## Coupling Particle Systems

Pierre E. Jacob (Harvard) joint work with
Fredrik Lindsten and Thomas Schön (Uppsala)

CRiSM Workshop on Estimating Constants April 20, 2016

## Outline

1 Motivation for coupling particle filters

2 How to couple two particle filters

3 A new smoothing algorithm

## Hidden Markov models



Figure: Graph representation of a general hidden Markov model.
$\left(X_{t}\right)$ : initial $\mu_{\theta}$, transition $f_{\theta} .\left(Y_{t}\right)$ given $\left(X_{t}\right)$ : measurement $g_{\theta}$.

## Hidden Markov models

- How to estimate/predict the latent process $\left(X_{t}\right)$ given the observations $\left(Y_{t}\right)$ and a fixed parameter $\theta$ ?
- How to estimate the parameter $\theta$ ?


## Example: Hidden Autoregressive

- Hidden process $X_{t}=A X_{t-1}+\varepsilon_{t}$, where $\varepsilon_{t} \sim \mathcal{N}_{d}(0, I)$, $X_{0} \sim \mathcal{N}_{d}(0, I)$.
- $A_{i j}=\theta^{|i-j|+1}$ for $i, j \in 1: d$.
- Observations $Y_{t}=X_{t}+\eta_{t}$, where $\eta_{t} \sim \mathcal{N}_{d}(0, I)$.
taken from Guarniero, Johansen \& Lee, 2015.


## Example: Phytoplankton-Zooplankton



Figure: A time series of 365 observations generated according to a phytoplankton-zooplankton model.

## Example: Phytoplankton-Zooplankton

■ Hidden process $\left(X_{t}\right)=\left(\alpha_{t}, p_{t}, z_{t}\right)$.

- At each (integer) time, $\alpha_{t} \sim \mathcal{N}\left(\mu_{\alpha}, \sigma_{\alpha}^{2}\right)$.
- Given $\alpha_{t}$,

$$
\begin{aligned}
\frac{d p_{t}}{d t} & =\alpha_{t} p_{t}-c p_{t} z_{t} \\
\frac{d z_{t}}{d t} & =e c p_{t} z_{t}-m_{l} z_{t}-m_{q} z_{t}^{2}
\end{aligned}
$$

- Observations: $\log Y_{t} \sim \mathcal{N}\left(\log p_{t}, \sigma_{y}^{2}\right)$.
- Initial distribution: $\left(\log p_{0}, \log z_{0}\right) \sim \mathcal{N}\left(\mu_{0}, \sigma_{0}^{2}\right)$.

■ Unknown parameters: $\theta=\left(\mu_{0}, \sigma_{0}, \mu_{\alpha}, \sigma_{\alpha}, \sigma_{y}, c, e, m_{l}, m_{q}\right)$.

## Particle filter



## Particle filter



## Particle filter



## Particle filter



## Particle filter



## Particle filter



## Particle filter



## Particle filter

At step $t=0$,
$\llbracket$ Sample $x_{0}^{k} \sim \mu_{\theta}\left(d x_{0}\right)$, for all $k \in 1: N$.

2 Set $w_{0}^{k}=N^{-1}$, for all $k \in 1: N$.

At step $t \geq 1$,
$\llbracket$ Sample ancestors $a_{t}^{1: N} \sim r\left(d a^{1: N} \mid w_{t-1}^{1: N}\right) . \quad \leftarrow$ resampling
2 Sample $x_{t}^{k} \sim f_{\theta}\left(d x_{t} \mid x_{t-1}^{a_{t}^{k}}\right)$, for all $k \in 1: N$.
3 Compute $w_{t}^{k}=g_{\theta}\left(y_{t} \mid x_{t}^{k}\right)$, for all $k \in 1: N$.

## Particle filter, rewritten

At step $t=0$,
$\boldsymbol{1}$ Sample $U_{M}^{k}$, compute $x_{0}^{k}=M\left(U_{M}^{k}, \theta\right)$, for all $k \in 1: N$.
2 Set $w_{0}^{k}=N^{-1}$, for all $k \in 1: N$.

