

Geometric ergodicity of Rao and Teh’s algorithm for Markov jump processes and CTBNs

Błażej Miasojedow*

Institute of Applied Mathematics and Mechanics, University of Warsaw

Institute of Mathematics, Polish Academy of Sciences

e-mail: bmia@mimuw.edu.pl

and

Wojciech Niemiro

Institute of Applied Mathematics and Mechanics, University of Warsaw

e-mail: wniem@mimuw.edu.pl

Abstract: Rao and Teh (2012, 2013) introduced an efficient MCMC algorithm for sampling from the posterior distribution of a hidden Markov jump process. The algorithm is based on the idea of sampling virtual jumps. In the present paper we show that the Markov chain generated by Rao and Teh’s algorithm is geometrically ergodic. To this end we establish a geometric drift condition towards a small set. A similar result is also proved for a special version of the algorithm, used for probabilistic inference in Continuous Time Bayesian Networks.

MSC 2010 subject classifications: Primary 65C40; secondary 65C05, 60J27.

Keywords and phrases: Continuous time Markov processes, MCMC, hidden Markov models, posterior sampling, geometric ergodicity, drift condition, small set, continuous time Bayesian network.

Received November 2016.

1. Introduction

Markov jump processes (MJP) are natural extensions of Markov chains to continuous time. They are useful models of numerous phenomena examined in chemical, biological, economic and other sciences. An important class of MJP are continuous time Bayesian networks (CTBN) introduced by Schweder (1970) under the name of composable Markov chains and then reinvented by Nodelman, Shelton and Koller (2002a) under the current name. Roughly, a CTBN is a multivariate MJP in which the dependence structure between coordinates can

*The work of Błażej Miasojedow is supported by Polish National Science Center grant no. 2015/17/D/ST1/01198.

be described by a graph. Such a graphical representation allows for decomposing a large intensity matrix into smaller conditional intensity matrices.

In many applications it is necessary to consider a situation where the trajectory of a Markov jump process is not observed directly, only partial and noisy observations are available. Typically, the posterior distribution over trajectories is then analytically intractable. In the literature there exist several approaches to the above mentioned problem: based on sampling (Boys, Wilkinson and Kirkwood, 2008; El-Hay, Friedman and Kupferman, 2008; Fan and Shelton, 2008; Golightly and Wilkinson, 2011; Golightly and Wilkinson, 2014; Golightly, Henderson and Sherlock, 2015; Nodelman, Shelton and Koller, 2002b; Rao and Teh, 2013, 2012), and also based on numerical approximations. To the best of our knowledge the most general efficient method for a finite state space is that proposed by Rao and Teh (2013), and extended to a more general class of continuous time discrete systems in Rao and Teh (2012). Their algorithm is based on introducing so-called virtual jumps and a thinning procedure for Poisson processes.

Geometric ergodicity is a key property of Markov chains which implies a Central Limit Theorem for sample averages. Recently Miasojedow and Niemiro (2016) proved geometric ergodicity of Rao and Teh's algorithm in a special case of the homogeneous MJPs observed at discrete moments and when the virtual jumps are introduced by uniformization procedure. In the present paper we generalise these results to a larger class of MJPs, more general observation models and also for more general class of state dependent thinning procedures. We also establish geometric ergodicity of a Gibbs sampler for CTBNs. Overall scheme of the proofs in the present paper is roughly similar to those in Miasojedow and Niemiro (2016). However, in the more general setting the arguments are much more subtle and require some new ideas, in particular in the case of CTBNs. Also, to tackle the nonhomogeneous case, Lemma 2 uses the strong mixing condition for discrete hidden Markov models, instead of the standard ergodic theorem for homogeneous Markov chains.

Note that in practice the parameters of the hidden MJP may be unknown and have to be estimated. Then for both Bayesian and frequentist statistical inference Rao and Teh's algorithm can be applied as a part of more complex algorithms. In the Bayesian approach, Rao and Teh's algorithm can be used within a Gibbs sampler or Metropolis-Hastings algorithm which updates unknown parameters, according to some posterior distribution. In the frequentist approach, Rao and Teh's algorithm can be applied to perform E-step of Monte Carlo or stochastic approximation version of EM algorithm. Such extended versions of Rao and Teh's algorithm are not considered in our paper. We assume that the probability law of a hidden MJP is known. However we strongly believe that geometric ergodicity of the Rao and Teh's algorithm for a given parameters of hidden process will be crucial in the theoretical analysis of such complex methods.

The rest of the paper is organised as follows. In Section 2 we briefly introduce hidden Markov jump processes, next in Section 3 we recall the dependent thinning procedure and Rao and Teh's algorithm. The main result is proved in Section 4 and extensions for CTBNs are given in Section 5.

Throughout this paper we use $p(X)$ as the generic notation for a probability density of a random object X , so it may denote different functions. Set $\{n, n+1, \dots, m\}$ is denoted by $[n:m]$ (for integer $n \leq m$).

2. Hidden Markov jump processes

Consider a continuous time Markov process $\{X(t), t^{\min} \leq t \leq t^{\max}\}$ on a finite state space $\mathcal{S}$. Its probability law is defined via the initial distribution $\nu(s) = \mathbb{P}(X(t^{\min}) = s)$ and the transition intensities

$$Q(t; s, s') = \lim_{h \rightarrow 0} \frac{1}{h} \mathbb{P}(X(t+h) = s' | X(t) = s)$$

for $s, s' \in \mathcal{S}$, $s \neq s'$. Let $Q(t; s) = \sum_{s' \neq s} Q(t; s, s')$ denote the intensity of leaving state s . In general, process X can be time-inhomogeneous, that is we allow the intensities to vary in time. For definiteness, assume that X has right continuous trajectories. We say X is a Markov jump process (MJP).

Suppose that process X cannot be directly observed but we can observe some random quantity Y with probability distribution $L(Y|X)$. Let us say Y is the data and L is the likelihood. The problem is to restore the hidden trajectory of X given Y . From the Bayesian perspective, the goal is to compute/approximate/sample from the posterior

$$p(X|Y) \propto p(X)L(Y|X).$$

Function L , transition intensities Q and initial distribution ν are assumed to be known. We consider two typical forms of noisy observation. In the first part of our paper we assume that the trajectory $X([t^{\min}, t^{\max}])$ is observed independently at k deterministic time points with some random errors. Formally, we observe $Y = (Y_1, \dots, Y_k)$ where

$$L(Y|X) = \prod_{j=1}^k L_j(Y_j | X(t_j^{\text{obs}})), \quad (1)$$

for some fixed known points $t^{\min} \leq t_1^{\text{obs}} < \dots < t_k^{\text{obs}} \leq t^{\max}$. Another type of evidence is considered later in Section 5, in the context of CTBNs. In Remarks 8 and 14 we mention some alternative assumptions about the form of evidence.

The obvious standing assumption in our paper is that $L(Y|X) > 0$ happens with nonzero probability if X is given by ν and Q . It means that the hidden MJP under consideration is “possible”, i.e. the data do not contradict the probabilistic model.

3. Dependent thinning and Rao and Teh's algorithm

The so-called “dependent thinning” is a useful representation of a Markov jump process in terms of potential times of jumps and the corresponding states (Rao

and Teh, 2012). The intensities are assumed to be uniformly bounded, so the process X has a finite number of jumps in the bounded interval $[t^{\min}, t^{\max}]$. Every trajectory $X([t^{\min}, t^{\max}])$ is right continuous and piecewise constant: $X(t) = S_{i-1}$ for $T_{i-1} \leq t < T_i$, where random variables T_i are such that $t^{\min} < T_1 < \dots < T_N < t^{\max} < T_{N+1}$. By convention, $T_0 = t^{\min}$. A random sequence of states $S = (S_0, S_1, \dots, S_N)$ such that $S_i = X(T_i)$ is called a skeleton. We do not assume that $S_{i-1} \neq S_i$, and therefore the two sequences

$$\begin{pmatrix} T \\ S \end{pmatrix} = \begin{pmatrix} T_0 & T_1 & \dots & T_i & \dots & T_N \\ S_0 & S_1 & \dots & S_i & \dots & S_N \end{pmatrix}$$

represent the process X in a redundant way: many pairs (T, S) correspond to the same trajectory $X([t^{\min}, t^{\max}])$. Let $J = \{i \in [1:N] : S_{i-1} \neq S_i\} \cup \{0\}$, so that $T_J = (T_i : i \in J)$ are moments of *true* jumps and $T_{-J} = T \setminus T_J = (T_i : i \notin J)$ are *virtual* jumps. By a harmless abuse of notation, we identify increasing sequences of points in $[t^{\min}, t^{\max}]$ with finite sets. Note that the trajectory of X is uniquely defined by (T_J, S_J) . Let us write $X \equiv (T_J, S_J)$ and also use the notation $J(X) = T_J$ for the set of true jumps.

