

**Manuscript version: Author's Accepted Manuscript**

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

**Persistent WRAP URL:**

<http://wrap.warwick.ac.uk/161983>

**How to cite:**

Please refer to published version for the most recent bibliographic citation information.

**Copyright and reuse:**

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

**Publisher's statement:**

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: [wrap@warwick.ac.uk](mailto:wrap@warwick.ac.uk).

# Zig-zag sampling for discrete structures and non-reversible phylogenetic MCMC

Jere Koskela

`j.koskela@warwick.ac.uk`

Department of Statistics

University of Warwick

Coventry CV4 7AL, UK

January 18, 2022

## Abstract

We construct a zig-zag process targeting a posterior distribution defined on a hybrid state space consisting of both discrete and continuous variables. The construction does not require any assumptions on the structure among discrete variables. We demonstrate our method on two examples in genetics based on the Kingman coalescent, showing that the zig-zag process can lead to efficiency gains of up to several orders of magnitude over classical Metropolis–Hastings algorithms, and that it is well suited to parallel computation. Our construction resembles existing techniques for Hamiltonian Monte Carlo on a hybrid state space, which suffers from implementationally and analytically complex boundary crossings when applied to the coalescent. We demonstrate that the continuous-time zig-zag process avoids these complications.

Keywords: Bayesian inference; coalescent; hybrid state space; piecewise deterministic Markov process; zig-zag process

## 1 Introduction

The zig-zag process is a non-reversible, piecewise deterministic Markov process introduced by [BR17, BFR19b] for continuous-time MCMC. It has several advantages over reversible methods such as Metropolis–Hastings [Has70] and Gibbs sampling [GS90]: it avoids diffusive backtracking which slows their mixing, and is rejection-free so that no computation is wasted on rejected moves.

In brief, the generator of the zig-zag process  $(\mathbf{x}_t, \mathbf{v}_t)_{t \geq 0}$  targeting the probability density (with respect to the product of the Lebesgue and counting measures)  $\tilde{\pi}(\mathbf{x}, \mathbf{v}) := \pi(\mathbf{x})/2^d$  on a state space  $X \times \{-1, 1\}^d \subseteq \mathbb{R}^d \times \{-1, 1\}^d$  is

$$Lf(\mathbf{x}, \mathbf{v}) = \sum_{i=1}^d v_i \partial_i f(\mathbf{x}, \mathbf{v}) + \sum_{i=1}^d \lambda_i(\mathbf{x}, \mathbf{v})(f(\mathbf{x}, F_i \mathbf{v}) - f(\mathbf{x}, \mathbf{v})),$$

where  $\partial_i$  is the derivative with respect to  $x_i$ , and  $F_i$  flips the sign of  $v_i$ . The flip rates

$$\lambda_i(\mathbf{x}, \mathbf{v}) := (-v_i \partial_i \log(\tilde{\pi}(\mathbf{x}, \mathbf{v})))^+ \tag{1}$$

with  $(x)^+ := \max\{x, 0\}$ , ensure that  $(\mathbf{x}_t, \mathbf{v}_t)_{t \geq 0}$  leaves  $\tilde{\pi}(\mathbf{x}, \mathbf{v})$  invariant [BFR19b, Theorem 2.2]. In words, the coordinates of  $\mathbf{x}_t$  move at constant velocities  $\mathbf{v}_t$  until a flip at coordinate  $i$ , when the corresponding velocity changes sign.

To date, the zig-zag processes have been applied to targets such as the Curie–Weiss model [BR17] and logistic regression [BFR19b], whose state spaces have simple geometric structures with natural notions of direction and velocity. Discrete variables (other than the velocities) have been restricted to cases with simple partial orders, such as model selection [CFS20, GD21]. We construct a zig-zag process on a general hybrid state space with both continuous and discrete coordinates, without imposing any structure on discrete coordinates. This is done by introducing a separate space of continuous variables for each value of the discrete variable, introducing boundaries into the continuous spaces, and updating the discrete variable when boundaries are hit. The strategy takes advantage of the full generality of piecewise-deterministic Markov processes [Dav93, Section 24], and is resembles similar work for Hamiltonian Monte Carlo (HMC) [DBZM17, NDL20]. Our method can also be seen of as a generalization of the zig-zag process on a restricted domain [BBCD<sup>+</sup>18] to a union of many restricted sub-domains, with jumps between sub-domains at boundary hitting events. We illustrate our method with an application to the coalescent [Kin82]: a tree-valued target with continuous branch lengths, discrete tree topologies with no natural partial order, and no canonical geometric structure.