At step $t \geq 1$,
1 Sample ancestors $a_{t}^{1: N} \sim r\left(d a^{1: N} \mid w_{t-1}^{1: N}\right)$.
$\leftarrow$ resampling
$\boldsymbol{2}$ Sample $U_{F, t}^{k}$, compute $x_{t}^{k}=F\left(x_{t-1}^{a_{t}^{k}}, U_{F, t}^{k}, \theta\right)$, for all $k \in 1: N$.
3 Compute $w_{t}^{k}=g_{\theta}\left(y_{t} \mid x_{t}^{k}\right)$, for all $k \in 1: N$.

## Output

- Approximation of the filtering distributions

$$
\forall t \in\{1, \ldots, T\} \quad p\left(d x_{t} \mid y_{1: t}, \theta\right)
$$

by

$$
\forall t \in\{1, \ldots, T\} \quad p^{N}\left(d x_{t} \mid y_{1: t}, \theta\right)=\sum_{k=1}^{N} w_{t}^{k} \delta_{x_{t}^{k}}\left(d x_{t}\right)
$$

- Approximation of the likelihood function $\mathcal{L}(\theta)=p\left(y_{1: T} \mid \theta\right)$ by

$$
p^{N}\left(y_{1: T} \mid \theta\right)=\prod_{t=1}^{T} \frac{1}{N} \sum_{k=1}^{N} w_{t}^{k}
$$

## Idea

- Particle filters are increasingly used as parts of encompassing algorithms. e.g. Particle MCMC, Iterated Filtering
- Some of these algorithms compare the outputs of multiple particle filters.
- Better algorithms can be obtained by correlating particle filters.
i.e. correlation helps comparison.


## Example: approximation of the likelihood



Figure: Estimates of the log-likelihood obtained by particle filters, in a hidden auto-regressive model, $T=100$ observations, $N=64$ particles.

See Pitt \& Malik, 2011, Lee 2008.

## Example: finite difference

A simple estimator of $\nabla \ell(\theta)=\nabla \log \mathcal{L}(\theta)$ is:

$$
\widehat{\nabla \ell}(\theta)=\frac{\log \hat{p}^{N}\left(y_{1: T} \mid \theta+h\right)-\log \hat{p}^{N}\left(y_{1: T} \mid \theta-h\right)}{2 h}
$$

The two log-likelihood estimators can be obtained using independent particle filters given $\theta+h$ and $\theta-h \ldots$
... but if we could positively correlate the two log-likelihood estimators, the variance of $\widehat{\nabla \ell}(\theta)$ would be smaller.

## Example: finite difference



Figure: Standard deviation of $\widehat{\nabla \ell}(\theta)$, for some $\theta$, in a hidden auto-regressive model, $T=100$ observations, $N=128$ particles.

## Example: finite difference



Figure: Same but in dimension 5.

## Example: finite difference



Figure: Same but in dimension 10.

## Example: Metropolis-Hastings

Assume we can compute the target density $\pi\left(\theta \mid y_{1: T}\right)$ pointwise.

1: Set some $\theta^{(1)}$.
2: for $i=2$ to $M$ do
3: Propose $\theta^{\star} \sim q\left(\cdot \mid \theta^{(i-1)}\right)$.
4: Compute the ratio:

$$
\alpha=\min \left(1, \frac{\pi\left(\theta^{\star}\right) p\left(y_{1: T} \mid \theta^{\star}\right)}{\pi\left(\theta^{(i-1)}\right) p\left(y_{1: T} \mid \theta^{(i-1)}\right)} \frac{q\left(\theta^{(i-1)} \mid \theta^{\star}\right)}{q\left(\theta^{\star} \mid \theta^{(i-1)}\right)}\right) .
$$

5: $\quad$ Set $\theta^{(i)}=\theta^{\star}$ with probability $\alpha$, otherwise set $\theta^{(i)}=\theta^{(i-1)}$.
6: end for

## Example: particle Metropolis-Hastings

Assume we can run a particle filter to get $p^{N}\left(y_{1: T} \mid \theta\right)$.