The state-dependent thinning procedure taken from Rao and Teh (2012) is the following. We choose a function $R(t; s) \geq Q(t; s)$, interpreted as intensity of an inhomogeneous Poisson process depending on state $s \in \mathcal{S}$. The first point of this Poisson process after time u , say W , has the probability density

$$f(w|u, s) = R(w; s) \exp \left[- \int_u^w R(t; s) dt \right] \quad (w > u). \quad (2)$$

Let

$$P(t; s, s') = \begin{cases} \frac{Q(t; s, s')}{R(t; s)} & \text{if } s \neq s'; \\ 1 - \frac{Q(t; s)}{R(t; s)} & \text{if } s = s'. \end{cases} \quad (3)$$

Sampling of (T, S) then proceeds as described in Algorithm 1.

Algorithm 1 Sampling by state-dependent thinning.

```

Let  $T_0 = t^{\min}$  and  $i = 0$ .
Sample  $S_0 \sim \nu(\cdot)$ .
while  $T_i < t^{\max}$  do
  Let  $i = i + 1$ .
  Sample  $T_i \sim f(\cdot | T_{i-1}, S_{i-1})$  { given by (2) }.
  Sample  $S_i \sim P(T_i; S_{i-1}, \cdot)$  { given by (3) }.
end while

```

By (2) and (3), the joint probability distribution of (T, S) is the following.

$$p(T, S) = \nu(S_0) \prod_{i=1}^N P(T_i; S_{i-1}, S_i) R(T_i; S_{i-1}) \exp \left[- \int_{T_{i-1}}^{T_i} R(t; S_{i-1}) dt \right] \\ \times \exp \left[- \int_{T_N}^{t^{\max}} R(t; S_N) dt \right]. \quad (4)$$

The last part of the above expression is equal to $\mathbb{P}(T_{N+1} > t^{\max} | T_N, S_N)$. The pair (T, S) produced by Algorithm 1 is a redundant representation of MJP X defined by ν and Q (probability distribution of X obtains if we “integrate out” virtual jumps).

Rao and Teh (2012) exploit dependent thinning to construct a special version of a Gibbs sampler which converges to the posterior $p(X|Y)$. The key facts behind their algorithm are the following. First, given the trajectory $X \equiv (T_J, S_J)$ the conditional distribution of virtual jumps T_{-J} is that of the inhomogeneous Poisson process with intensity $R(t; X(t)) - Q(t; X(t)) \geq 0$. Second, this distribution does not change if we add the likelihood. Indeed, $L(Y|X) = L(Y|T_J, S_J)$, so Y and T_{-J} are conditionally independent and thus

$$p(T_{-J} | T_J, S_J, Y) = p(T_{-J} | T_J, S_J).$$

Third, the conditional distribution $p(S|T, Y)$ is that of a hidden discrete time Markov chain and can be efficiently sampled from using the algorithm FFBS (Forward Filtering-Backward Sampling, Carter and Kohn (1994); Frühwirth-Schnatter (1994)). Indeed, from (4) and (1) it follows that

$$p(S|T, Y) \propto p(T, S) L(Y|T, S) \propto \nu(S_0) g_0(S_0) \prod_{i=1}^N P_i(S_{i-1}, S_i) g_i(S_i), \quad (5)$$

where $P_i(s, s') = P(T_i, s, s')$ is the stochastic matrix defined by (3) and

$$g_{i-1}(s) = R(T_i, s) \exp \left[- \int_{T_{i-1}}^{T_i} R(t; s) dt \right] \times \prod_{j: T_{i-1} \leq t_j^{\text{obs}} < T_i} L_j(Y_j | s) \\ \text{for } i = 1, \dots, N; \quad (6) \\ g_N(s) = \exp \left[- \int_{T_N}^{t^{\max}} R(t; s) dt \right] \times \prod_{j: T_N \leq t_j^{\text{obs}} < t^{\max}} L_j(Y_j | s).$$

(We use the convention that $\prod_{j: T_{i-1} \leq t_j^{\text{obs}} < T_i} \dots = 1$ if there is no observation between T_{i-1} and T_i .) Note that functions g_i include not only the likelihood but also a part due to the prior distribution $p(T, S)$.

Rao and Teh's algorithm generates a Markov chain $X_0, X_1, \dots, X_m, \dots$ where $X_m = X_m([t^{\min}, t^{\max}])$ is a trajectory of a MJP, convergent to $p(X|Y)$. A single

Algorithm 2 Single step of Rao and Teh's algorithm.

input: previous state $(T_J, S_J) \equiv X$ and observation Y .

(V) Sample a Poisson process V with intensity $R(t; X(t)) - Q(t; X(t))$ on $[t^{\min}, t^{\max}]$.
Let $T' = T_J \cup V$ {new set of potential times of jumps}.

(S) Draw new skeleton S' from the conditional distribution $p(\cdot | T', Y)$ by FFBS. The new allocation of virtual and true jumps is via $J' = \{i : S'_{i-1} \neq S'_i\} \cup \{0\}$ {we discard new virtual jumps $T'_{-J'}$ }.

return new state $(T'_{J'}, S'_{J'}) \equiv X'$.

step, that is the rule of transition from $X_{m-1} = X$ to $X_m = X'$ is described in Algorithm 2.

Convergence of the algorithm has been shown by its authors in Rao and Teh (2012). It follows from the fact that the chain has the stationary distribution $p(X|Y)$ and is irreducible and aperiodic, provided that $R(t; s) > Q(t; s)$.

4. Main result

Let A be the transition kernel of the Markov chain X_m generated by Rao and Teh's algorithm. Let Π be the target distribution. (It is the posterior distribution of X given Y . In this paper we consider only Monte Carlo randomness, so Y is fixed and can be omitted in notation.)

Theorem 1. *Consider a hidden MJP in which the evidence is of the form (1). Assume that*

1. *there exists an irreducible matrix $Q^{\min}$ such that $Q^{\min}(s, s') \leq Q(t; s, s')$ for all $s, s' \in \mathcal{S}$, $s \neq s'$, $t \in [t^{\min}, t^{\max}]$,*
2. *there exists $\eta > 0$ such that $Q(t; s)/R(t; s) \leq 1 - \eta$ for all $s \in \mathcal{S}$, $t \in [t^{\min}, t^{\max}]$,*
3. *there exists $r^{\max} < \infty$ such that $R(t; s) \leq r^{\max}$ for all $s \in \mathcal{S}$, $t \in [t^{\min}, t^{\max}]$.*

Then the chain X_m produced by Rao and Teh's Algorithm 2 is geometrically ergodic, i.e. there exist constant $\gamma < 1$ and function M such that for every initial trajectory X such that $L(Y|X) > 0$ and every m ,

$$\|A^m(X, \cdot) - \Pi(\cdot)\|_{\text{tv}} \leq \gamma^m M(X).$$

We begin with some auxiliary results. In Lemmas 2 and 5 we consider an inhomogeneous Markov chain $S_0, S_1, \dots, S_n$ on a finite state space $\mathcal{S}$ with the joint probability distribution given by

$$\mathbb{P}(S_0 = s_0, S_1 = s_1, \dots, S_n = s_n) \propto \nu(s_0)g_0(s_0) \prod_{i=1}^n P_i(s_{i-1}, s_i)g_i(s_i),$$

where P_i are stochastic matrices and g_i are non-negative functions (formula (5) shows that the conditional distribution of skeleton is of this form). Assuming

that $0 < n_0 \leq i < n$, we define

$$P_{i-n_0:i}(s_{i-n_0}, s) = \sum_{\{s_{i-n_0+1}, \dots, s_{i-1}\}} \left(\prod_{l=i-n_0+1}^{i-1} P_l(s_{l-1}, s_l) \right) P_i(s_{i-1}, s). \quad (7)$$

and

$$P_{i-n_0:i}^g(s_{i-n_0}, s) = \sum_{\{s_{i-n_0+1}, \dots, s_{i-1}\}} \left(\prod_{l=i-n_0+1}^{i-1} P_l(s_{l-1}, s_l) g_l(s_l) \right) P_i(s_{i-1}, s).$$

Lemma 2. Consider fixed n_0 and i such that $1 \leq n_0 \leq i \leq n-1$. Assume that

1. for some $\xi > 0$ inequality $P_{i-n_0:i}(s_{i-n_0}, s) \geq \xi$ holds for all $s_{i-n_0}, s \in \mathcal{S}$,
2. for some $\eta > 0$ inequality $P_{i+1}(s, s) \geq \eta$ holds for all $s \in \mathcal{S}$,
3. for some $g_l^{\min}$ and $g_l^{\max}$ we have $g_l^{\min} \leq g_l(s) \leq g_l^{\max}$ for all $s \in \mathcal{S}$, $l \in [i-n_0+1:i]$.

