

Conditional particle filters with bridge backward sampling

Matti Vihola (Jyväskylä)

MCM @ Paris

28 June 2023



Joint work with: Santeri Karppinen (Jyväskylä)
Sumeetpal S. Singh (Wollongong)
Nicolas Chopin (ENSAE)
Tomás Soto (LUT)





Introduction: Smoothing & conditional particle filters

Problematic case: Time-discretised path integral model

Solution: CPF with bridge backward sampling

Introduction: Smoothing & conditional particle filters


The smoothing problem



Hidden Markov model (a.k.a. general state space model):

- Latent Markov chain X_k with initial law $m_1(\cdot)$ and transitions $m_k(x_{k-1}, \cdot)$
- Conditionally independent observations y_k with densities $g_k(\cdot | X_k)$

We aim at inferring $\pi_T = \mathcal{L}(X_{1:T} | Y_{1:T} = y_{1:T})$ (smoothing).




The smoothing problem



Hidden Markov model (a.k.a. general state space model):

- Latent Markov chain X_k with initial law $m_1(\cdot)$ and transitions $m_k(x_{k-1}, \cdot)$
- Conditionally independent observations y_k with densities $g_k(\cdot | X_k)$

We aim at inferring $\pi_T = \mathcal{L}(X_{1:T} | Y_{1:T} = y_{1:T})$ (smoothing).



The Feynman–Kac notation/generalisation:

- Proposal distributions: $M_1(\cdot)$, $M_k(x_{k-1}, \cdot)$ e.g. $M_k \equiv m_k$
- Potential (or weight) functions: $G_k(\cdot) \geq 0$ e.g. $G_k(\cdot) \equiv g_k(y_k | \cdot)$

\rightsquigarrow different 'proposal distributions', 'lookaheads' and more general models...

Conditional particle filter (CPF)




Markov transition $x_{1:T}^* \rightarrow (X_1^{B_1}, \dots, X_T^{B_T})$ leaving π_T invariant $\forall N \geq 2!$ ¹

1. Reference path $x_{1:T}^*$
2. Sample remaining $N - 1$ particles sequentially from M_1, M_2, \dots, M_T ;
conditional resample proportional to $G_k(\cdot)$
3. Pick particle B_T at last time T ;
Ancestor trace B_{T-1}, \dots, B_1 and output $(X_1^{B_1}, \dots, X_T^{B_T})$


¹Andrieu, Doucet & Holenstein (*J. Roy. Statist. Soc. Ser. B.*, 2010)

Iterated CPF on noisy AR(1)




- Reference $x_{1:T}^*$, Output $X_{1:T}^{B_{1:T}}$

Iterated CPF on noisy AR(1)




- Reference $x_{1:T}^*$, Output $X_{1:T}^{B_{1:T}}$

Iterated CPF on noisy AR(1)



- Reference $x_{1:T}^*$, Output $X_{1:T}^{B_{1:T}}$

Iterated CPF on noisy AR(1)



- Reference $x_{1:T}^*$, Output $X_{1:T}^{B_{1:T}}$

Conditional particle filter with backward sampling (CPF-BS)



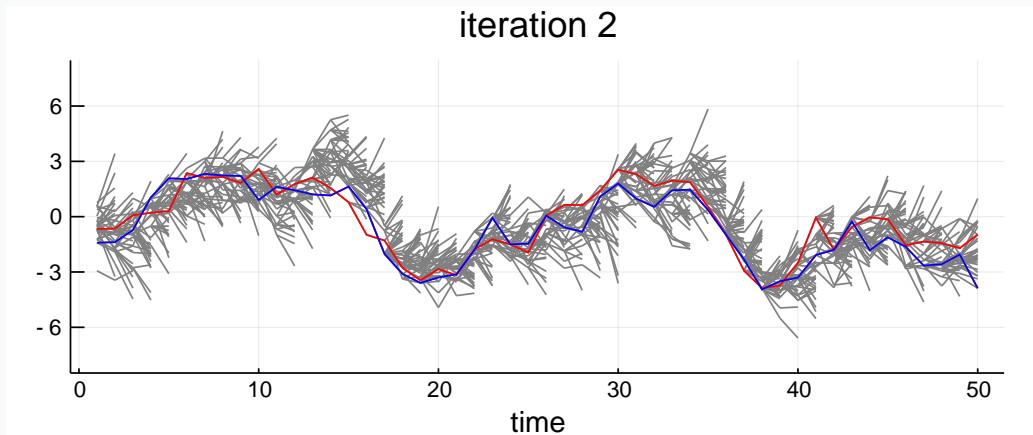
Efficient MCMC smoothing: π_T -invariant transition $x_{1:T}^* \rightarrow (X_1^{B_1}, \dots, X_T^{B_T})^2$

