_{ik}^2 ((y_i - X_i\boldsymbol{\beta})^2 - \nu\sigma^2)}{(\nu\sigma^2 + (y_i - X_i\boldsymbol{\beta})^2)^2} - \frac{1}{C\sigma_{\beta_k}^2}, \quad (104)$$

$$\frac{\partial^2 \log \pi(\boldsymbol{\beta}|X, \mathbf{y})}{\partial \beta_k \partial \beta_l} = (\nu+1) \sum_{i=1}^n \frac{X_{ik} X_{il} ((y_i - X_i\boldsymbol{\beta})^2 - \nu\sigma^2)}{(\nu\sigma^2 + (y_i - X_i\boldsymbol{\beta})^2)^2} \text{ for } k \neq l, \quad (105)$$

for $k, l = 0, \dots, p$. We can use these derivatives directly to compute ϕ_c given in (6).

Following in the same approach as Section J.1.2, we can compute the bounds of ϕ_c (in (6)) by utilising the bounds provided in (17) and (18). As noted in Section J.1.2, we must be able to find an upper bound on the matrix norm of $\nabla^2 \log f_c^{(z)}(\mathbf{z})$ where $\mathbf{z} := \boldsymbol{\Lambda}_c^{-\frac{1}{2}} \boldsymbol{\beta}$, i.e. compute (84). To do so, we can compute the matrix norm of the matrix which bounds $\nabla^2 \log f_c^{(z)}(\mathbf{z})$ element-wise. Now, let $\mathbf{z} = \boldsymbol{\Lambda}_c^{-\frac{1}{2}} \boldsymbol{\beta}$ then $f_c^{(z)}(\mathbf{z}) = \pi(\boldsymbol{\beta}|X, \mathbf{y})|J|$, where $J = \boldsymbol{\Lambda}_c^{-\frac{1}{2}}$ is the Jacobian matrix, so we have

$$\begin{aligned} f_c^{(z)}(\mathbf{z}) &:= \pi(\boldsymbol{\beta}|X, \mathbf{y}) \cdot |\boldsymbol{\Lambda}_c^{-\frac{1}{2}}| \\ &= \left[\prod_{i=1}^n \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\pi\nu\sigma}} \left(1 + \frac{1}{\nu} \left(\frac{y_i - X_i(\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})}{\sigma}\right)^2\right)^{-\left(\frac{\nu+1}{2}\right)} \right] \\ &\quad \cdot \left[\prod_{j=0}^p \frac{1}{\sqrt{2\pi C\sigma_{\beta_j}^2}} \exp\left(-\frac{\left((\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})_j - \mu_j\right)^2}{2C\sigma_{\beta_j}^2}\right) \right] \cdot |\boldsymbol{\Lambda}_c^{-\frac{1}{2}}|. \end{aligned} \quad (106)$$

so

$$\begin{aligned} \log f_c^{(z)}(\mathbf{z}) &= - \left(\frac{\nu+1}{2} \right) \sum_{i=1}^n \log \left(1 + \frac{1}{\nu\sigma^2} \left(y_i - X_i(\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}) \right)^2 \right) \\ &\quad - \sum_{j=0}^p \frac{\left((\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})_j - \mu_j \right)^2}{2C\sigma_{\beta_j}^2} + \text{constant}. \end{aligned} \quad (107)$$

Recall from (94) and (95), then The first derivative of the log transformed posterior with respect to β_k for $k = 0, \dots, p$ is given by

$$\begin{aligned} \frac{\partial \log f_c^{(z)}(\mathbf{z})}{\partial z_k} &= \left(\frac{\nu+1}{2} \right) \sum_{i=1}^n \frac{-\frac{2(X\boldsymbol{\Lambda}_c^{\frac{1}{2}})_{ik}}{\nu\sigma^2} \left(y_i - (X_i\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}) \right)}{1 + \frac{1}{\nu\sigma^2} \left(y_i - (X_i\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}) \right)^2} - \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \left((\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})_j - \mu_j \right)}{C\sigma_{\beta_j}^2} \\ &= (\nu+1) \sum_{i=1}^n \frac{(X\boldsymbol{\Lambda}_c^{\frac{1}{2}})_{ik} \left(y_i - (X_i\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}) \right)}{\nu\sigma^2 + \left(y_i - (X_i\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}) \right)^2} - \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \left((\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})_j - \mu_j \right)}{C\sigma_{\beta_j}^2} \end{aligned} \quad (108)$$