1: Set some $\theta^{(1)}$ and sample $p^{N}\left(y_{1: T} \mid \theta^{(1)}\right)$.
2: for $i=2$ to $M$ do
3: Propose $\theta^{\star} \sim q\left(\cdot \mid \theta^{(i-1)}\right)$ and sample $p^{N}\left(y_{1: T} \mid \theta^{\star}\right)$.
4: Compute the ratio:

$$
\alpha=\min \left(1, \frac{\pi\left(\theta^{\star}\right) p^{N}\left(y_{1: T} \mid \theta^{\star}\right)}{\pi\left(\theta^{(i-1)}\right) p^{N}\left(y_{1: T} \mid \theta^{(i-1)}\right)} \frac{q\left(\theta^{(i-1)} \mid \theta^{\star}\right)}{q\left(\theta^{\star} \mid \theta^{(i-1)}\right)}\right) .
$$

5: $\quad \operatorname{Set} \theta^{(i)}=\theta^{\star}$ with probability $\alpha$, otherwise set $\theta^{(i)}=\theta^{(i-1)}$. 6: end for

Andrieu, Doucet \& Holenstein, 2010.

## Example: particle Metropolis-Hastings

- The acceptance ratio involves a ratio of particle filter estimators:

$$
\alpha=\min \left(1, \frac{\pi\left(\theta^{\star}\right) p^{N}\left(y_{1: T} \mid \theta^{\star}\right)}{\pi\left(\theta^{(i-1)}\right) p^{N}\left(y_{1: T} \mid \theta^{(i-1)}\right)} \frac{q\left(\theta^{(i-1)} \mid \theta^{\star}\right)}{q\left(\theta^{\star} \mid \theta^{(i-1)}\right)}\right) .
$$

- If we positively correlate $p^{N}\left(y_{1: T} \mid \theta^{\star}\right)$ and $p^{N}\left(y_{1: T} \mid \theta^{(i-1)}\right)$, the ratio of estimators becomes more precise.

Deligiannidis, Doucet, Pitt \& Kohn, 2015: epic improvements for the case large $T / \operatorname{small} N$.

## Outline

## 1 Motivation for coupling particle filters

2 How to couple two particle filters

## 3 A new smoothing algorithm

## Coupled particle filter

By coupled particle filters we mean ...

- two particle systems, $\left(w_{t}^{k}, x_{t}^{k}\right)_{k=1}^{N}$ and $\left(\tilde{w}_{t}^{k}, \tilde{x}_{t}^{k}\right)_{k=1}^{N}$,
- conditioned on $\theta$ and $\tilde{\theta}$ respectively,

■ using common random numbers $U_{M}$ and $U_{F}$ for the initial and propagation steps.

One still has the freedom to choose a "coupled resampling" scheme.

## Coupled particle filter

At step $t=0$,
1 Sample $U_{M}^{k}$, and compute, for all $k \in 1: N$,

$$
x_{0}^{k}=M\left(U_{M}^{k}, \theta\right) \text { and } \tilde{x}_{0}^{k}=M\left(U_{M}^{k}, \tilde{\theta}\right)
$$

2 Set $w_{0}^{k}=N^{-1}$ and $\tilde{w}_{0}^{k}=N^{-1}$, for all $k \in 1: N$.
At step $t \geq 1$,
1 Sample ancestors:

$$
\left(a_{t}, \tilde{a}_{t}\right) \sim \bar{r}\left(\cdot \mid w_{t-1}^{1: N}, \tilde{w}_{t-1}^{1: N}\right) . \quad \leftarrow \text { coupled resampling }
$$

2 Sample $U_{F, t}^{k}$, and compute, for all $k \in 1: N$,

$$
x_{t}^{k}=F\left(x_{t-1}^{a_{t}^{k}}, U_{F, t}^{k}, \theta\right) \text { and } \tilde{x}_{t}^{k}=F\left(\tilde{x}_{t-1}^{\tilde{a}_{t}^{k}}, U_{F, t}^{k}, \tilde{\theta}\right)
$$

3 Compute $w_{t}^{k}=g\left(y_{t} \mid x_{t}^{k}, \theta\right)$ and $\tilde{w}_{t}^{k}=g\left(y_{t} \mid \tilde{x}_{t}^{k}, \tilde{\theta}\right)$.