1. Reference path $x_{1:T}^*$
2. Sample remaining $N - 1$ particles sequentially from M_1, M_2, \dots, M_T ;
conditional resample proportional to $G_k(\cdot)$
3. Pick particle B_T at last time T ;
Backward sample B_{T-1}, \dots, B_1 and output $(X_1^{B_1}, \dots, X_T^{B_T})$

²Whiteley (J. Roy. Statist. Soc. Ser. B., 2010);


Algorithmic variant: Lindsten, Jordan & Schön (J. Mach. Learn. Res., 2014) ancestor sampling CPF

Iterated CPF-BS on noisy AR(1)




- Reference $x_{1:T}^*$, Output $X_{1:T}^{B_{1:T}}$

Iterated CPF-BS on noisy AR(1)




- Reference $x_{1:T}^*$, Output $X_{1:T}^{B_{1:T}}$

Iterated CPF-BS on noisy AR(1)



- Reference $x_{1:T}^*$, Output $X_{1:T}^{B_{1:T}}$

Iterated CPF-BS on noisy AR(1)



- Reference $x_{1:T}^*$, Output $X_{1:T}^{B_{1:T}}$

Why to use CPF-BS?



- Hard high-dimensional inference problem: $X_k \in \mathbb{R}^d \implies X_{1:T} \in \mathbb{R}^{dT}$.
- Impressive empirical performance even with very large T
- Theoretical guarantees under general conditions³:
 - Number of particles N sufficiently large (but **constant in T !**)
 - Mixing time at most **linear in T**
- CPF-BS can be plugged into **particle Gibbs**, where the model has unknown parameters θ :
 - Update **$X_{1:T} \mid \theta$ by CPF-BS** (using $(M_{1:T}^\theta, G_{1:T}^\theta)$)
 - Update $\theta \mid X_{1:T}$ by your favourite MCMC...

³Lee, Singh & V (*Ann. Statist.*, 2020)

Problematic case: Time-discretised path integral model

The path integral model



- A continuous-time analogue of the HMM:

$$\Pi_\tau(A) = \frac{1}{\mathcal{Z}} \mathbb{E} \left[\overbrace{1(Z_{[0,\tau]} \in A)}^{\text{Markov process law}} \overbrace{\exp \left(- \int_0^\tau V(Z_u) du \right)}^{\text{Weight}} \right],$$

where $(Z_t)_{t \geq 0} \sim \mathbb{M}$ is a Markov process (typically diffusion, from SDE) and $\mathcal{Z} > 0$ is a normalising constant.

- Potential function $V \geq 0$ modulates the law \mathbb{M} — the ‘environment’
 - $V(x)$ has small ‘penalty’ value \rightsquigarrow preferred state x
 - $V(x)$ has high ‘penalty’ value \rightsquigarrow unlikely state x

Time-discretised path integral model



- Time-discretisation $0, \Delta, 2\Delta, \dots, \Delta T = \tau$ and Riemann sum approximation:

$$\pi_T(x_{1:T}) = \frac{1}{\mathcal{Z}} \left(M_1(x_1) \prod_{k=2}^T M_k^\Delta(x_k | x_{k-1}) \right) \left(\prod_{k=1}^{T-1} \overbrace{e^{-\Delta V(x_k)}}^{G_k^\Delta(x_k)} \right),$$

where

- $X_k = Z_{\Delta(k-1)}$
- M_k^Δ are the (approximate) transition probabilities from $Z_{\Delta(k-1)}$ to $Z_{\Delta k}$
- When $\Delta \rightarrow 0$:
 - $M_k^\Delta(x_k | x_{k-1}) \rightarrow \delta_{x_{k-1}}(x_k)$
 - $G_k^\Delta(x) \approx 1 - \Delta V(x) \rightarrow 1$
 - $T = \tau/\Delta \rightarrow \infty$
- Problem should not get 'harder' as $\Delta \rightarrow 0$, but 'standard' CPF-BS degenerates...