Then the second order derivatives are given by

$$\frac{\partial^2 \log f_c^{(z)}(\mathbf{z})}{\partial z_k \partial z_l} = (\nu+1) \sum_{i=1}^n \frac{(X\boldsymbol{\Lambda}_c^{\frac{1}{2}})_{ik} (X\boldsymbol{\Lambda}_c^{\frac{1}{2}})_{il} \left(\left(y_i - (X_i\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}) \right)^2 - \nu\sigma^2 \right)}{\left(\left(y_i - (X_i\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}) \right)^2 + \nu\sigma^2 \right)^2} - \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \Lambda_{jl}^{\frac{1}{2}}}{C\sigma_{\beta_j}^2} \quad (109)$$

for $k, l = 0, \dots, p$.

J.2.1 GLOBAL BOUNDS OF P^{Λ_c}

To compute P^{Λ_c} for this example, first note that we can write

$$\frac{\partial^2 \log f_c^{(z)}(\mathbf{z})}{\partial z_k \partial z_l} = (\nu+1) \sum_{i=1}^n (X\boldsymbol{\Lambda}_c^{\frac{1}{2}})_{ik} (X\boldsymbol{\Lambda}_c^{\frac{1}{2}})_{il} \left[\frac{1}{E_i + b} - \frac{2b}{(E_i + b)^2} \right] - \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \Lambda_{jl}^{\frac{1}{2}}}{C\sigma_{\beta_j}^2}, \quad (110)$$

for $k, l = 0, \dots, p$, where $b = \nu\sigma^2$ and $E_i = \left(y_i - (X_i\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}) \right)^2$. Now let

$$K(E_i) = \frac{1}{E_i + b} - \frac{2b}{(E_i + b)^2},$$

then the derivative is given by

$$K'(E_i) = -\frac{1}{(E_i + b)^2} + \frac{4b}{(E_i + b)^3}.$$

Setting $K'(E_i) = 0$ gives $E_i = 3b$, and we have $K(E_i = 3b) = \frac{1}{8b}$. So the supremum of the second derivative is given by

$$\sup \left[\left| \frac{\partial^2 \log f_c^{(z)}(\mathbf{z})}{\partial z_k \partial z_l} \right| \right] = \frac{(\nu + 1)}{8\nu\sigma^2} \sum_{i=1}^n |X\Lambda_c^{\frac{1}{2}}|_{ik} \cdot |X\Lambda_c^{\frac{1}{2}}|_{il} - \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \Lambda_{jl}^{\frac{1}{2}}}{C\sigma_{\beta_j}^2}. \quad (111)$$

We can therefore use this to compute P^{Λ_c} to compute bounds for ϕ_c as per (17) and (18).

J.3 Negative Binomial Regression

In Section 5.3, we considered a negative Binomial regression example with Gaussian prior distributions for the parameters. In particular, our sub-posterior densities were given by the posterior density with $\mathcal{N}_d(\mu_j, C\sigma_{\beta_j}^2)$ priors for β_j for $j = 0, \dots, p$, is given by