The coalescent examples illustrate the need for methods which are implementable on complex state spaces. They are also of interest because existing MCMC algorithms for coalescents tend to mix slowly. The key difficulty lies in designing Metropolis–Hastings proposal distributions which combine efficient exploration with a high acceptance rate [MV05, HDPD08, LVH<sup>+</sup>08]. Workarounds consist of empirical searches for efficient proposals [HD12, ASR16] or Metropolis-coupled MCMC [Gey92]. The former does not scale to problems for which empirical optimization is infeasible. The latter helps mixing between modes, but does not overcome low acceptance rates or the backtracking behavior of reversible MCMC.

The zig-zag process has some similarities with HMC [Nea10], which augments the state space with momentum and uses Hamiltonian dynamics to propose large steps which are accepted with high probability, though they are not rejection-free. Like the zig-zag process, HMC requires gradient information and a suitable geometric embedding of the target. [DBZM17] provided those for the coalescent using an orthant complex construction of phylogenetic tree space [BHV01]. Our examples differ from the method of [DBZM17] in three ways. Firstly, we replace the embedding of [BHV01] with  $\tau$ -space [GD16], which is better suited to ultrametric trees. Secondly, the zig-zag process is readily implementable on  $\tau$ -space via Poisson thinning without a numerical integrator such as *leap-prog* [DBZM17, Algorithm 1]. Finally, the zig-zag process has simple boundary behavior between orthants and does not require boundary smoothing [DBZM17, Section 3.3], chiefly because discontinuous gradients are easier to handle on continuous rather than discretized paths.

The rest of the paper is structured as follows. Section 2 defines the zig-zag algorithm with discrete and continuous variables, and proves that it has the desired invariant distribution. Section 3 recalls the coalescent and the  $\tau$ -space embedding. In Sections 4 and 5 we recall the popular infinite and finite sites models of mutation, derive zig-zag processes for each, and demonstrate their performance via simulation studies. Section 6 concludes with a discussion. The algorithms and data sets used in the simulation studies are available at <https://github.com/JereKoskela/tree-zig-zag>.

## 2 Zig-zag on hybrid spaces

The definition of our zig-zag process follows [Dav93, Section 24]. Let  $\mathbb{F}$  be a countable set. For each  $m \in \mathbb{F}$ , let  $\Omega_m^o$  be an open subset of  $\mathbb{R}^d$ ,  $\bar{\Omega}_m^o$  be its closure, and  $\partial\Omega_m^* := \bar{\Omega}_m^o \setminus \Omega_m^o$  be its boundary. We assume that  $\{\partial\Omega_m^*\}_{m \in \mathbb{F}}$  are piecewise Lipschitz and denote by  $\partial\Omega_m$  the restriction of  $\partial\Omega_m^*$  to non-corner points. Let  $\Omega^o := \cup_{m \in \mathbb{F}} \Omega_m^o = \{(m, \mathbf{x}) : m \in \mathbb{F}, \mathbf{x} \in \Omega_m^o\}$  and  $\partial\Omega := \cup_{m \in \mathbb{F}} \partial\Omega_m$ . For a point  $(m, \mathbf{x}) \in \partial\Omega_m$ , let  $n(m, \mathbf{x})$  be the outward unit normal, and

let  $\Gamma^\pm(m, \mathbf{x}) := \{\mathbf{v} \in \{-1, 1\}^d : \pm(\mathbf{v} \cdot \mathbf{n}(m, \mathbf{x})) > 0\}$  be the sets of velocities with which a zig-zag process can exit (+) or enter (-)  $\Omega_m^o$  at  $\mathbf{x}$ . Zig-zag dynamics imply  $\mathbf{v} \in \Gamma^+(m, \mathbf{x}) \Leftrightarrow -\mathbf{v} \in \Gamma^-(m, \mathbf{x})$ . We also define  $\Gamma^\pm(\partial\Omega) := \cup_{(m, \mathbf{x}) \in \partial\Omega} (\{(m, \mathbf{x})\} \times \Gamma^\pm(m, \mathbf{x}))$  and  $\Omega^* := (\Omega^o \times \{-1, 1\}^d) \cup \Gamma^-(\partial\Omega)$ . Integrals over  $\Omega^o$  and  $\partial\Omega$ , or subsets thereof, are taken to incorporate discrete sums in the  $m \in \mathbb{F}$  coordinate.