Then

$$\mathbb{P}(S_i = s | S_{i+1} = s) \geq \delta_i, \quad \text{where} \quad \delta_i = \xi \eta \prod_{l=i-n_0+1}^i \frac{g_l^{\min}}{g_l^{\max}}.$$

Proof. We condition additionally on $S_{i-n_0} = s_{i-n_0}$ and use two-sided Markov property to obtain

$$\begin{aligned} \mathbb{P}(S_i = s | S_{i-n_0} = s_{i-n_0}, S_{i+1} = s) &= \frac{P_{i-n_0:i}^g(s_{i-n_0}, s) g_i(s) P_{i+1}(s, s)}{\sum_{s'} P_{i-n_0:i}^g(s_{i-n_0}, s') g_i(s') P_{i+1}(s', s)} \\ &\geq \frac{\left(\prod_{l=i-n_0+1}^i g_l^{\min} \right) P_{i-n_0:i}(s_{i-n_0}, s) P_{i+1}(s, s)}{\left(\prod_{l=i-n_0+1}^i g_l^{\max} \right) \sum_{s'} P_{i-n_0:i}(s_{i-n_0}, s') P_{i+1}(s', s)} \geq \xi \eta \prod_{l=i-n_0+1}^i \frac{g_l^{\min}}{g_l^{\max}}, \end{aligned}$$

because $\sum_{s'} P_{i-n_0:i}(s_{i-n_0}, s') = 1$ and $P_{i+1}(s', s) \leq 1$. Since the right hand side of the obtained bound does not depend on s_{i-n_0} , by taking the sum over all possible values of s_{i-n_0} we obtain the desired result. Finally let us remark that conclusion of the lemma remains trivially true if $g_l^{\min} = 0$ for some index $l \in [i-n_0+1:i]$ and thus $\delta_i = 0$. $\square$

Remark 3. In Lemma 2 we bound from below the backward transition probability used by the FFBS algorithm. However, the identical inequality is true also for the forward transition probability $\mathbb{P}(S_i = s | S_{i-1} = s)$.

Remark 4. In the time-homogeneous case, when $P_i = P$, the first assumption of Lemma 2 is essentially equivalent to irreducibility and aperiodicity of matrix P . Note also that the two constants ξ and η play different roles in Rao and Teh's algorithm.

Lemma 5. Let the assumptions of Lemma 2 hold for all $i \in [n_0:n-1]$. Then

$$\mathbb{E}|J| \leq n+1 - \sum_{i=n_0}^{n-1} \delta_i,$$

where $J = \{i \in [1:n]: S_{i-1} \neq S_i\} \cup \{0\}$.

Proof. Note that $|J| = 1 + \sum_{i=0}^{n-1} \mathbb{I}(S_i \neq S_{i+1})$. We apply Lemma 2 to each $i \in [n_0 : n-1]$ to obtain $\mathbb{P}(S_i = S_{i+1} | S_{i+1}) \geq \delta_i$ and consequently

$$\mathbb{E}\mathbb{I}(S_i \neq S_{i+1}) = \mathbb{E}\mathbb{P}(S_i \neq S_{i+1} | S_{i+1}) \leq 1 - \delta_i.$$

For $i < n_0$ we apply the trivial bound $\mathbb{E}\mathbb{I}(S_i \neq S_{i+1}) \leq 1$. $\square$

In the next proposition we establish a geometric drift condition for the Markov chain $X_0, X_1, \dots, X_m, \dots$. Consider a single step, that is transition from $X_{m-1} = X$ to $X_m = X'$. The dependence on the input trajectory X (and also on Y) is implicitly assumed but indicated only when necessary. Recall that $|J(X)|$ is the number of true jumps of the trajectory $X([t^{\min}, t^{\max}])$.

Proposition 6 (Drift Condition). *Under the assumptions of Theorem 1, there exist $\delta > 0$ and $c < \infty$ such that in a single step of Rao and Teh's algorithm, $\mathbb{E}(|J(X')||X) \leq (1 - \delta)|J(X)| + c$.*

Proof. Let us analyse what happens in both two stages (V) and (S) of Algorithm 2. The initial X is fixed. In stage (V) we add a new set V of potential jumps. Since $|V|$ has the Poisson distribution with intensity

$$\int_{t^{\min}}^{t^{\max}} (R(t; X(t)) - Q(t; X(t))) dt,$$

we have $\mathbb{E}|V| \leq \mu := r^{\max}(t^{\max} - t^{\min})$. Thus we obtain T' with

$$\mathbb{E}(|T'| | X) \leq |J(X)| + \mu. \quad (8)$$

In stage (S) the set T' is “thinned” to T'' . We now condition on T' and write $n+1 = |T'|$. Our next goal is to bound the expected value of $|J'|$ using Lemma 5. Recall that P_i and g_i are defined by (5) and (6). We proceed to verifying the assumptions of Lemma 2. First we use Assumption 1 of Theorem 1 to choose n_0 such that $(Q^{\min})^{n_0}(s, s') \geq \varepsilon$ for some $\varepsilon > 0$ and for all $s, s' \in \mathcal{S}$. Without loss of generality we can assume that the diagonal of $Q^{\min}$ is $\eta^{r^{\max}}$. Then $Q^{\min}$ is aperiodic, irreducible and satisfies $Q^{\min}(s, s')/r^{\max} \leq P_i(s, s')$ for all s, s' and all i . If we assume that $n > n_0$ then $P_{i-n_0:i}$ is defined by (7) for all $i \in [n_0 : n-1]$ and satisfies

$$P_{i-n_0:i}(s_{i-n_0}, s) \geq \frac{(Q^{\min})^{n_0}(s_{i-n_0}, s)}{(r^{\max})^{n_0}} \geq \xi := \varepsilon/(r^{\max})^{n_0}, \quad (9)$$

so Condition 1 of Lemma 2 holds. Condition 2 of Lemma 2 follows directly from Assumption 2 of Theorem 1. Condition 3 of Lemma 2, that is the two-sided bound $g_l^{\min} \leq g_l(s) \leq g_l^{\max}$, is fulfilled with $g_l^{\min}$ and $g_l^{\max}$ defined as follows. The upper bound is uniform for all l , $g_l^{\max} = g^{\max}$, where

$$g^{\max} := r^{\max} \max_{s,j} L(y_j | s)^k.$$

As for the lower bound, we consider two cases. Note that formula (6) for $g_l(s)$ includes the “likelihood factor” $\prod L_j(Y_j | s)$ for at most k indices l , simply because

there are k points t_j^{obs} . For these indices l we use the trivial bound $g_l^{\min} = 0$. For the remaining $n - k$ indices, by (6), we have $g_l^{\min} = g^{\min}$, where $g^{\min}$ does not depend on l . Indeed, we can choose $g^{\min} = q^{\min} \exp[-r^{\max}(t^{\max} - t^{\min})]$, where $q^{\min} = \min_s Q^{\min}(s) = \min_s \sum_{s' \neq s} Q^{\min}(s, s')$. Assumption 1 of Theorem 1 entails $R(t; s) \geq q^{\min} > 0$. We now use Lemma 5, with $n + 1 = |T'|$ to obtain

$$\mathbb{E}(|J'| | T') \leq n + 1 - \sum_{i=n_0}^{n-1} \delta_i, \quad (10)$$

on the event $|T'| - 1 > n_0$. Notice that $\delta_i = 0$ only if $g_l^{\min} = 0$ for some index $l \in [i - n_0 + 1 : i]$, so there are at most kn_0 zeros among δ_i s. For at least $n - n_0 - kn_0$ nonzero δ_i s we have a lower bound

$$\delta_i = \xi \eta \left(\frac{g^{\min}}{g^{\max}} \right)^{n_0} =: \delta > 0.$$