CPF resampling for $\Delta \rightarrow 0$



- Original CPF & CPF-BS used conditional **multinomial resampling**⁴
 - Degenerates ✗
- Two other conditional resamplings for CPF⁵:
 - **residual** degenerates ✗
 - **systematic** seems stable ✓ (empirically...)
- Can we say something more about resamplings in the limit?

⁴Andrieu, Doucet & Holenstein (*J. Roy. Statist. Soc. Ser. B.*, 2010)

⁵Chopin & Singh (*Bernoulli*, 2015)

⁶Chopin, Singh, Soto & V, *Ann. Statist.* (2022)

CPF resampling for $\Delta \rightarrow 0$



- Original CPF & CPF-BS used conditional **multinomial resampling**⁴
 - Degenerates ✗
- Two other conditional resamplings for CPF⁵:
 - **residual** degenerates ✗
 - **systematic** seems stable ✓ (empirically...)
- Can we say something more about resamplings in the limit?
- At least for ~~conditional~~ particle filters⁶:
 - 'Killing' ✓
 - 'Stratified' (with mean partition) ✓
 - 'Systematic' (w.m.p) ✓
 - 'Srinivasan Sampling process' (w.m.p.) (SSP) ✓

⁴Andrieu, Doucet & Holenstein (*J. Roy. Statist. Soc. Ser. B.*, 2010)

⁵Chopin & Singh (*Bernoulli*, 2015)

⁶Chopin, Singh, Soto & V, *Ann. Statist.* (2022)

Interlude (for experts): Particle filter limit as $\Delta \rightarrow 0$



Assumption: Instantaneous resampling rate

For all $v^{1:N} \geq 0$ and all $a^{1:N} \in \{1:N\}^N$, $a^{1:N} \neq 1:N$ the limit:

$$\iota(a^{1:N}, v^{1:N}) = \lim_{\Delta \rightarrow 0+} \frac{1}{\Delta} r(a^{1:N} \mid \exp(-\Delta v^1), \dots, \exp(-\Delta v^N))$$

exists and the term inside the limit is uniformly bounded for bounded potentials $v^{1:N} \leq v^* < \infty$ and $\Delta \in (0, 1)$.

Theorem (Chopin, Singh, Soto, V, 2022)

Under Assumption & technical conditions, the discrete-time particle filter converges to Markov process with generator

$$\mathcal{L}f(x^{1:N}) = \mathcal{L}^{\mathbb{M}}f(x^{1:N}) + \sum_{a^{1:N} \neq 1:N} \iota(a^{1:N}, (V(x^1), \dots, V(x^N))) (f(x^{a^{1:N}}) - f(x^{1:N}))$$

Interlude (cont): Limiting overall death rates



- Simple expressions for the overall death rate:

$$\iota^*(v^{1:N}) = \sum_{a^{1:N} \neq 1:N} \iota(a^{1:N}, v^{1:N}),$$

which captures $\mathbb{P}(A_k^{1:N} \neq 1:N) \approx \Delta \iota^*(V(X_k^1), \dots, V(X_k^N))$

Proposition

$$\iota_{\text{killing}}^*(v^{1:N}) = \frac{N-1}{N} \sum_{i=1}^N (v^i - v_{\min}) = (N-1)(\bar{v} - v_{\min}) \quad v_{\min} = \min_j v^j$$

$$\iota_{\text{stratified}}^*(v^{1:N}) = \sum_{j=1}^N j(\bar{v} - v^{\varpi(j)}) \quad \bar{v} = \frac{1}{N} \sum_{i=1}^N v^i$$

$$\iota_{\text{systematic}}^*(v^{1:N}) = \frac{1}{2} \sum_{i=1}^N |\bar{v} - v^i| = \iota_{\text{ssp}}^*(v^{1:N})$$

Interlude (cont.): Ordering the overall death rates



Theorem (Chopin, Singh, Soto, V, 2022)