## Coupled resampling

Given two particle systems, $\left(w^{k}, x^{k}\right)_{k=1}^{N}$ and $\left(\tilde{w}^{k}, \tilde{x}^{k}\right)_{k=1}^{N} \ldots$

- We want (?) to sample $a^{1: N}$ and $\tilde{a}^{1: N}$ in $\{1, \ldots, N\}^{N}$ such that

$$
\forall k \quad \forall j \quad \mathbb{P}\left(a^{k}=j\right)=w^{j} \text { and } \mathbb{P}\left(\tilde{a}^{k}=j\right)=\tilde{w}^{j}
$$

- Equivalently, we want to sample $\left(a^{k}, \tilde{a}^{k}\right)_{k=1}^{N}$ from a probability matrix $P$ such that

$$
P \mathbb{1}=w \quad \text { and } \quad P^{T} \mathbb{1}=\tilde{w}
$$

Independent resampling corresponds to $P=w \tilde{w}^{T}$. What else?

## Transport resampling

- Suppose that we want to sample a couple $(a, \tilde{a})$, from some probability matrix $P$, such that the resampled particles, $x^{a}$ and $\tilde{x}^{\tilde{a}}$, are as similar as possible.
- Similarity can be encoded by a distance $d$ on the space of $x$.
- The expected distance between $x^{a}$ and $\tilde{x}^{\tilde{a}}$, conditional upon the particles, is given by

$$
\mathbb{E}\left[d\left(x^{a}, \tilde{x}^{\tilde{a}}\right)\right]=\sum_{i=1}^{N} \sum_{j=1}^{N} P_{i j} d\left(x^{i}, \tilde{x}^{j}\right) .
$$

## Transport resampling

- Introduce $\mathcal{J}(w, \tilde{w})$, the set of matrices satisfying

$$
P \mathbb{1}=w \quad \text { and } \quad P^{T} \mathbb{1}=\tilde{w}
$$

■ Compute $D=\left(d\left(x^{i}, \tilde{x}^{j}\right)\right)_{i, j=1}^{N}$, for a cost of $\mathcal{O}\left(N^{2}\right)$.

■ Optimal transport problem: solving

$$
P^{\star}=\inf _{P \in \mathcal{J}(w, \tilde{w})} \sum_{i=1}^{N} \sum_{j=1}^{N} P_{i j} D_{i j} .
$$

## Transport resampling

- Sampling from the optimal $P^{\star}$ minimizes the expected distance between the two sets of particles, under the marginal constraint.
- (Which is not exactly the same as maximizing the correlation between e.g. likelihood estimators).
- Computing $P^{\star}$ requires $\mathcal{O}\left(N^{3}\right)$ operations, but efficient approximations have been proposed (Cuturi 2013 and following work) in $\mathcal{O}\left(N^{2}\right)$.
- The cost is linear in the dimension of $x$, and independent of the model.


## Index-matching resampling

- At the initial step,

$$
x_{0}^{k}=M\left(U_{M}^{k}, \theta\right) \text { and } \tilde{x}_{0}^{k}=M\left(U_{M}^{k}, \tilde{\theta}\right)
$$

so that particles with the same index are similar (if $M$ is continuous in $\theta$ ).

- At subsequent steps, the same random numbers $U_{F}^{k}$ are used to propagate $x^{a^{k}}$ and $\tilde{x}^{\tilde{a}^{k}}$.
- Standard resampling breaks the correspondence between similarity and indices: $x^{a^{k}}$ and $\tilde{x}^{\tilde{a}^{k}}$ might not be similar.


## Index-matching resampling

- To preserve the correspondence between indices, we want to maximize the probability of drawing couples of ancestors such that $a=\tilde{a}$.
- ...i.e. we want $P$ in $\mathcal{J}(w, \tilde{w})$ with maximum trace.
- We can define:

$$
P=\operatorname{diag}(\min \{w, \tilde{w}\})+(1-\alpha) r \tilde{r}^{T},
$$

with

$$
\begin{aligned}
\alpha & =\sum_{k=1}^{N} \min \left\{w^{k}, \tilde{w}^{k}\right\}, \\
r & =(w-\min \{w, \tilde{w}\}) /(1-\alpha), \\
\tilde{r} & =(\tilde{w}-\min \{w, \tilde{w}\}) /(1-\alpha) .
\end{aligned}
$$

- $P$ has maximum trace in $\mathcal{J}(w, \tilde{w})$ : we cannot augment its diagonal without violating the marginal constraints.


## Index-matching resampling



## Index-matching resampling

Distance between 1,000 pairs of particles, sampled independently and then propagated with common random numbers.


Hidden $\mathrm{AR}, \theta=0.30, \tilde{\theta}=0.31$.