The zig-zag process  $(m_t, \mathbf{x}_t, \mathbf{v}_t)_{t \geq 0}$  is defined on  $\Omega^*$ , with target  $\tilde{\pi}(m, \mathbf{x}, \mathbf{v}) := \pi(m, \mathbf{x})/2^d$  on  $\Omega^* \cup \Gamma^+(\partial\Omega)$  for a given density  $\pi(m, \mathbf{x})$ . At  $(m, \mathbf{x}, \mathbf{v}) \in \Omega^*$ , the process moves with velocity  $\mathbf{v}$  and each coordinate  $v_i$  flips at rate  $\lambda(m, \mathbf{x}, \mathbf{v})$ , defined as in (1) since  $m$  is fixed between boundary hitting events. When  $(m, \mathbf{x}, \mathbf{v}) \in \Gamma^+(\partial\Omega)$ , the process jumps according to a Markov kernel  $Q : \Gamma^+(\partial\Omega) \mapsto \mathcal{P}(\Gamma^-(\partial\Omega))$ , where  $\mathcal{P}(A)$  denotes the set of probability measures on  $(A, \mathcal{B}(A))$ . We assume that  $Q$  and  $\tilde{\pi}$  satisfy the skew-detailed balance condition

$$\tilde{\pi}(m, \mathbf{x}, \mathbf{v})Q(m, \mathbf{x}, \mathbf{v}; j, d\mathbf{y}, d\mathbf{w}) = \tilde{\pi}(j, \mathbf{y}, -\mathbf{w})Q(j, \mathbf{y}, -\mathbf{w}; m, d\mathbf{x}, -d\mathbf{v}) \quad (2)$$

for any  $(m, \mathbf{x}, \mathbf{v}) \in \Gamma^+(\partial\Omega)$  and  $(j, \mathbf{y}, \mathbf{w}) \in \Gamma^-(\partial\Omega)$ , as well as

$$\int_{(j, \mathbf{y}) \in \partial\Omega} \sum_{\mathbf{w} \in \Gamma^-(j, \mathbf{y})} (\mathbf{w} \cdot \mathbf{n}(j, \mathbf{y}))Q(m, \mathbf{x}, \mathbf{v}; j, d\mathbf{y}, d\mathbf{w}) = -\mathbf{v} \cdot \mathbf{n}(m, \mathbf{x}), \quad (3)$$

for any  $(\mathbf{x}, \mathbf{v}) \in \Gamma^+(\partial\Omega)$ , and exclude jumps to paths pointing into corners by assuming

$$\int_{(m, \mathbf{x}) \in \partial\Omega} \sum_{\mathbf{v} \in \Gamma^+(m, \mathbf{x})} \int_{(j, \mathbf{y}) \in \partial\Omega} \sum_{\mathbf{w} \in \Gamma^-(j, \mathbf{y})} \mathbb{1}_{(\cup_{m \in \mathbb{F}} \partial\Omega^*) \setminus \partial\Omega}(j, \mathbf{y} + T_{\Gamma^+(\partial\Omega)}(j, \mathbf{y}, \mathbf{w})\mathbf{w}) \times Q(m, \mathbf{x}, \mathbf{v}; j, d\mathbf{y}, d\mathbf{w}) d\mathbf{x} = 0, \quad (4)$$

where  $T_{\Gamma^+(\partial\Omega)}(j, \mathbf{y}, \mathbf{w})$  is the time the line  $(j, \mathbf{y} + t\mathbf{w})_{t \geq 0}$  hits  $\Gamma^+(\partial\Omega)$ . We will abuse notation and use  $\tilde{\pi}(m, \mathbf{x}, \mathbf{v})d\mathbf{x}$  as the target density and Lebesgue measure on  $\Omega^o$ , and as their restrictions to the surface  $\partial\Omega$ , on which  $\tilde{\pi}$  is not a probability density.