Therefore the right hand side of (10) satisfies the following inequality:

$$\begin{aligned} n + 1 - \sum_{i=n_0}^{n-1} \delta_i &\leq n + 1 - (n - n_0 - kn_0)\delta \\ &= (1 - \delta)(n + 1) + (kn_0 + n_0 + 1)\delta \\ &\leq (1 - \delta)(n + 1) + (k + 1)(n_0 + 1). \end{aligned}$$

Using (10) for $|T'| - 1 > n_0$ and the trivial bound $|J'| \leq |T'| \leq n_0 + 1$ for $|T'| - 1 \leq n_0$ we obtain

$$\mathbb{E}(|J'| | T') \leq (1 - \delta)|T'| + (k + 1)(n_0 + 1).$$

Finally, by (8) we get

$$\mathbb{E}(|J(X')| | X) = \mathbb{E}(\mathbb{E}(|J'| | T') | X) \leq (1 - \delta)(|J(X)| + \mu) + (k + 1)(n_0 + 1).$$

The conclusion of the proposition follows. $\square$

Proposition 7 (Small Set Condition). *The set $\{X : |J(X)| \leq h\}$ is 1-small for every h , i.e. there exists a probability measure Φ and a constant $\beta > 0$ such that $A(X, dX') \geq \beta \Phi(dX')$, whenever $|J(X)| \leq h$.*

Recall that A denotes the transition kernel of the Markov chain defined via Algorithm 2. Φ is called a regeneration measure.

Proof. The scheme of our proof is the following. We will define a sequence of states $s^* = (s_0^*, \dots, s_n^*)$ and a sequence of times $t^* = (t_0^*, \dots, t_n^*)$. Both these sequences are deterministic and fixed. The regeneration measure $\Phi(dX')$ is described in terms of s^* and t^* as follows:

$$\begin{aligned} T'_i &\sim \text{Uniform}(t_{i-1}^*, t_i^*) \text{ independently for } i = 1, \dots, n; \\ S'_i &= s_i^* \text{ for } i = 0, 1, \dots, n. \end{aligned} \quad (11)$$

Trajectory X' is determined by (T', S') as described in Section 3. Note that the skeleton S' is deterministic and random vector $(T'_1, \dots, T'_n)$ has the uniform distribution on the set

$$\mathcal{T} = \{(t_1, \dots, t_n) : t_{i-1}^* \leq t_i \leq t_i^* \text{ for } i = 1, \dots, n\}.$$

We will show that Algorithm 2 can be equivalently executed in such a way that the resulting X' is distributed according to Φ with probability at least $\beta > 0$, provided that $|J(X)| \leq h$ (β must not depend on X ; it will be defined in the course of our proof).

Now we proceed to details of our construction. To define s^* and t^* , let us first choose a sequence $s^\dagger = (s_1^\dagger, s_2^\dagger, \dots, s_k^\dagger)$ such that

$$\prod_{j=1}^k L_j(Y_j | s_j^\dagger) =: L^\dagger > 0.$$

Now we are going to use Assumptions 1 and 3 of Theorem 1. By irreducibility of matrix $Q^{\min}$ we can embed $s^\dagger$ in a skeleton s^* , which has probability bounded below for the chain with transition matrices P_i appearing in (5) (whatever the choice of the times of jumps, on which these matrices depend). Put differently, we define a sequence $s^* = (s_0^*, \dots, s_n^*)$ for some $n \geq k$ such that $s^\dagger$ is a subsequence of s^* , $s_{i-1}^* \neq s_i^*$ and, uniformly in t , we have

$$\nu(s_0^*) \prod_{i=1}^n P(t; s_i^*, s_{i+1}^*) \geq \nu(s_0^*) \prod_{i=1}^n \frac{Q^{\min}(s_i^*, s_{i+1}^*)}{r^{\max}} =: \beta_1^* > 0.$$

To get a sequence of times “compatible with” the skeleton s , we embed the sequence $t^{\text{obs}} = (t_1^{\text{obs}}, \dots, t_k^{\text{obs}})$ in a longer sequence $t^* = (t_0^*, t_1^*, \dots, t_n^*)$. More precisely, we choose a sequence $t^{\min} = t_0^* < t_1^* < \dots < t_n^* < t^{\max}$ such that $t_{ij}^* = t_j^{\text{obs}}$ implies $s_{ij}^* = s_j^\dagger$ for $j = 1, \dots, k$.

Fix X with $|J(X)| \leq h$. We are going to describe a special way in which Algorithm 2 can be executed. Recall that Assumptions 1 and 2 of Theorem 1 ensure that $R(t; s) \geq Q(t; s) = \sum_{s' \neq s} Q(t; s, s') \geq \sum_{s' \neq s} Q^{\min}(s, s') \geq q^{\min} > 0$. Using again Assumption 2 we obtain

$$R(t; s) - Q(t; s) \geq \eta R(t; s) \geq \eta q^{\min} =: \epsilon > 0.$$

In stage (V) we can independently sample two Poisson processes on the interval $[t^{\min}, t^{\max}]$, say V^0 and V^{rest} , with intensities ϵ and $R(t; s) - Q(t; s) - \epsilon$, respectively. Next let $V = V^0 \cup V^{\text{rest}}$ and $T' = J(X) \cup V$. Note that

$$\mathbb{P}(V^0 \in \mathcal{T}) =: \beta^0 > 0.$$

Moreover, since V^{rest} is a Poisson process with intensity bounded by $r^{\max}$ we have

$$\mathbb{P}(V^{\text{rest}} = \emptyset) \geq \exp[-r^{\max}(t^{\max} - t^{\min})] =: \beta^{\text{rest}} > 0.$$

In stage (S) of the Algorithm 2 we construct skeleton S' . Although the actual sampling from $p(S'|T', Y)$ is by FFBS, an equivalent result can be obtained via rejection sampling. We have $\prod_{i=0}^{|T'|} g_i(S'_i) \leq (r^{\max})^{|T'|} L^{\max}$, where

$$L^{\max} = \prod_{j=1}^k \max_s L_j(Y_j|s).$$

The rejection sampling proceeds as follows.

(S1) Simulate Markov chain S' (of length $|T'|$) with transition matrices

$$P_i = P(T'_i; \cdot, \cdot)$$

and initial distribution ν , c.f. (5).

(S2) Accept the skeleton S' with probability $\prod_{i=0}^{|T'|} g_i(S'_i) / [(r^{\max})^{|T'|} L^{\max}]$. If the skeleton is not accepted then go to (S1).

(Of course the rejection method is highly inefficient and is considered only to clarify presentation.)

We consider the following random events $\mathcal{E}_i$:

- $\mathcal{E}_0$: in stage (V) we obtain $T' = J(X) \cup V^0$ and $V^0 \in \mathcal{T}$.
- $\mathcal{E}_1$: in stage (S1) all points belonging to $J(X)$ are changed to virtual jumps, while jumps at V^0 form the skeleton s^* .
- $\mathcal{E}_2$: in stage (S2) we accept the skeleton obtained in stage (S1).

We can see that

- $\mathcal{E}_0$ happens with probability at least $\beta^0 \beta^{\text{rest}}$.
- Given that $\mathcal{E}_0$ has happened, the probability of $\mathcal{E}_1$ is at least

$$\beta_1^* \eta^{J(X)} \geq \beta_1^* \eta^h =: \beta_1,$$

because $P(T'_i, s, s) \geq \eta$.

- Given that $\mathcal{E}_0$ and $\mathcal{E}_1$ have happened, the probability of $\mathcal{E}_2$ is at least

$$\left(\frac{q^{\min}}{r^{\max}} \right)^{h+n} \exp[-r^{\max}(t^{\max} - t^{\min})] \frac{L^\dagger}{L^{\max}} =: \beta_2,$$

in view of (6), because $|T'| = |J(X)| + |V^0| \leq h + n$.