$$\begin{aligned} f_c(\boldsymbol{\beta}) &:= \pi(\boldsymbol{\beta}|X, \mathbf{y}) \\ &= \left[\prod_{i=1}^n \frac{\Gamma(y_i + r)}{y_i! \Gamma(r)} \left(\frac{\mu_i}{\mu_i + r} \right)^{y_i} \left(\frac{r}{\mu_i + r} \right)^r \right] \cdot \left[\prod_{j=0}^p \frac{1}{\sqrt{2\pi C\sigma_{\beta_j}^2}} \exp\left(-\frac{(\beta_j - \mu_j)^2}{2C\sigma_{\beta_j}^2} \right) \right] \\ &= \left[\prod_{i=1}^n \frac{\Gamma(y_i + r)}{y_i! \Gamma(r)} \frac{\exp(X_i \boldsymbol{\beta} \cdot y_i) \cdot r^r}{(\exp(X_i \boldsymbol{\beta}) + r)^{y_i + r}} \right] \cdot \left[\prod_{j=0}^p \frac{1}{\sqrt{2\pi C\sigma_{\beta_j}^2}} \exp\left(-\frac{(\beta_j - \mu_j)^2}{2C\sigma_{\beta_j}^2} \right) \right] \end{aligned} \quad (112)$$

The log-posterior is given by

$$\log f_c(\boldsymbol{\beta}) = \sum_{i=1}^n [X_i \boldsymbol{\beta} \cdot y_i - (y_i + r) \log(\exp(X_i \boldsymbol{\beta}) + r)] - \sum_{j=0}^p \frac{(\beta_j - \mu_j)^2}{2C\sigma_{\beta_j}^2} + \text{constant}. \quad (113)$$

The first order derivative of the log-posterior with respect to β_k for $k = 0, \dots, p$, is given by

$$\begin{aligned} \frac{\partial \log f_c(\boldsymbol{\beta})}{\partial \beta_k} &= \sum_{i=1}^n \left[X_{ik} y_i - \frac{(y_i + r) X_{ik} \exp(X_i \boldsymbol{\beta})}{\exp(X_i \boldsymbol{\beta}) + r} \right] - \frac{(\beta_k - \mu_k)}{C\sigma_{\beta_k}^2}, \\ &= \sum_{i=1}^n \left[X_{ik} \cdot \left(y_i - \frac{(y_i + r) \exp(X_i \boldsymbol{\beta})}{\exp(X_i \boldsymbol{\beta}) + r} \right) \right] - \frac{(\beta_k - \mu_k)}{C\sigma_{\beta_k}^2}, \end{aligned} \quad (114)$$

and the second order derivatives of the log-posterior are given by

$$\frac{\partial^2 \log f_c(\boldsymbol{\beta})}{\partial \beta_k^2} = -\sum_{i=1}^n \frac{(y_i + r) r X_{ik}^2 \exp(X_i \boldsymbol{\beta})}{(\exp(X_i \boldsymbol{\beta}) + r)^2} - \frac{1}{C\sigma_{\beta_k}^2}, \quad (115)$$

$$\frac{\partial^2 \log f_c(\boldsymbol{\beta})}{\partial \beta_k \partial \beta_l} = - \sum_{i=1}^n \frac{(y_i + r) r X_{ik} X_{il} \exp(X_i \boldsymbol{\beta})}{(\exp(X_i \boldsymbol{\beta}) + r)^2} \quad \text{for } k \neq l, \quad (116)$$

for $k, l = 0, \dots, p$. We can use these directly to compute ϕ_c given in (6).

Following in the same approach as Section J.1.2, we can compute the bounds of ϕ_c (in (6)) by utilising the bounds provided in (17) and (18). As noted in Section J.1.2, we must compute (84). To do so, we can compute the matrix norm of the matrix which bounds $\nabla^2 \log f_c^{(z)}(\mathbf{z})$ element-wise. We have

$$\begin{aligned} f_c^{(z)}(\mathbf{z}) &:= \pi(\boldsymbol{\beta} | X, \mathbf{y}) \cdot |\boldsymbol{\Lambda}_c^{-\frac{1}{2}}| \\ &= \left[\prod_{i=1}^n \frac{\Gamma(y_i + r)}{y_i! \Gamma(r)} \frac{\exp\left(X_i(\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}) \cdot y_i\right) \cdot r^r}{\left(\exp\left(X_i(\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})\right) + r\right)^{y_i+r}} \right] \\ &\quad \cdot \left[\prod_{j=0}^p \frac{1}{\sqrt{2\pi C \sigma_{\beta_j}^2}} \exp\left(-\frac{\left((\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})_j - \mu_j\right)^2}{2C \sigma_{\beta_j}^2}\right) \right] \cdot |\boldsymbol{\Lambda}_c^{-\frac{1}{2}}|, \end{aligned} \quad (117)$$