The overall resampling intensities of killing, stratified and systematic/SSP resampling satisfy

$$\iota_{\text{killing}}^*(v^{1:N}) \geq \iota_{\text{systematic}}^*(v^{1:N}), \quad \text{and} \quad \iota_{\text{stratified}}^*(v^{1:N}) \geq \iota_{\text{systematic}}^*(v^{1:N}),$$

for all potential values $v^{1:N}$. However, ι_{killing}^ and $\iota_{\text{stratified}}^*$ do not satisfy such order in general.*

↪ Systematic/SSP with mean partition preferable

- Karppinen, Singh & V (to appear):
 - Conditional killing resampling
 - Conditional systematic with mean partition ✓✓

Backward sampling with stiff dynamics M_k




As $\Delta \rightarrow 0$:

- Continuous paths $\implies M_k^\Delta(\cdot | x) \rightarrow \delta_x(\cdot)$.
- Backward sampling step:

$$\mathbb{P}(B_k = i | B_{k+1}) \propto G_k^\Delta(X_k^i) \overbrace{M_{k+1}^\Delta(X_{k+1}^{B_{k+1}} | X_k^i)}^{\approx 0 \text{ for } i \neq A_{k+1}^{B_{k+1}}}$$

$\implies B_k = A_{k+1}^{B_{k+1}}$ with high probability.

- \therefore CPF-BS degenerates to CPF with ancestor tracing;
the benefits of backward sampling are lost...




**Solution: CPF with bridge backward
sampling**

CPF with bridge backward sampling (CPF-BBS)




Generalisation of backward sampling to blocks of length > 1

- We use auxiliary ‘bridge CPF’ steps
- Requires tractable M_k (can calculate conditionals, such as linear-Gaussian)
- Initialisation: run CPF with reference $x_{1:T}^*$ and choose a particle $X_T^{(B_T)}$




Bridge CPF: block from $\ell = 10$ to $u = 20$



- 1 Trace back from $X_u^{B_u}$
→ 'block reference' $\hat{X}_{\ell:u}$

Bridge CPF: block from $\ell = 10$ to $u = 20$




1 Trace back from $X_u^{B_u}$
→ 'block reference' $\hat{X}_{\ell:u}$

2 Simulate particles $\tilde{X}_{\ell+1:u-1}^{1:N}$ with auxiliary CPF:

- reference $\hat{X}_{\ell:u}$
- conditional transitions $M_{k|k-1,u}(\cdot \mid x_{k-1}, \hat{X}_u)$
- modified potentials $G_k(x_{k-1}, x_k) M_{u|\ell}(\hat{X}_u \mid x_\ell)^{\frac{1}{u-\ell}},$

Bridge CPF: block from $\ell = 10$ to $u = 20$



1 Trace back from $X_u^{B_u}$
 \rightarrow 'block reference' $\hat{X}_{\ell:u}$

2 Simulate particles $\tilde{X}_{\ell+1:u-1}^{1:N}$ with auxiliary CPF:

- reference $\hat{X}_{\ell:u}$
- conditional transitions $M_{k|k-1,u}(\cdot \mid x_{k-1}, \hat{X}_u)$
- modified potentials $G_k(x_{k-1}, x_k) M_{u|\ell}(\hat{X}_u \mid x_\ell)^{\frac{1}{u-\ell}},$

3 Output trajectory $\tilde{X}_{\ell:u}^{\tilde{B}_{\ell:u}}$

Choosing the blocks



- CPF-BBS efficient if blocks chosen appropriately
- Natural rule: maximise ‘probability of local update’ (PLU):

$$\mathbb{P}(\tilde{X}_{\ell}^{\tilde{B}_{\ell}} \neq \hat{X}_{\ell})$$

Choosing the blocks



- CPF-BBS efficient if blocks chosen appropriately
- Natural rule: maximise ‘probability of local update’ (PLU):

$$\mathbb{P}(\tilde{X}_\ell^{\tilde{B}_\ell} \neq \hat{X}_\ell) \approx \text{PLU}_M(\ell, u) \text{PLU}_G(\ell, u) \left(1 - \frac{1}{N}\right)^{-1}$$


- Based on two approximate estimators for PLU:

$$\text{PLU}_M(\ell, u) = 1 - \frac{M_{u|\ell}(x_u^* | x_\ell^*)}{\sum_{j=1}^N M_{u|\ell}(x_u^* | x_\ell^{(j)})} \text{ (short blocks)}$$

$$\text{PLU}_G(\ell, u) = \left(1 - \frac{1}{N}\right) \prod_{k=\ell}^{u-1} \left(1 - \frac{p_k N}{(N-1)^2}\right) \text{ (long blocks),}$$

where $p_k = \frac{1}{2} \sum_{i=1}^N |W_k^i - \frac{1}{N}|$ systematic ‘resampling rate’