## Index-matching resampling

Distance between 1,000 pairs of particles, sampled independently and then propagated with common random numbers.


Hidden $\mathrm{AR}, \theta=0.30, \tilde{\theta}=0.40$.

## Sorting and resampling

Univariate setting: $x$ is of dimension 1 .

- Sort the two systems $\left(x^{k}\right)_{k=1}^{N}$ and $\left(\tilde{x}^{k}\right)_{k=1}^{N}$.
- Perform e.g. systematic resampling on each sorted system, using the same random numbers.
- Thus if $a^{k}$ selects $x^{j}$ in the first system, $\tilde{a}^{k}$ is likely to select a $\tilde{x}^{i}$ close to $x^{j}$.

This can be extended to multivariate settings by sorting the particles according to the Hilbert space-filling curve. See Deligiannidis, Doucet, Pitt \& Kohn, 2015, and Pitt \& Malik, 2011, Gerber \& Chopin, 2015.

## Outline

## 1 Motivation for coupling particle filters

2 How to couple two particle filters

3 A new smoothing algorithm

## Who cares about unbiased estimators?

- Coupled resampling leads to a practical unbiased estimator $H_{u}$ of the smoothing quantity

$$
\int h\left(x_{0: T}\right) p\left(d x_{0: T} \mid y_{1: T}, \theta\right)
$$

for a test function $h$ (for fixed $\theta$ ).
■ Computing $H_{u}^{(1)}, \ldots, H_{u}^{(R)}$ in parallel, we obtain

$$
\bar{H}_{u}=\frac{1}{R} \sum_{r=1}^{R} H_{u}^{(r)}
$$

along with a CLT-based error estimate.

- By contrast, for existing smoothing techniques, parallelism is not trivial, nor is the construction of error estimates.


## Proposed estimator

- Use Rhee \& Glynn (2014) trick to turn a Conditional Particle Filter kernel into an unbiased estimator of smoothing functionals.
- Coupled resampling schemes are instrumental in this construction.
- Instead of two particle systems given $\theta$ and $\tilde{\theta}$, we consider two particle systems with same $\theta$ but different "reference trajectories".


## Trajectories from particle filters

Upon running a particle filter, we get trajectories $x_{0: T}^{1: N}$ with weights $w_{T}^{1: N}$.


Figure: Hidden auto-regressive model, $T=100$ observations, $N=128$.

## Conditional particle filter

Input: a trajectory $\check{x}_{0: T}$.
At step $t=0$,
1 Sample $x_{0}^{k} \sim \mu_{\theta}\left(d x_{0}\right)$, for all $k \in 1: N-1$, set $x_{0}^{N}=\check{x}_{0}$.
2 Set $w_{0}^{k}=N^{-1}$, for all $k \in 1: N$.
At step $t \geq 1$,
$\boldsymbol{1}$ Sample ancestors $a_{t}^{1: N-1} \sim r\left(d a^{1: N-1} \mid w_{t-1}^{1: N}\right)$, set $a_{t}^{N}=N$.
2 Sample $x_{t}^{k} \sim f_{\theta}\left(d x_{t} \mid x_{t-1}^{a_{t}^{k}}\right)$, for all $k \in 1: N-1$, set

$$
x_{t}^{N}=\check{x}_{t} .
$$

3 Compute $w_{t}^{k}=g_{\theta}\left(y_{t} \mid x_{t}^{k}\right)$, for all $k \in 1: N$.
Output: sample a trajectory, $x_{0: T}^{k}$ with probability $w_{T}^{k}$.

## Conditional particle filter



Figure: $M=100$ paths, for the hidden auto-regressive model, $T=100$ observations, $N=128$.

## Conditional particle Filter



Figure: $M=100$ samples for $x_{0}$, for the hidden auto-regressive model, $T=100$ observations, $N=128$.

## Conditional particle Filter



Figure: $M=100$ samples for $x_{100}$, for the hidden auto-regressive model, $T=100$ observations, $N=128$.