and

$$\begin{aligned} \log f_c^{(z)}(\mathbf{z}) &= \sum_{i=1}^n \left[(X_i \boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z}) \cdot y_i - (y_i + r) \log\left(\exp\left((X_i \boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})\right) + r\right) \right] \\ &\quad - \sum_{j=0}^p \frac{\left((\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})_j - \mu_j\right)^2}{2C \sigma_{\beta_j}^2} + \text{constant}. \end{aligned} \quad (118)$$

Recall from (94) and (95) that we have $\frac{\partial(X_i \boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})}{\partial z_k} = (X \boldsymbol{\Lambda}_c^{\frac{1}{2}})_{ik}$ and $\frac{\partial(\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})_i}{\partial z_k} = \boldsymbol{\Lambda}_c^{\frac{1}{2}}_{ik}$. Then first derivative of the log transformed posterior with respect to β_k is given by

$$\begin{aligned} \frac{\partial \log f_c^{(z)}(\mathbf{z})}{\partial z_k} &= \sum_{i=1}^n \left[(X \boldsymbol{\Lambda}_c^{\frac{1}{2}})_{ik} \cdot \left(y_i - \frac{(y_i + r) \exp\left((X_i \boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})\right)}{\exp\left((X_i \boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})\right) + r} \right) \right] \\ &\quad - \sum_{j=0}^p \frac{\boldsymbol{\Lambda}_c^{\frac{1}{2}}_{jk} \left((\boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})_j - \mu_j\right)}{C \sigma_{\beta_j}^2} \end{aligned} \quad (119)$$

and the second order derivatives are given by

$$\frac{\partial^2 \log f_c^{(z)}(\mathbf{z})}{\partial z_k \partial z_l} = - \sum_{i=1}^n \frac{(y_i + r) r (X \boldsymbol{\Lambda}_c^{\frac{1}{2}})_{ik} (X \boldsymbol{\Lambda}_c^{\frac{1}{2}})_{il} \exp\left((X_i \boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})\right)}{\left(\exp\left((X_i \boldsymbol{\Lambda}_c^{\frac{1}{2}} \mathbf{z})\right) + r\right)^2} - \sum_{j=0}^p \frac{\boldsymbol{\Lambda}_c^{\frac{1}{2}}_{jk} \boldsymbol{\Lambda}_c^{\frac{1}{2}}_{jl}}{C \sigma_{\beta_j}^2}. \quad (120)$$

To find bounds for r_c , we must now try to find bounds on the second derivatives given above and compute the matrix norm of the matrix made up of these bounds (which ultimately bounds $\nabla^2 \log f_c^{(z)}(\mathbf{z})$ element-wise). For this example, we can find global and lower bounds of the second derivatives. Note however, we typically will expect better performance with the local bounds on P^{Λ_c} (as this will typically lead to the expected number of points we need to evaluate while performing Poisson thinning, κ_c , to be lower) despite these bounds being slightly more expensive to compute in practice.

J.3.1 GLOBAL BOUNDS OF P^{Λ_c}

Note that $\frac{e^{ax}}{(e^{ax}+r)^2} \leq \frac{1}{4r}$ for all x (where a is some constant), so we can use this to obtain global bounds on the matrix norm in the transformed space. Note that this maximum occurs at $x = \frac{1}{a} \log(r)$. To find global bounds, we can use

$$\sup \left[\left| \frac{\partial^2 \log f_c^{(z)}(\mathbf{z})}{\partial z_k \partial z_l} \right| \right] = \sum_{i=1}^n \frac{(y_i + r)r |X \Lambda_c^{\frac{1}{2}}|_{ik} \cdot |X \Lambda_c^{\frac{1}{2}}|_{il}}{4r} + \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \Lambda_{jl}^{\frac{1}{2}}}{C \sigma_{\beta_j}^2}. \quad (121)$$