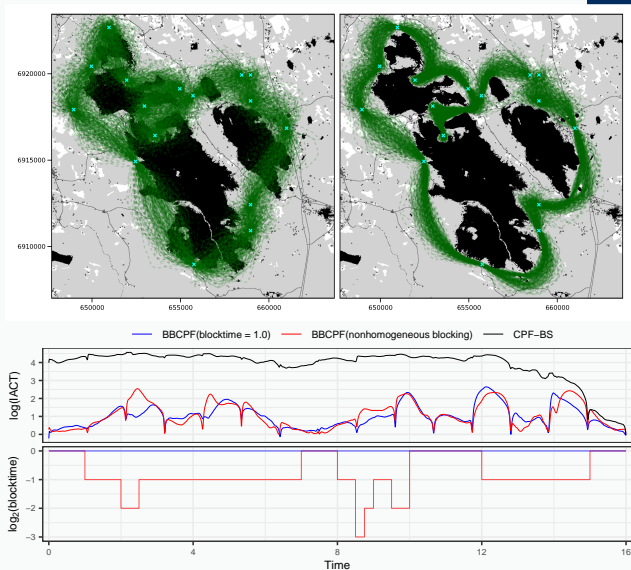
- One trial run, can evaluate with multiple candidate $\ell, u...$



Example: Interpolation with terrain preferences



- Ornstein–Uhlenbeck velocity conditioned on linear Gaussian **observations** \times (left)
- Path integral model: avoid dark water bodies (right)^a
- Hand-tuned **constant size blocking**
- PLU-based automatic **'optimised' blocking**



^aAnimation:  @MattiVihola

Concluding remarks



CPF can work with weakly informative G_k :

- Use conditional **systematic resampling with mean partition** with CPF

CPF-BS can be extended to work with stiff M_k (if tractable):

- Use **bridge backward sampling** CPF
- Good heuristics to choose **appropriate blocking**

Discussion:

- Weak potentials and stiff dynamics arise with HMMs, too!
- The method requires tractable $M_{1:T}$, which can be from a Gaussian approximation...
- Blocked particle Gibbs (Singh, Lindsten & Moulines, *Biometrika*, 2015) can also be used, but requires block size & *block overlap* calibration...



- **C. Andrieu, A. Doucet and R. Holenstein.**
Particle Markov chain Monte Carlo methods.
J. R. Stat. Soc. Ser. B. Stat. Methodol., 2010.
- **N. Whiteley.**
Discussion on “PMCMC methods”.
J. R. Stat. Soc. Ser. B. Stat. Methodol., 2010.
- **F. Lindsten, M. I. Jordan and T. B. Schön.**
Particle Gibbs with ancestor sampling.
J. Mach. Learn. Res., 2014.
- **N. Chopin and S. S. Singh**
On particle Gibbs sampling.
Bernoulli, 2015.
- **A. Lee, S. S. Singh and M. Vihola.**
Coupled conditional backward sampling particle filter.
Ann. Statist., 2020.
- **S. S. Singh, F. Lindsten and E. Moulines.**
Blocking strategies and stability of particle Gibbs samplers.
Biometrika, 2017.
- **N. Chopin, S. S. Singh, T. Soto and M. Vihola**
On resampling schemes for particle filters with weakly informative observations.
Ann. Statist., 2022.
- **S. Karppinen, S. S. Singh and M. Vihola.**
Conditional particle filters with bridge backward sampling
J. Comput. Graph. Statist., to appear. (arXiv:2205.13898)

Julia package <https://github.com/skarppinen/Resamplings.jl>
& codes <https://github.com/skarppinen/cpf-bbs>

Mean partition





Definition

Suppose that $g^{1:N} \in \mathbb{R}^N$. A permutation ϖ of $1:N$ is a mean partition (order) for $g^{1:N}$ if the re-indexed $g_{\varpi}^i = g^{\varpi(i)}$ satisfy

$$g_{\varpi}^1, \dots, g_{\varpi}^m \leq \bar{g} \quad \text{and} \quad g_{\varpi}^{m+1}, \dots, g_{\varpi}^N > \bar{g},$$

where the pivot \bar{g} is the mean $\bar{g} = \frac{1}{N} \sum_{i=1}^N g^i$.