(Of course, all the probabilities in the above statements are conditional on X .) Putting everything together, $\mathcal{E} := \mathcal{E}_0 \cap \mathcal{E}_1 \cap \mathcal{E}_2$ happens with probability at least $\beta := \beta^0 \beta^{\text{rest}} \beta_1 \beta_2 > 0$. If $\mathcal{E}$ happens then the output X' of Algorithm 2 is independent of the input X and has the probability distribution $\Phi(dX')$ described in the beginning of this proof. $\square$

Theorem 1 immediately follows from Propositions 6 and 7, see for example Roberts and Rosenthal (2004, Th. 9). (In the cited paper the drift condition appears in a form slightly different than in our Proposition 6; however (Johnson, 2009, Lem. 2.1) shows that both formulations are equivalent.)

Remark 8. Note that Theorem 1 holds also for other observation models than given by (1). For example, we could consider the observed object Y of quite general nature but assume that $L(Y|X) \geq \varepsilon > 0$. The proofs of the drift condition and especially of the small set condition would be then much simpler. However, assumption (1) does not imply $L(Y|X) \geq \varepsilon > 0$, and we think that this latter condition is less realistic in applications.

5. Continuous time Bayesian networks

Let $(\mathcal{N}, \mathcal{K})$ denote a directed graph with possible cycles. We write $w \rightarrow u$ instead of $(w, u) \in \mathcal{K}$. For every node $w \in \mathcal{N}$ consider a corresponding space $\mathcal{S}_w$ of possible states. Assume that each space $\mathcal{S}_w$ is finite. We consider a continuous time homogeneous Markov process on the product space $\mathcal{S} = \prod_{w \in \mathcal{N}} \mathcal{S}_w$. Thus a state $\underline{s} \in \mathcal{S}$ is a configuration $\underline{s} = (s_w) = (s_w)_{w \in \mathcal{N}}$, where $s_w \in \mathcal{S}_w$. If $\mathcal{W} \subseteq \mathcal{N}$ then we write $s_{\mathcal{W}} = (s_w)_{w \in \mathcal{W}}$ for configuration $\underline{s}$ restricted to nodes in $\mathcal{W}$. We also use notation $\mathcal{S}_{\mathcal{W}} = \prod_{w \in \mathcal{W}} \mathcal{S}_w$, so that we can write $s_{\mathcal{W}} \in \mathcal{S}_{\mathcal{W}}$. The set $\mathcal{W} \setminus \{w\}$ will be denoted by $\mathcal{W} - w$ and $\mathcal{N} \setminus \{w\}$ simply by $-w$. We define the set of parents of node w by

$$\text{pa}(w) = \{u \in \mathcal{N} : u \rightarrow w\},$$

and we define the set of children of node w by

$$\text{ch}(w) = \{u \in \mathcal{N} : w \rightarrow u\}.$$

Suppose we have a family of functions $Q_w : \mathcal{S}_{\text{pa}(w)} \times (\mathcal{S}_w \times \mathcal{S}_w) \rightarrow [0, \infty)$. For fixed $c \in \mathcal{S}_{\text{pa}(w)}$, we consider $Q_w(c; \cdot, \cdot)$ as a conditional intensity matrix (CIM) at node w (only off-diagonal elements of this matrix have to be specified, the diagonal ones are irrelevant). The state of a CTBN at time t is a random element $X(t)$ of the space $\mathcal{S}$ of configurations. Let $X_w(t)$ denote its w th coordinate. The process $\{(X_w(t))_{w \in \mathcal{N}}, t \geq 0\}$ is assumed to be Markov and its evolution can be described informally as follows. Transitions at node w depend on the current configuration of the parent nodes. If the state of some parent changes, then node w switches to other transition probabilities. Formally, CTBN is a *time-homogeneous* MJP with transition intensities given by

$$Q(\underline{s}, \underline{s}') = \begin{cases} Q_w(s_{\text{pa}(w)}, s_w, s'_w) & \text{if } s_{-w} = s'_{-w} \text{ and } s_w \neq s'_w \text{ for some } w; \\ 0 & \text{if } s_{-w} \neq s'_{-w} \text{ for all } w, \end{cases}$$

for $\underline{s} \neq \underline{s}'$. Define also $Q_w(c; s) = \sum_{s' \neq s} Q_w(c; s, s')$ for $c \in \mathcal{S}_{\text{pa}(w)}$, $s \in \mathcal{S}_w$.

For a CTBN, the density of sample path $X = X([t^{\min}, t^{\max}])$ in a bounded time interval $[t^{\min}, t^{\max}]$ decomposes as follows:

$$p(X) = \nu(X(t^{\min})) \prod_{w \in \mathcal{N}} p(X_w \| X_{\text{pa}(w)}), \quad (12)$$

where ν is the initial distribution on $\mathcal{S}$ and $p(X_w \| X_{\text{pa}(w)})$ is the density of piecewise homogeneous MJP with intensity matrix equal to $Q_w(c; \cdot, \cdot)$ in every time

sub-interval such that $X_{\text{pa}(w)} = c$. Formulas for the density of CTBN appear e.g. in Nodelman, Shelton and Koller (2002b, Sec. 3.1), Fan, Xu and Shelton (2010, Eq. 2), Fan and Shelton (2008, Eq. 1) and Miasojedow et al. (2014). Our notation $p(X_w \| X_{\text{pa}(w)})$ – with double vertical lines – is consistent with the notion of “conditioning by intervention” see e.g. Pearl (1994) or Lauritzen (2001). In fact, $p(X_w \| X_{\text{pa}(w)})$ coincides with the usual conditional density $p(X_w | X_{\text{pa}(w)})$ in a modified network in which all arrows leading to $\text{pa}(w)$ have been removed. For details see e.g. Lauritzen (2001) in the context of static Bayesian Networks and Miasojedow et al. (2014) in the context of CTBNs. Below we explicitly write an expression for $p(X_w \| X_{\text{pa}(w)})$. We need the following notations:

Let $\iota_w^X(c; s, s')$ denote the number of jumps from $s \in \mathcal{S}_w$ to $s' \in \mathcal{S}_w$ at node w , which occurred when the parent configuration was $c \in \mathcal{S}_{\text{pa}(w)}$. Let $\tau_w^X(c; s)$ be the length of time that node w spent in state $s \in \mathcal{S}_w$ when the parent configuration was $c \in \mathcal{S}_{\text{pa}(w)}$.

With these notations we can write

$$p(X_w \| X_{\text{pa}(w)}) = \left\{ \prod_{c \in \mathcal{S}_{\text{pa}(w)}} \prod_{s \in \mathcal{S}_w} \prod_{\substack{s' \in \mathcal{S}_w \\ s' \neq s}} Q_w(c; s, s')^{\iota_w^X(c; s, s')} \right\} \times \left\{ \prod_{c \in \mathcal{S}_{\text{pa}(w)}} \prod_{s \in \mathcal{S}_w} \exp[-Q_w(c; s) \tau_w^X(c; s)] \right\}. \quad (13)$$

Let us also write $p^{\text{jump}}(X_w \| X_{\text{pa}(w)})$ and $p^{\text{stay}}(X_w \| X_{\text{pa}(w)})$ for the first and second expression in (13), respectively, to facilitate future references.

The problem of probabilistic reasoning for a CTBN can be formulated as follows. Assume that the available evidence is the *complete* observation of some nodes, say $X_{\mathcal{O}}([t^{\min}, t^{\max}])$ for some set $\mathcal{O} \subset \mathcal{N}$. We are to compute the posterior distribution over unobserved nodes, i.e. on the trajectories of $X_{\mathcal{W}}([t^{\min}, t^{\max}])$, where $\mathcal{W} = \mathcal{N} \setminus \mathcal{O}$. The basic idea, proposed in (Rao and Teh, 2013) is to use “Algorithm 2 within Gibbs sampler”. Let us fix a node $w \in \mathcal{W}$. By (12), the full conditional distribution is the following.

$$p(X_w | X_{-w}) \propto \nu(X_w(t^{\min}) | X_{-w}(t^{\min})) p(X_w \| X_{\text{pa}(w)}) \prod_{u \in \text{ch}(w)} p(X_u \| X_{\text{pa}(u)}). \quad (14)$$

The density $\nu(X_w(t^{\min}) | X_{-w}(t^{\min})) p(X_w \| X_{\text{pa}(w)})$ corresponds to a piecewise homogeneous Markov process and can be treated as the prior distribution in Algorithm 2. The expression $\prod_{u \in \text{ch}(w)} p(X_u \| X_{\text{pa}(u)})$ can be treated as likelihood. In Algorithm 2 we can use “instrumental” intensities $R_w(t; s)$ different for different nodes w , and possibly time-inhomogeneous (in practical implementations however, R_w will usually be time-homogeneous).