J.3.2 LOCAL BOUNDS OF P^{Λ_c}

Local bounds can be obtained if we can find local bounds for

$$G_r(\mathbf{z}) := \frac{\exp\left((X_i \Lambda_c^{\frac{1}{2}}) \mathbf{z}\right)}{\left(\exp\left((X_i \Lambda_c^{\frac{1}{2}}) \mathbf{z}\right) + r\right)^2} \quad (122)$$

for $\mathbf{z} \in R^{(z)}$ and $i = 1, \dots, n$. In that case, we have

$$\sup_{\mathbf{z} \in R^{(z)}} \left[\left| \frac{\partial^2 \log f_c^{(z)}(\mathbf{z})}{\partial z_k \partial z_l} \right| \right] = \sum_{i=1}^n \left[(y_i + r)r |X \Lambda_c^{\frac{1}{2}}|_{ik} \cdot |X \Lambda_c^{\frac{1}{2}}|_{il} \cdot \max_{\mathbf{z} \in R^{(z)}} \{G_r(\mathbf{z})\} \right] + \sum_{j=0}^p \frac{\Lambda_{jk}^{\frac{1}{2}} \Lambda_{jl}^{\frac{1}{2}}}{C \sigma_{\beta_j}^2}. \quad (123)$$

We can obtain bounds for $G_r(\mathbf{z})$ by noting that $\frac{\exp(x)}{(r+\exp(x))^2} \leq \frac{1}{4r}$ for all x and this maximum is attained at $x = \log(r)$. Further note that $\frac{\exp(x)}{(r+\exp(x))^2} \leq \frac{1}{4r}$ is a uni-modal function (with mode at $x = \log(r)$ as noted). Now let

$$F_i(\mathbf{z}) := (X_i \Lambda_c^{\frac{1}{2}}) \mathbf{z} = \sum_{j=1}^d (X_i \Lambda_c^{\frac{1}{2}})_j \mathbf{z}_j, \quad (124)$$

then let $F_i^\downarrow := \min_{\mathbf{z} \in R^{(z)}} F_i(\mathbf{z})$ and $F_i^\uparrow := \max_{\mathbf{z} \in R^{(z)}} F_i(\mathbf{z})$ denote the minimum and maximum of $F_i(\mathbf{z})$ for $\mathbf{z} \in R^{(z)}$ respectively. Then we note that this can simply be computed in

with a linear cost with d . Now, noting that $F_i(\mathbf{z})$ is linear and $\frac{\exp(x)}{(r+\exp(x))^2} \leq \frac{1}{4r}$ is uni-modal, after computing F_i^\downarrow and F_i^\uparrow , there are two cases:

1. If we have $\log(r) \in [F_i^\downarrow, F_i^\uparrow]$, then we know that for this hypercube $R^{(z)}$, we will attain the maximum $\frac{1}{4r}$.
2. If $\log(r) \notin [F_i^\downarrow, F_i^\uparrow]$, then the maximum of $G_r(x)$ occurs at which ever point is the closest to $\log(r)$.

Therefore local bounds can be obtained by minimising and maximising $F_i(\mathbf{z})$ for $\mathbf{z} \in R^{(z)}$. If this interval includes $\log(r)$, then the local maximum attains the global maximum, otherwise, the local maximum occurs at either of these intervals (which ever is closer to $\log(r)$).