Stratified/Systematic/SSP with mean partition



Slightly modified resampling schemes:

1. Mean partition order the $g^{1:N} \rightarrow g_{\varpi}^{1:N}$
 - costs $O(N)$, like the first stage of Quicksort...
2. Apply the Strat./Syst./SSP resampling:
 $\tilde{A}^{1:N} \sim r(\cdot \mid g_{\varpi}^{1:N})$
3. Map the $\tilde{A}^{1:N}$ back to 'original indexing'
 $A^{1:N}$:


$$A^{\varpi(i)} = \varpi(\tilde{A}^i)$$



Weak potentials

The potentials G_k are ‘weakly informative’ if the weights $G_k(X_k^1), \dots, G_k(X_k^N)$ are ‘nearly constant’


\leadsto Normalised weights $W_k^i = \frac{G_k(X_k^i)}{\sum_{j=1}^N G_k(X_k^j)} \approx \frac{1}{N}$



Weak potentials

The potentials G_k are ‘weakly informative’ if the weights $G_k(X_k^1), \dots, G_k(X_k^N)$ are ‘nearly constant’

\leadsto Normalised weights $W_k^i = \frac{G_k(X_k^i)}{\sum_{j=1}^N G_k(X_k^j)} \approx \frac{1}{N}$



Examples of weakly informative scenarios:


1. If $G_k(x_k) = g_k(y_k | x_k)$ in the HMM context, Y_k have high variability wrt. M_k



Weak potentials

The potentials G_k are ‘weakly informative’ if the weights $G_k(X_k^1), \dots, G_k(X_k^N)$ are ‘nearly constant’

\leadsto Normalised weights $W_k^i = \frac{G_k(X_k^i)}{\sum_{j=1}^N G_k(X_k^j)} \approx \frac{1}{N}$



Examples of weakly informative scenarios:


1. If $G_k(x_k) = g_k(y_k | x_k)$ in the HMM context, Y_k have high variability wrt. M_k
2. When M_k form a good approximation of the smoothing distribution (e.g. Laplace approximation with linear-Gaussian m_k and nonlinear g_k ...)



Weak potentials

The potentials G_k are ‘weakly informative’ if the weights $G_k(X_k^1), \dots, G_k(X_k^N)$ are ‘nearly constant’

\leadsto Normalised weights $W_k^i = \frac{G_k(X_k^i)}{\sum_{j=1}^N G_k(X_k^j)} \approx \frac{1}{N}$



Examples of weakly informative scenarios:

1. If $G_k(x_k) = g_k(y_k | x_k)$ in the HMM context, Y_k have high variability wrt. M_k
2. When M_k form a good approximation of the smoothing distribution (e.g. Laplace approximation with linear-Gaussian m_k and nonlinear $g_k \dots$)
3. When M_k, G_k come from a continuous-time path integral model

BRIDGECPF($\underline{x}_\ell, b_{\ell:u-1}^*, x_{\ell+1:u}^*$)