Algorithm 3 Random scan Gibbs sampler for CTBN.

input: previous trajectory $X_{\mathcal{W}}$ and evidence $X_{\mathcal{O}}$.
 Choose at random $w \in \mathcal{W}$ (according to some probability distribution on $\mathcal{W}$).
 Update X_w to X'_w . Apply Algorithm 2 with the target distribution $p(X_w|X_{-w})$ given by (14) $\{ X_{-w}$ includes all nodes $u \in \mathcal{O}$ as well as $u \in \mathcal{W}$, $u \neq w$ $\}$.
return new trajectory $X'_{\mathcal{W}} = (X_{\mathcal{W}-w}, X'_w)$.

Theorem 9. Consider a CTBN in which we observe trajectories

$$X_{\mathcal{O}}([t^{\min}, t^{\max}]).$$

Assume that

1. for every $w \in \mathcal{W} = \mathcal{N} \setminus \mathcal{O}$ there exists an irreducible matrix $Q_w^{\min}$ such that $Q_w^{\min}(s, s') \leq Q_w(c; s, s')$ for all $s, s' \in \mathcal{S}_w$, $s \neq s'$, $c \in \mathcal{S}_{\text{pa}(w)}$,
2. there exists $\eta > 0$ such that $Q_w(c; s)/R_w(t; s) \leq 1 - \eta$ for all $w \in \mathcal{W}$, $s \in \mathcal{S}_w$, $c \in \mathcal{S}_{\text{pa}(w)}$ and $t \in [t^{\min}, t^{\max}]$,
3. there exists $r^{\max} < \infty$ such that $R_w(t; s) \leq r^{\max}$ for all $w \in \mathcal{W}$, $s \in \mathcal{S}_w$, $t \in [t^{\min}, t^{\max}]$,
4. for every $w \in \mathcal{N}$, the supports of $Q_w(c; \cdot, \cdot)$ do not depend on the parent configuration $c \in \mathcal{S}_{\text{pa}(w)}$, i.e. $Q_w(c; s, s') > 0$ implies $Q_w(c'; s, s') > 0$.

Then the chain $(X_{\mathcal{W}})_m$ produced by Algorithm 3 is geometrically ergodic.

As in Section 4, we will show a drift condition towards a small set. The Lyapunov function in the CTBN setting will be the global number of jumps $|J(X_{\mathcal{W}})| = \sum_{w \in \mathcal{W}} |J(X_w)|$. Let us begin with an elementary fact, needed in the proof of the drift condition.

Lemma 10. Let $0 \leq \varrho < 1$. For any $x_1, \dots, x_n$ such that $\sum_{i=1}^n x_i \leq nb$,

$$\sum_{i=1}^n \varrho^{x_i} \geq n\varrho^b.$$

Proof. The function $x \mapsto \varrho^x$, for $0 \leq \varrho < 1$ is convex and decreasing, so that

$$\frac{1}{n} \sum_{i=1}^n \varrho^{x_i} \geq \varrho^{\frac{1}{n} \sum_{i=1}^n x_i} \geq \varrho^b. \quad \square$$

Proposition 11 (Drift Condition). Under the assumptions of Theorem 9, there exist $\varepsilon > 0$ and $c < \infty$ such that in a single step of Algorithm 3,

$$\mathbb{E}(|J(X'_{\mathcal{W}})| | X) \leq (1 - \varepsilon)|J(X_{\mathcal{W}})| + c.$$

Proof. For a given $X_{\mathcal{W}}$, there always exists a node $w \in \mathcal{W}$ such that

$$|J(X_w)| \geq |J(X_{\mathcal{W}})|/|\mathcal{W}|,$$

e.g. a node with maximum number of jumps. From now on, node w is fixed. We will prove that for some c_0 and $\varepsilon_0 > 0$,

$$\mathbb{E}(|J(X'_w)| | X, \text{node } w \text{ is updated}) \leq (1 - \varepsilon_0)|J(X_w)| + c_0. \quad (15)$$

This is a “local version” of the drift condition. The conclusion of the proposition will easily follow from (15). Indeed, for every node u we have

$$\mathbb{E}(|J(X'_u)||X, \text{node } u \text{ is updated}) \leq |J(X_u)| + \mu,$$

with $\mu := r^{\max}(t^{\max} - t^{\min})$, as in the proof of Proposition 6. Let $p^{\min}$ denote the minimum probability of selecting a node for update in Algorithm 3. The singled out node w is chosen with probability at least $p^{\min}$. Therefore

$$\begin{aligned} \mathbb{E}(|J(X'_w)||X) &\leq p^{\min} ((1 - \varepsilon_0)|J(X_w)| + c_0 + |J(X_{\mathcal{W}-w})|) \\ &\quad + (1 - p^{\min}) (|J(X_{\mathcal{W}})| + \mu) \\ &\leq p^{\min} ((1 - \varepsilon_0)|J(X_w)| + |J(X_{\mathcal{W}})| - |J(X_w)|) \\ &\quad + (1 - p^{\min})|J(X_{\mathcal{W}})| + c_0 + \mu \\ &\leq |J(X_{\mathcal{W}})| - p^{\min}\varepsilon_0|J(X_w)| + c_0 + \mu \\ &\leq |J(X_{\mathcal{W}})| \left(1 - p^{\min} \frac{\varepsilon_0}{|\mathcal{W}|}\right) + c_0 + \mu, \end{aligned}$$

which is the desired conclusion.

It remains to show (15). Since node w is fixed, we will omit subscript w in notation whenever the context permits. The reasoning leading to (15) is similar to the proof of Proposition 6, but more delicate due to a different form of the likelihood. The role of observed Y is now played by X_{-w} . We can write the conditional distribution of the skeleton S (at node w) given times of possible jumps T (at node w) in the same form as (5), namely

$$p(S|T, X_{-w}) \propto \tilde{\nu}(S_0)g_0(S_0) \prod_{i=1}^N P_i(S_{i-1}, S_i)g_i(S_i), \quad (16)$$

where $N + 1 = |T|$, $\tilde{\nu}(X_w(t^{\min})) := \nu(X_w(t^{\min})|X_{-w}(t^{\min}))$ and P_i s are defined by $P_i(s, s') = P(T_i, s, s')$, as in (5), with

$$P(t; s, s') = \begin{cases} \frac{Q_w(X_{\text{pa}(w)}(t); s, s')}{R_w(t; s)} & \text{if } s \neq s'; \\ 1 - \frac{Q_w(X_{\text{pa}(w)}(t); s)}{R_w(t; s)} & \text{if } s = s'. \end{cases}$$

(The “prior distribution” at node w is that of a piecewise-homogeneous MJP). However, the expressions for the g_i s are different than (6). We now have

$$\begin{aligned} g_{i-1}(s) &= R_w(T_i, s) \exp \left[- \int_{T_{i-1}}^{T_i} R_w(t; s) dt \right] \times \prod_{u \in \text{ch}(w)} p(X_u^{(i)} \| X_{\text{pa}(u)}^{(i, w, s)}), \\ &\quad \text{for } i = 1, \dots, N, \\ g_N(s) &= \exp \left[- \int_{T_N}^{t^{\max}} R_w(t; s) dt \right] \times \prod_{u \in \text{ch}(w)} p(X_u^{(N)} \| X_{\text{pa}(u)}^{(N, w, s)}), \end{aligned} \quad (17)$$

where $X^{(i)} = X([T_{i-1}, T_i])$ is the process restricted to an inter-jump-at- w interval and $X_{\text{pa}(u)}^{(i,w,s)}$ denotes the trajectories of $X_{\text{pa}(u)}^{(i)}$ with $X_w^{(i)}$ replaced by $s \in \mathcal{S}_w$. The likelihood parts $p(\cdot|\cdot)$ in equation (17) can be decomposed into $p(\cdot|\cdot) = p^{\text{jump}}(\cdot|\cdot) \times p^{\text{stay}}(\cdot|\cdot)$, c.f. (13). If we write

$$g_i(s) = h_i(s) \prod_{u \in \text{ch}(w)} p^{\text{jump}}(X_u^{(i)} \| X_{\text{pa}(u)}^{(i,w,s)}),$$

then h_i is easy to bound (at least qualitatively), because by (13) we have

$$\exp[-r^{\max}(t^{\max} - t^{\min})|\mathcal{N}|] \leq \prod_{u \in \text{ch}(w)} p^{\text{stay}}(X_u^{(i)} \| X_{\text{pa}(u)}^{(i,w,s)}) \leq 1.$$

The part of h_i which corresponds to the prior can be bounded analogously as in the proof of Proposition 6 and thus we obtain

$$q^{\min} \exp[-r^{\max}(t^{\max} - t^{\min})] \leq h_i(s) \leq r^{\max}, \quad (18)$$

with $q^{\min} = \min_w \min_s \sum_{s' \neq s} Q_w^{\min}(s, s')$, c.f. Assumption 1 of Theorem 9.