## Unbiased estimators

Von Neumann \& Ulam ( $\sim 1950$ ), Kuti ( $\sim 1980$ ), Rychlik ( $\sim 1990$ ), McLeish ( $\sim 2010$ ), Rhee \& Glynn (2012, 2013, 2014).
Introduce

- a sequence of random variables $\left(H^{(n)}\right)$ with

$$
\mathbb{E}\left[H^{(n)}\right] \underset{n \rightarrow \infty}{ } \int h\left(x_{0: T}\right) p\left(d x_{0: T} \mid y_{1: T}, \theta\right)
$$

e.g. $H^{(n)}=h\left(X^{(n)}\right)$, with $\left(X^{(n)}\right)$ generated by CPF,

- a sequence $\left(\Delta^{(n)}\right)$ such that

$$
\begin{aligned}
& \mathbb{E}\left[\Delta^{(n)}\right]=\mathbb{E}\left[H^{(n)}-H^{(n-1)}\right] \\
& \mathbb{E}\left[\sum_{n=0}^{\infty}\left|\Delta^{(n)}\right|\right]<\infty
\end{aligned}
$$

with $H^{(-1)}=0$ by convention.

## Unbiased estimators

- Then

$$
\begin{aligned}
\mathbb{E} \sum_{n=0}^{\infty} \Delta^{(n)} & =\sum_{n=0}^{\infty} \mathbb{E}\left[\Delta^{(n)}\right]=\sum_{n=0}^{\infty} \mathbb{E}\left[H^{(n)}-H^{(n-1)}\right] \\
& =\lim _{n \rightarrow \infty} \mathbb{E}\left[H^{(n)}\right]=\int h\left(x_{0: T}\right) p\left(d x_{0: T} \mid y_{1: T}, \theta\right)
\end{aligned}
$$

Thus, consider

$$
H_{u}=\sum_{n=0}^{K} \frac{\Delta^{(n)}}{\mathbb{P}(K \geq n)}
$$

where $K$ is an integer-valued random variable. Then

$$
\mathbb{E}\left[H_{u}\right]=\mathbb{E}\left[\sum_{n=0}^{\infty} \frac{\Delta^{(n)} \mathbb{1}(K \geq n)}{\mathbb{P}(K \geq n)}\right]=\int h\left(x_{0: T}\right) p\left(d x_{0: T} \mid y_{1: T}, \theta\right)
$$

## Unbiased estimators

Idea from Rhee \& Glynn, 2014. Write

$$
X^{(n)}=\varphi_{n}\left(X^{(n-1)}\right)=\varphi_{n} \circ \varphi_{n-1} \circ \ldots \circ \varphi_{1}\left(X^{(0)}\right)
$$

Introduce
$\tilde{X}^{(0)} \triangleq X^{(0)}, \quad \tilde{X}^{(1)}=\varphi_{2}\left(\tilde{X}^{(0)}\right), \quad \ldots, \quad \tilde{X}^{(n)}=\varphi_{n+1} \circ \ldots \circ \varphi_{2}\left(X^{(0)}\right)$.
Then $\Delta^{(n)}=h\left(X^{(n)}\right)-h\left(\tilde{X}^{(n-1)}\right)$ is such that

$$
\mathbb{E}\left[\Delta^{(n)}\right]=\mathbb{E}\left[H^{(n)}-H^{(n-1)}\right]
$$

and we might have

$$
\mathbb{E}\left[\sum_{n=0}^{\infty}\left|\Delta^{(n)}\right|\right]<\infty
$$

## Unbiased estimators based on CPF chains

- Start from $X^{(0)}$ and $\tilde{X}^{(0)}$ generated by two particle filters.
- Apply one step of CPF kernel to $X^{(0)}$, to get $X^{(1)}$.
- For $n \geq 2$, apply the CPF kernel to both $X^{(n-1)}$ and $\tilde{X}^{(n-2)}$, with the same random numbers, to get $X^{(n)}$ and $\tilde{X}^{(n-1)}$.
- We can see each step as a joint CPF acting on pairs of trajectories, and use coupled resampling ideas.
- Can we expect $\Delta^{(n)}=h\left(X^{(n)}\right)-h\left(\tilde{X}^{(n-1)}\right)$ to decrease to zero in average?


## Norm of $\Delta^{(n)}$ with independent resampling



Hidden auto-regressive model, $T=20$ observations, $N=32$.

## Norm of $\Delta^{(n)}$ with independent resampling


$T=100$ observations, $N=128$.

## Norm of $\Delta^{(n)}$ with index-matching resampling


$T=20$ observations, $N=32$.