This method for finding local bounds requires two optimisations of $F_i(\mathbf{z})$, but we note that we can actually obtain the bounds by only performing one optimisation. In particular, we can evaluate $F_i(\mathbf{z})$ at any arbitrary value $\hat{\mathbf{z}} \in R^{(z)}$ (we can simply take this to be the centre of the hypercube). If we have $F_i(\hat{\mathbf{z}}) > \log(r)$, then we just need only need minimise the function $F_i(\mathbf{z})$, since if we have $F_i^\downarrow < \log(r)$, then we know that $\log(r) \in [F_i^\downarrow, F_i^\uparrow]$, so the global maximum is attained. If $F_i^\downarrow > \log(r)$, then the maximum of $G(x)$ just occurs at F_i^\downarrow and we can avoid the need to maximise the function $F_i(\mathbf{z})$. However, if conversely, we evaluate $F_i(\mathbf{z})$ at $\mathbf{z} = \hat{\mathbf{z}}$ and we have $F_i(\hat{\mathbf{z}}) < \log(r)$, then we just need to maximise $F_i(\mathbf{z})$ for $\mathbf{z} \in R^{(z)}$ and apply the inverse of the same trick. To summarise, in order to find $\max_{\mathbf{z} \in R^{(z)}} \{G_r(\mathbf{z})\}$, we can apply Algorithm 8.

Algorithm 8 Computing the local bounds of $G_r(\mathbf{z})$ given in (122) for $\mathbf{z} \in R^{(z)}$.

1. Compute $F_i(\hat{\mathbf{z}})$ at some arbitrary value $\hat{\mathbf{z}} \in R^{(z)}$.
 2. If $F_i(\hat{\mathbf{z}}) > \log(r)$:
 - (a) Compute $F_i^\downarrow := \min_{\mathbf{z} \in R^{(z)}} F_i(\mathbf{z})$.
 - (b) $\max_{\mathbf{z} \in R^{(z)}} \left\{ \frac{\exp\left((X_i \Lambda_c^{\frac{1}{2}}) \mathbf{z}\right)}{\left(\exp\left((X_i \Lambda_c^{\frac{1}{2}}) \mathbf{z}\right) + r\right)^2} \right\} = \begin{cases} \frac{1}{4r} & \text{if } F_i^\downarrow < \log(r), \\ G(F_i^\downarrow) = \frac{\exp(F_i^\downarrow)}{(\exp(F_i^\downarrow) + r)^2} & \text{otherwise.} \end{cases}$
 3. Else (if $F_i(\hat{\mathbf{z}}) < \log(r)$):
 - (a) Compute $\max_{\mathbf{z} \in R^{(z)}} F_i(\mathbf{z})$.
 - (b) $\max_{\mathbf{z} \in R^{(z)}} \left\{ \frac{\exp\left((X_i \Lambda_c^{\frac{1}{2}}) \mathbf{z}\right)}{\left(\exp\left((X_i \Lambda_c^{\frac{1}{2}}) \mathbf{z}\right) + r\right)^2} \right\} = \begin{cases} \frac{1}{4r} & \text{if } F_i^\uparrow > \log(r), \\ G(F_i^\uparrow) = \frac{\exp(F_i^\uparrow)}{(\exp(F_i^\uparrow) + r)^2} & \text{otherwise.} \end{cases}$
-