We are now left with a task of bounding the expression with p^{jump} . This is more difficult, because this part of the likelihood depends on $|J(X_{\mathcal{W}})|$. By Assumptions 2 and 3 of Theorem 9, we have $Q_u(c; s, s') \leq r^{\max}$ for all $s \neq s'$, $s, s' \in \mathcal{S}_u$, $c \in \mathcal{S}_{\text{pa}(u)}$, for all u . To obtain a lower bound, we define $q_+^{\min}$ as the minimum of *nonzero* values of $Q_u(c; s, s')$, $s \neq s'$, $s, s' \in \mathcal{S}_u$, $c \in \mathcal{S}_{\text{pa}(u)}$, all u . From (13) it follows that

$$(q_+^{\min})^{|J(X_{\text{ch}(w)}^{(i)})|} \leq \prod_{u \in \text{ch}(w)} p^{\text{jump}}(X_u^{(i)} \| X_{\text{pa}(u)}^{(i,w,s)}) \leq (r^{\max})^{|J(X_{\text{ch}(w)}^{(i)})|}. \quad (19)$$

Note that Assumption 4 of Theorem 9 is needed to justify the lower bound above. Indeed, under the obvious assumption that X is possible, i.e. $p(X) > 0$, the jumps of $X_u^{(i)}$ are possible under the configuration $X_{\text{pa}(u)}^{(i)}$. By Assumption 4 of Theorem 9, they must be possible also under the configuration $X_{\text{pa}(u)}^{(i,w,s)}$. Combining (18) and (19) we obtain

$$a_1 \varrho_1^{|J(X_{\text{ch}(w)}^{(i)})|} \leq g_i(s) \leq a_2 \varrho_2^{|J(X_{\text{ch}(w)}^{(i)})|}$$

for some constants a_1, a_2 and ϱ_1, ϱ_2 which depend only on the parameters of the network and on the instrumental intensity R (they depend neither on w nor on i). Of course, $|J(X_{\text{ch}(w)}^{(i)})| = |J(X_{\text{ch}(w) \cap \mathcal{W}}^{(i)})| + |J(X_{\text{ch}(w) \cap \mathcal{O}}^{(i)})|$. Since $X_{\mathcal{O}}$ is fixed, the parts with the exponent $|J(X_{\text{ch}(w) \cap \mathcal{O}}^{(i)})|$ can be absorbed in constants, which leads to the bound

$$\tilde{a}_1 \varrho_1^{|J(X_{\text{ch}(w) \cap \mathcal{W}}^{(i)})|} \leq g_i(s) \leq \tilde{a}_2 \varrho_2^{|J(X_{\text{ch}(w) \cap \mathcal{W}}^{(i)})|}. \quad (20)$$

Assume that stage (V) of Algorithm 2 has been completed, resulting in a new set T' of potential times of jumps at node w . Just as in the proof of Proposition 6, we infer that $\mathbb{E}(|T'| | X, w \text{ is updated}) \leq |J(X_w)| + \mu$, where

$$\mu = r^{\max}(t^{\max} - t^{\min}).$$

Now consider stage (S) of Algorithm 2. After sampling a new skeleton S' at node w , the set T' is “thinned” of to the set of true jumps $J(X'_w)$. We are to bound $|J(X'_w)|$ from above. We are going to apply Lemmas 2 and 5 to the skeleton chain S' at node w which has the probability distribution $p(S'|T', X_{-w})$ given by (16). We now decompose the process X into inter-jump parts according to T' and use the notation $X^{(i)} := X([T'_{i-1}, T'_i])$. From (20) we infer that the conclusion of Lemma 2 holds with

$$\delta_i = a\varrho^{|J(X_{\text{ch}(w) \cap \mathcal{W}}^{(i)})|} \quad (21)$$

for some constants a, ϱ, n_0 and for $i \geq n_0$. Indeed, we can choose, analogously as in the proof of Proposition 6, the diagonal of $Q^{\min}$, n_0 and ξ such that $(Q_w^{\min})^{n_0}(s, s')/(r^{\max})^{n_0} \geq \xi$, c.f. Assumption 1 of Theorem 9. Then put

$$\varrho = (\varrho_2/\varrho_1)^{n_0}$$

and

$$a = (\tilde{a}_2/\tilde{a}_1)\xi\eta/|\mathcal{S}_w|,$$

where η appears in Assumption 2 of Theorem 9.

We are now prepared to use Lemma 10. To verify its assumption, it is necessary to bound the sum of the exponents in (21). Recall that $J(X_u^{(i)})$ is the set of jumps at u in the time interval between consecutive potential jumps at w , i.e. $[T'_{i-1}, T'_i]$. Therefore

$$\begin{aligned} \sum_{i=n_0}^{|T'|-2} |J(X_{\text{ch}(w) \cap \mathcal{W}}^{(i)})| &\leq |J(X_{\mathcal{W}})| \\ &\leq |J(X_w)| \cdot |\mathcal{W}| \leq |T'| \cdot |\mathcal{W}| \leq 2|T' - n_0 - 1||\mathcal{W}|, \end{aligned}$$

where the last inequality holds for $|J(X_w)| \geq 2(n_0 + 1)$, Lemma 10 implies that

$$\sum_{i=n_0}^{|T'|-2} \delta_i \geq a(|T'| - n_0 - 1)\varrho^{2|\mathcal{W}|}.$$

Lemma 5 implies that

$$\begin{aligned} \mathbb{E}(|J(X'_w)|T') &\leq \max \left\{ 2(n_0 + 1), |T'| - \sum_{i=n_0}^{|T'|-2} \delta_i \right\} \\ &\leq |T'| (1 - a\varrho^{2|\mathcal{W}|}) + (a + 2)(n_0 + 1). \end{aligned}$$

It is now enough to use $\mathbb{E}(|T'| | X, w \text{ is updated}) \leq |J(X_w)| + \mu$ to complete the proof. $\square$

Note that the evidence, i.e. the trajectory $X_{\mathcal{O}}([t^{\min}, t^{\max}])$ is considered as fixed. In particular the constants in Proposition 11 may depend on $X_{\mathcal{O}}$.

The following proposition is an analogue of Proposition 7 in the CTBN setting. Now A denotes the transition kernel of the Markov chain defined via Algorithm 3.

Proposition 12. *The set $\{X_{\mathcal{W}} : |J(X_{\mathcal{W}})| \leq h\}$ is $|\mathcal{W}|$ -small for every h , i.e. there exists a probability measure Φ and a constant $\beta > 0$ such that*

$$A^{|\mathcal{W}|}(X_{\mathcal{W}}, dX'_{\mathcal{W}}) \geq \beta \Phi(dX'_{\mathcal{W}}),$$

whenever $|J(X_{\mathcal{W}})| \leq h$.

Proof. In contrast to the drift condition, the proof of the small set condition is easier in the present setting. Under the assumptions of Theorem 9, for the regeneration measure Φ we can take the measure concentrated at a deterministic, constant trajectory $s_{\mathcal{W}}^* = (s_u^*, u \in \mathcal{W}) \in \mathcal{S}_{\mathcal{W}}$. The only requirement is that $\tilde{\nu}(s_{\mathcal{W}}^*) > 0$, where $\tilde{\nu}$ is the posterior initial distribution of $X_{\mathcal{W}}(t^{\min})$, given $X_{\mathcal{O}}(t^{\min})$. We are to bound from below the probability that $X'_{\mathcal{W}}(t) = s_{\mathcal{W}}^*$, for $t \in [t^{\min}, t^{\max}]$, where $X'_{\mathcal{W}}$ is the result of $|\mathcal{W}|$ steps of Algorithm 3, starting from an arbitrary $X_{\mathcal{W}}$ such that $|J(X_{\mathcal{W}})| \leq h$.