## Norm of $\Delta^{(n)}$ with index-matching resampling


$T=100$ observations, $N=128$.

## Coupled conditional particle Filter

- We consider coupled conditional particle filters, acting on pairs of trajectories:

$$
\left(X^{(n)}, \tilde{X}^{(n-1)}\right)=\bar{\varphi}_{n}\left(X^{(n-1)}, \tilde{X}^{(n-2)}\right)
$$

- A coupled CPF kernel uses common random numbers for both systems, and a coupled resampling scheme.
- We focus on index-matching resampling.
- We see that after a number of coupled CPF steps, $X^{(n)}=\tilde{X}^{(n)}$ exactly, and thus $\Delta^{(n)}=0$.
■ We can thus stop early in the computation of

$$
H_{u}=\sum_{n=0}^{K} \frac{\Delta^{(n)}}{\mathbb{P}(K \geq n)}
$$

## Proposed estimator

Case $h=I d$ : we estimate the smoothing means.

- Sample an integer-valued random variable $K$.
- Sample $\varphi_{1}$, draw $X^{(0)}$ and set $X^{(1)}=\varphi_{1}\left(X^{(0)}\right)$.
- Compute $\Delta^{(0)}=X^{(0)}$, set $H_{u} \leftarrow \Delta^{(0)}$.
- Sample $\tilde{X}^{(0)} \triangleq X^{(0)}$, compute $\Delta^{(1)}=X^{(1)}-\tilde{X}^{(0)}$.
- Set $H_{u} \leftarrow H_{u}+\Delta^{(1)} / \mathbb{P}(K \geq 1)$.
- For $n=2, \ldots, K$,
- Sample $\bar{\varphi}_{n}$, set $\left(X^{(n)}, \tilde{X}^{(n-1)}\right)=\bar{\varphi}_{n}\left(X^{(n-1)}, \tilde{X}^{(n-2)}\right)$.
- Compute $\Delta^{(n)}=X^{(n)}-\tilde{X}^{(n-1)}$.
- Stop if $\Delta^{(n)}=0$.
- Set $H_{u} \leftarrow H_{u}+\Delta^{(n)} / \mathbb{P}(K \geq n)$.
- Return $H_{u}$.


## Example: Phytoplankton-Zooplankton



Figure: Phytoplankton-Zooplankton model, $T=365, N=1,024$, $R=1,000$ estimators, with a Geometric truncation with mean 100 .

## Example: Phytoplankton-Zooplankton



Figure: Smoothing means of $P, T=365, N=1,024, R=1,000$ estimators, with a Geometric truncation with mean 100.

The bars represent $\pm 2 \sigma$ around the estimated means. The blue line is obtained from a long CPF run.

## Example: Phytoplankton-Zooplankton



Figure: Smoothing means of $Z, T=365, N=1,024, R=1,000$ estimators, with a Geometric truncation with mean 100.

The bars represent $\pm 2 \sigma$ around the estimated means. The blue line is obtained from a long CPF run.

## Example: Phytoplankton-Zooplankton



Figure: Smoothing means of $Z$, first $65 / 365$ time steps, $N=1,024$, $R=1,000$ estimators, with a Geometric truncation with mean 100 .

## Example: Phytoplankton-Zooplankton



Figure: Smoothing means of $Z$, last $65 / 365$ time steps, $N=1,024$, $R=1,000$ estimators, with a Geometric truncation with mean 100 .

## Example: Phytoplankton-Zooplankton



Figure: Trace of relative variance of the smoothing mean estimator, $T=365, N=1,024, R=1,000$ estimators, with a Geometric truncation with mean 100 .

## Discussion

- Coupled resampling schemes can be used to improve a variety of particle-based algorithms.
- New estimator of smoothing functionals, easy to parallelize and with error estimates.
- Benefits greatly from ancestor sampling:



## The End

## Thank you for listening!

Soon on arXiv...
PJ, Fredrik Lindsten, Thomas Schön, Coupling Particle Filters.

- Pitt \& Malik, 2011, Particle filters for continuous likelihood evaluation and maximisation, J. of Econometrics.
- Rhee \& Glynn, 2014, Exact estimation for markov chain equilibrium expectations, arXiv.

■ Deligiannidis, Doucet, Pitt \& Kohn, 2015, The correlated pseudo-marginal method, arXiv.