References

- Isabelle Albert, Sophie Donnet, Chantal Guihenneuc-Jouyaux, Samantha Low-Choy, Kerrie Mengersen, and Judith Rousseau. Combining Expert Opinions in Prior Elicitation. *Bayesian Analysis*, 7(3):503–532, 2012.
- Christophe Andrieu and Gareth O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697–725, 2009.
- Jack Baker, Paul Fearnhead, Emily B. Fox, and Christopher Nemeth. Control variates for stochastic gradient MCMC. *Statistics and Computing*, 29(3):599–615, 2019.
- James O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer, New York, 1980.
- Alexandros Beskos, Omiros Papaspiliopoulos, and Gareth O. Roberts. Retrospective exact simulation of diffusion sample paths with applications. *Bernoulli*, 12(6):1077–1098, 2006.
- Alexandros Beskos, Omiros Papaspiliopoulos, and Gareth O. Roberts. A Factorisation of Diffusion Measure and Finite Sample Path Constructions. *Methodology and Computing in Applied Probability*, 10(1):85–104, 2008.
- Joris Bierkens, Paul Fearnhead, and Gareth O. Roberts. The Zig-Zag Process and Super-Efficient Sampling for Bayesian Analysis of Big Data. *The Annals of Statistics*, 47(3):1288–1320, 2019.
- Alexandre Bouchard-Côté, Sebastian J. Vollmer, and Arnaud Doucet. The Bouncy Particle Sampler: A Non-Reversible Rejection-Free Markov Chain Monte Carlo Method. *Journal of the American Statistical Association*, 113(522):855–867, 2018.
- Bob Carpenter, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), 2017.
- Didier Dacunha-Castelle and Danielle Florens-Zmirou. Estimation of the Coefficients of a Diffusion from Discrete Observations. *Stochastics: An International Journal of Probability and Stochastic Processes*, 19(4):263–284, 1986.
- Hongsheng Dai, Murray Pollock, and Gareth O. Roberts. Monte Carlo Fusion. *Journal of Applied Probability*, 56(1):174–191, 2019.
- Hongsheng Dai, Murray Pollock, and Gareth O. Roberts. Bayesian Fusion: Scalable unification of distributed statistical analyses. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(1):84–107, 2023.
- Arnaud Doucet and Anthony Lee. Sequential Monte Carlo Methods. In *Handbook of Graphical Models*, pages 165–188. CRC Press, 2018.
- Dirk Eddelbuettel. *Seamless R and C++ Integration with Rcpp*. Springer, New York, 2013.

- Hadi Fanaee-T and Joao Gama. Event labeling combining ensemble detectors and background knowledge. *Progress in Artificial Intelligence*, 2(2):113–127, 2014.
- Paul Fearnhead, Omiros Papaspiliopoulos, and Gareth O. Roberts. Particle Filters for Partially Observed Diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(4):755–777, 2008.
- Axel Finke, Arnaud Doucet, and Adam M. Johansen. Limit Theorems for Sequential MCMC Methods. *Advances in Applied Probability*, 52(2):377–403, 2020.
- Joseph L. Fleiss. The statistical basis of meta-analysis. *Statistical Methods in Medical Research*, 2(2):121–145, 1993.
- Christian Genest and James V. Zidek. Combining Probability Distributions: A Critique and an Annotated Bibliography. *Statistical Science*, 1(1):114–135, 1986.
- Mathieu Gerber, Nicolas Chopin, and Nick Whiteley. Negative association, ordering and convergence of resampling methods. *Annals of Statistics*, 47(4):2236–2260, 2019.
- Robert J.B. Goudie, Anne M. Presanis, David Lunn, Daniela De Angelis, and Lorenz Wernisch. Joining and splitting models with Markov Melding. *Bayesian Analysis*, 14(1):81–109, 2019.
- Tomoyuki Higuchi. Monte Carlo filter using the genetic algorithm operators. *Journal of Statistical Computation and Simulation*, 59(1):1–23, 1997.
- Richard A. Johnson. Asymptotic expansions associated with posterior distributions. *The Annals of Mathematical Statistics*, pages 851–864, 1970.
- Heysem Kaya, Pınar Tüfekci, and Fikret S. Gürgen. Local and global learning methods for predicting power of a combined gas & steam turbine. In *Proceedings of the International Conference on Emerging Trends in Computer and Electronics Engineering ICETCEE*, pages 13–18, 2012.
- Augustine Kong, Jun S. Liu, and Wing Hung Wong. Sequential Imputations and Bayesian Missing Data Problems. *Journal of the American Statistical Association*, 89(425):278–288, 1994.
- Xiangshun Kong and Wei Zheng. Design Based Incomplete U-Statistics. *Statistica Sinica*, 31:1593–1618, 2021.
- Juan Kuntz, Francesca R. Crucinio, and Adam M. Johansen. Product-form estimators: exploiting independence to scale up Monte Carlo. *Statistics and Computing*, 32(12):1–22, 2022.
- Juan Kuntz, Francesca R. Crucinio, and Adam M. Johansen. The divide-and-conquer sequential Monte Carlo algorithm: Theoretical properties and limit theorems. *Annals of Applied Probability*, 2023. In press.
- Lucien Le Cam. *Asymptotic Methods in Statistical Decision Theory*. Springer Science & Business Media, New York, 1986.