With probability at least $(p^{\min})^{|\mathcal{W}|}$, Algorithm 3 in $|\mathcal{W}|$ steps will visit and update all nodes belonging to $\mathcal{W}$. Let us now consider a single step, in which $w \in \mathcal{W}$ is updated via Algorithm 2. In the rest of the proof w is arbitrary but fixed. (T', S') denote the times of potential jumps and the skeleton of X'_w . Since Assumptions 1, 2 and 3 of Theorem 9 are, for a fixed w , essentially the same as Assumptions 1, 2 and 3 of Theorem 1, the reasoning is similar as in the proof of Proposition 7. We assume that stage (S) of Algorithm 2 is executed in a way described in that proof, via rejection sampling. We consider the following random events $\mathcal{E}_i$:

- $\mathcal{E}_0$: in stage (V) we obtain $V = \emptyset$ so $T' = J(X_w)$.
- $\mathcal{E}_1$: in stage (S1) all points belonging to $J(X_w)$ are changed to virtual jumps.
- $\mathcal{E}_2$: in stage (S2) we accept the skeleton obtained in stage (S1).

It is easy to obtain the following lower bounds.

- $\mathcal{E}_0$ happens with probability at least $\exp[-r^{\max}(t^{\max} - t^{\min})] =: \beta_0$, because V is a Poisson process with intensity $R(t; s) - Q(t; s) \leq r^{\max}$.
- Given that $\mathcal{E}_0$ has happened, the probability of $\mathcal{E}_1$ is at least

$$\eta^{|J(X_w)|} \geq \eta^h =: \beta_1,$$

because $P(T'_i, s, s) \geq \eta$.

- Given that $\mathcal{E}_0$ and $\mathcal{E}_1$ have happened, the probability of $\mathcal{E}_2$ is at least

$$\left(\frac{\tilde{a}_1 \varrho_1}{\tilde{a}_2 \varrho_2} \right)^h =: \beta_2,$$

where $\tilde{a}_1, \varrho_1, \tilde{a}_2, \varrho_2$ are constants appearing in (20). Indeed, the acceptance criterion can be $\prod g_i(S'_i) \leq (\tilde{a}_2 \varrho_2)^h$ and it is always true that

$$\prod g_i(S'_i) \geq (\tilde{a}_2 \varrho_1)^h$$

(by (20); note that $\sum_i |J(X_{\text{ch}(w) \cap \mathcal{W}}^{(i)})| \leq |J(X_{\mathcal{W}})| \leq h$ and also the number of summands is at most $|J(X_w)| \leq h$).

Of course, if $\mathcal{E}_0 \cap \mathcal{E}_1 \cap \mathcal{E}_2$ happens then $X'_w(t) = s_w^*$ for $t \in [t^{\min}, t^{\max}]$. Putting everything together, we get $X'_w([t^{\min}, t^{\max}]) = s_w^*$ with probability at least

$$\beta := p^{\min} \beta_0 \beta_1 \beta_2 > 0.$$

In $|\mathcal{W}|$ steps we get $X'_{\mathcal{W}}([t^{\min}, t^{\max}]) = s_{\mathcal{W}}^*$ with probability at least $\tilde{\nu}(s_{\mathcal{W}}^*) \beta^{|\mathcal{W}|}$. The proof is complete. $\square$

Theorem 9 follows from Propositions 11 and 12.

Remark 13. In this paper the focus is on qualitative results. The constants in our bounds are chosen in a way which makes presentation clearer, and we did not attempt to optimize them.

Remark 14. For clarity of presentation we have proved Theorem 9 under the assumption that some nodes of CTBN are fully observed. However, by minor modification of the proofs we can establish geometric ergodicity of Rao and Teh's algorithm in a more general case. Our results remain true if we assume that some nodes are only partially observed with random noise at discrete moments, just as in (1). Clearly, for a drift condition, the likelihood part can be treated in the same way as in proof of Proposition 6. For a small set condition, we can repeat the construction from the proof of Proposition 7 for every node $w \in \mathcal{N}$, and then define the regeneration measure for whole network as a product of regeneration measures for single nodes. The proofs of the key propositions in the more general case are not essentially different but become notationally complicated and awkward. For this reason they are omitted.

Acknowledgements

We would like to thank anonymous reviewers for their comments, which helped us to improve the paper.

References

- BOYS, R. J., WILKINSON, D. J. and KIRKWOOD, T. B. (2008). Bayesian inference for a discretely observed stochastic kinetic model. *Statistics and Computing* **18** 125–135. [MR2390814](#)
- CARTER, C. K. and KOHN, R. (1994). On Gibbs Sampling for State Space Models. *Biometrika* **81** 541–553. [MR1311096](#)
- EL-HAY, T., FRIEDMAN, N. and KUPFERMAN, R. (2008). Gibbs Sampling in Factorized Continuous-Time Markov Processes. In *Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)* 169–178. AUAI Press, Corvallis, Oregon.
- FAN, Y. and SHELTON, C. R. (2008). Sampling for Approximate Inference in Continuous Time Bayesian Networks. In *Tenth International Symposium on Artificial Intelligence and Mathematics*.

- FAN, Y., XU, J. and SHELTON, C. R. (2010). Importance Sampling for Continuous Time Bayesian Networks. *Journal of Machine Learning Research* **11** 2115–2140. [MR2719849](#)
- FRÜHWIRTH-SCHNATTER, S. (1994). Data augmentation and dynamic linear models. *Journal of Time Series Analysis* **15** 183–202. [MR1263889](#)
- GOLIGHTLY, A., HENDERSON, D. and SHERLOCK, C. (2015). Delayed acceptance particle MCMC for exact inference in stochastic kinetic models. *Statistics and Computing* **25** 1039–1055. [MR3375634](#)
- GOLIGHTLY, A. and WILKINSON, D. J. (2011). Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus*.
- GOLIGHTLY, A. and WILKINSON, D. J. (2014). Bayesian inference for Markov jump processes with informative observations. *ArXiv e-prints*. [MR3331772](#)
- JOHNSON, A. A. (2009). Geometric Ergodicity of Gibbs Samplers. *PHD thesis, Univeristy of Minnesota*. [MR2713500](#)
- LAURITZEN, S. L. (2001). Causal inference from graphical models. *Complex stochastic systems* 63–107. [MR1893411](#)
- MIASOJEDOW, B. and NIEMIRO, W. (2016). Geometric ergodicity of Rao and Teh’s algorithm for homogeneous Markov jump processes. *Statistics & Probability Letters* **113** 1–6. [MR3480387](#)
- MIASOJEDOW, B., NIEMIRO, W., NOBLE, J. and OPALSKI, K. (2014). Metropolis-type algorithms for Continuous Time Bayesian Networks. *arXiv preprint arXiv:1403.4035*.
- NODELMAN, U., SHELTON, C. R. and KOLLER, D. (2002a). Continuous time Bayesian networks. In *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence* 378–387.
- NODELMAN, U., SHELTON, C. R. and KOLLER, D. (2002b). Learning continuous time Bayesian networks. In *Proceedings of the Nineteenth conference on Uncertainty in Artificial Intelligence* 451–458. Morgan Kaufmann Publishers Inc.
- PEARL, J. (1994). A Probabilistic Calculus of Actions. *Proceedings of the Tenth International conference on Uncertainty in Artificial Intelligence*.
- RAO, V. and TEH, Y. W. (2012). MCMC for continuous-time discrete-state systems. In *Advances in Neural Information Processing Systems* 701–709.
- RAO, V. and TEH, Y. W. (2013). Fast MCMC sampling for Markov jump processes and extensions. *Journal of Machine Learning Research* **14** 3207–3232. [MR3144463](#)
- ROBERTS, G. O. and ROSENTHAL, J. S. (2004). General state space Markov chains and MCMC algorithms. *Probability Surveys* **1** 20–71. [MR2095565](#)
- SCHWEDER, T. (1970). Composable markov processes. *Journal of applied probability* **7** 400–410. [MR0264755](#)