- 1: $W_\ell^{(1:N)} \leftarrow M_{u|\ell}(x_u^* \mid \underline{x}_\ell)^{\frac{1}{u-\ell}}; \tilde{\mathbf{X}}_\ell \leftarrow \underline{x}_\ell$
- 2: **for** $v = \ell + 1 : u - 1$ **do**
- 3: $\tilde{A}_{v-1}^{(1:N)} \leftarrow r^{(b_{v-1}^*, b_v^*)}(\cdot \mid G_{v-1}(\tilde{\mathbf{X}}_{v-1}^{(1:N)})W_{v-1}^{(1:N)})$
- 4: Draw $\tilde{X}_v^{(i)} \sim \bar{M}_v(\cdot \mid \tilde{X}_{v-1}^{(\tilde{A}_{v-1}^{(i)})}, \tilde{X}_u^*)$ for $i \neq b_v^*$ and set $\tilde{X}_v^{(b_v^*)} = x_v^*$
- 5: Set $\tilde{\mathbf{X}}_v^{(i)} \leftarrow (X_{v-1}^{(\tilde{A}_{v-1}^{(i)})}, \tilde{X}_v^{(i)})$ for $i \in \{1:N\}$.
- 6: $W_v^{(1:N)} \leftarrow W_{v-1}^{(\tilde{A}_{v-1}^{(1:N)})}$
- 7: **end for**
- 8: Draw $\tilde{B}_{u-1} \sim \text{Categorical}(\tilde{\omega}_{u-1}^{(1:N)})$ where $\tilde{\omega}_{u-1}^{(j)} =$
 $G_{u-1}(\tilde{\mathbf{X}}_{u-1}^{(j)})G_u(\tilde{X}_{u-1}^{(j)}, x_u^*)W_{u-1}^{(j)}$
- 9: $\tilde{B}_{\ell:u-2} \leftarrow \text{ANCESTORTRACE}(\tilde{A}_{\ell:u-2}, \tilde{B}_{u-1})$
- 10: **output** $((x_\ell^{(\tilde{B}_\ell)}, \tilde{X}_{\ell+1:u-1}^{(\tilde{B}_{\ell+1:u-1})}), \tilde{B}_{\ell:u-1})$

Conditional particle filter (CPF)



Markov transition $x_{1:T}^* \rightarrow (X_1^{B_1}, \dots, X_T^{B_T})$ leaving π_T invariant $\forall N \geq 2!$

CONDITIONALPARTICLEFILTER($M_{1:T}, G_{1:T}, r, N, x_{1:T}^*, b_{1:T}^*$)

```
1:  $X_1^{b_1^*} \leftarrow x_1^*$  and  $X_1^i \sim M_1(\cdot)$  for  $i \neq b_1$ 
2: for  $k = 2, \dots, T$  do
3:    $A_k^{1:N} \sim r^{(b_{k-1}^*, b_k^*)}(G_{k-1}(X_{k-1}^1), \dots, G_{k-1}(X_{k-1}^N))$  Conditional resample
4:    $X_k^{b_k^*} \leftarrow x_k^*$  and  $X_k^i \sim M_k(\cdot \mid X_{k-1}^{A_k^i})$  for  $i \neq b_k^*$ 
5: end for
6: Draw  $B_T \sim \text{Categorical}(G_T(X_T^1), \dots, G_T(X_T^N))$  Pick particle at end time  $T$ 
7: for  $k = T - 1, \dots, 1$  do
8:    $B_k \leftarrow A_{k+1}^{B_{k+1}}$  Ancestor trace
9: end for
10: output  $(X_1^{B_1}, \dots, X_T^{B_T})$ 
```

⁶Andrieu, Doucet & Holenstein (J. Roy. Statist. Soc. Ser. B., 2010)

Conditional particle filter with backward sampling (CPF-BS)



Efficient MCMC smoothing: π_T -invariant transition $x_{1:T}^* \rightarrow (X_1^{B_1}, \dots, X_T^{B_T})^7$

CPF-BS($M_{1:T}, G_{1:T}, r, N, x_{1:T}^*, b_{1:T}^*$)

- 1: $X_1^{b_1^*} \leftarrow x_1^*$ and $X_1^i \sim M_1(\cdot)$ for $i \neq b_1$
- 2: **for** $k = 2, \dots, T$ **do**
- 3: $A_k^{1:N} \sim r^{(b_{k-1}^*, b_k^*)}(G_{k-1}(X_{k-1}^1), \dots, G_{k-1}(X_{k-1}^N))$ Conditional resample
- 4: $X_k^{b_k^*} \leftarrow x_k^*$ and $X_k^i \sim M_k(\cdot \mid X_{k-1}^{A_k^i})$ for $i \neq b_k^*$
- 5: **end for**
- 6: Draw $B_T \sim \text{Categorical}(G_T(X_T^1), \dots, G_T(X_T^N))$ Pick index at end
- 7: **for** $k = T - 1, \dots, 1$ **do**
- 8: $B_k \sim \text{Categorical}(w_k^{1:N})$, where $w_k^i = G_k(X_k^i)M_{k+1}(X_{k+1}^{(B_{k+1})} \mid X_k^{(i)})$
- 9: **end for**
- 10: **output** $(X_1^{B_1}, \dots, X_T^{B_T})$

⁷Whiteley (J. Roy. Statist. Soc. Ser. B., 2010);