- Lucien Le Cam and Grace Lo Yang. *Asymptotics in Statistics: Some Basic Concepts*. Springer Science & Business Media, New York, 2000.
- Fredrik Lindsten, Adam M. Johansen, Christian A. Naesseth, Bonnie Kirkpatrick, Thomas B. Schön, John A.D. Aston, and Alexandre Bouchard-Côté. Divide-and-Conquer with Sequential Monte Carlo. *Journal of Computational and Graphical Statistics*, 26(2): 445–458, 2017.
- Jun S. Liu and Rong Chen. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association*, 93(443):1032–1044, 1998.
- Stanislav Minsker, Sanvesh Srivastava, Lizhen Lin, and David B. Dunson. Scalable and Robust Bayesian Inference via the Median Posterior. In *International Conference on Machine Learning*, pages 1656–1664, 2014.
- Alexey Miroshnikov and Erin M. Conlon. ParallelMCMCcombine: an R package for Bayesian Methods for Big Data and Analytics. *PloS one*, 9(9):e108425, 2014.
- Willie Neiswanger, Chong Wang, and Eric P. Xing. Asymptotically Exact, Embarrassingly Parallel MCMC. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, UAI’14, page 623–632, Arlington, Virginia, USA, 2014. AUAI Press.
- Christopher Nemeth and Chris Sherlock. Merging MCMC Subposteriors through Gaussian-process Approximations. *Bayesian Analysis*, 13(2):507–530, 2018.
- Murray Pollock, Adam M. Johansen, and Gareth O. Roberts. On the exact and ε -strong simulation of (jump) diffusions. *Bernoulli*, 22(2):794–856, 2016.
- Murray Pollock, Paul Fearnhead, Adam M. Johansen, and Gareth O. Roberts. Quasi-stationary Monte Carlo and the ScaLE algorithm. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 82(5):1167–1221, 2020.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2022. URL <http://www.R-project.org/>.
- L. Chris G. Rogers and David Williams. *Diffusions, Markov processes and martingales: Volume 2, Itô calculus*, volume 2. Cambridge University Press, Cambridge, 2000.
- Steven L. Scott, Alexander W. Blocker, Fernando V. Bonassi, Hugh A. Chipman, Edward I. George, and Robert E. McCulloch. Bayes and Big Data: The Consensus Monte Carlo Algorithm. *International Journal of Management Science and Engineering Management*, 11(2):78–88, 2016.
- Sanvesh Srivastava, Volkan Cevher, Quoc Dinh, and David B. Dunson. WASP: Scalable Bayes via barycenters of subset posteriors. In *Artificial Intelligence and Statistics*, pages 912–920, 2015.
- Pınar Tüfekci. Prediction of full load electrical power output of a base load operated combined cycle power plant using machine learning methods. *International Journal of Electrical Power & Energy Systems*, 60:126–140, 2014.

- A. W. Van der Vaart. *Asymptotic Statistics*. Cambridge University Press, Cambridge, 1998.
- Andrew M. Walker. On the Asymptotic Behaviour of Posterior Distributions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 31(1):80–88, 1969.
- Xiangyu Wang and David B. Dunson. Parallelizing MCMC via Weierstrass Sampler. *Statistics e-print 1312.4605*, arXiv, 2013.
- Darrell Whitley. A genetic algorithm tutorial. *Statistics and Computing*, 4(2):65–85, 1994.
- Hadley Wickham. *nycflights13: Flights that Departed NYC in 2013*, 2021.
- Sinan Yıldırım and Beyza Ermiş. Exact MCMC with differentially private moves. *Statistics and Computing*, 29(5):947–963, 2019.