By [Dav93, Theorem 26.14], the zig-zag process with dynamics defined above is a piecewise-deterministic Markov process with extended generator

$$Lf(m, \mathbf{x}, \mathbf{v}) = \sum_{i=1}^d v_i \partial_i f(m, \mathbf{x}, \mathbf{v}) + \sum_{i=1}^d \lambda_i(m, \mathbf{x}, \mathbf{v})(f(m, \mathbf{x}, F_i \mathbf{v}) - f(m, \mathbf{x}, \mathbf{v})), \quad (5)$$

whose domain  $D(L)$  consists of measurable functions  $f(m, \mathbf{x}, \mathbf{v})$  on  $\Omega^*$  satisfying:

1. For each  $(m, \mathbf{x}, \mathbf{v}) \in \Gamma^+(\partial\Omega)$ , the limit  $\lim_{t \rightarrow 0} f(m, \mathbf{x} - t\mathbf{v}, \mathbf{v}) =: f(m, \mathbf{x}, \mathbf{v})$  exists.
2. For each  $(m, \mathbf{x}, \mathbf{v}) \in \Omega^*$ , the function  $t \mapsto f(m, \mathbf{x} + t\mathbf{v}, \mathbf{v})$  is absolutely continuous on  $t \in [0, T_{\Gamma^+(\partial\Omega)}(m, \mathbf{x}, \mathbf{v})]$ .
3. For  $(m, \mathbf{x}, \mathbf{v}) \in \Gamma^+(\partial\Omega)$ ,

$$f(m, \mathbf{x}, \mathbf{v}) = \int_{(j, \mathbf{y}) \in \partial\Omega} \sum_{\mathbf{w} \in \Gamma^-(j, \mathbf{y})} f(j, \mathbf{y}, \mathbf{w})Q(m, \mathbf{x}, \mathbf{v}; j, d\mathbf{y}, d\mathbf{w}). \quad (6)$$

4. The random variable  $\sum_{k=1}^n f(m_{T_k}, \mathbf{x}_{T_k}, \mathbf{v}_{T_k}) - f(m_{T_k}, \mathbf{x}_{T_k-}, \mathbf{v}_{T_k-})$  is integrable for each  $n \in \mathbb{N}$ , where  $\{T_k\}_{k \geq 1}$  are the successive jump times (both velocity flips and jumps due to hitting a boundary) of  $(m_t, \mathbf{x}_t, \mathbf{v}_t)_{t \geq 0}$ .

For a set  $A$ , let  $B(A)$  and  $C^1(A)$  be the respective spaces of bounded and continuously differentiable functions on  $A$ . Let  $C_b^1(A) := B(A) \cap C^1(A)$ . For  $t > 0$ , we define

$$N_t := \#\{\text{velocity flips and boundary jumps in } (m_s, \mathbf{x}_s, \mathbf{v}_s)_{s=0}^t\}.$$

**Theorem 1.** Suppose  $\mathbb{F}$  is finite,  $\tilde{\pi}(m, \cdot, \mathbf{v}) \in C^1(\Omega_m^o)$  for each  $\mathbf{v} \in \{-1, 1\}^d$  and  $m \in \mathbb{F}$ ,  $\tilde{\pi} > 0$  on  $\Omega^o \times \{-1, 1\}^d$ , that  $Q(m, \mathbf{x}, \mathbf{v}; \cdot)$  has compact support for each  $(m, \mathbf{x}, \mathbf{v}) \in \Gamma^+(\partial\Omega)$ , and for each  $t > 0$  and  $(m, \mathbf{x}, \mathbf{v}) \in \Omega^*$ ,

$$\mathbb{E}[N_t | (m_0, \mathbf{x}_0, \mathbf{v}_0) = (m, \mathbf{x}, \mathbf{v})] < \infty. \quad (7)$$

Suppose the initial distribution of  $(m, \mathbf{x})$  has a density on  $\Omega^*$  and that (2), (3), and (4) hold. Then the zig-zag process with generator (5) and domain  $D(L)$  as described above has stationary distribution  $\tilde{\pi}$ .

*Proof.* Provided in the appendix. □

**Remark 1.** In addition to having the right invariant distribution, a practical algorithm needs to be ergodic. To discuss ergodicity of the zig-zag process  $(m_t, \mathbf{x}_t, \mathbf{v}_t)_{t \geq 0}$  from Theorem 1, let  $\{(m_t^j, \mathbf{x}_t^j, \mathbf{v}_t^j)_{t \geq 0}\}_{j \in \mathbb{F}}$  be zig-zag processes restricted to respective spaces  $\Omega_j^o$  by boundary jump kernels  $Q^j : \cup_{\mathbf{x} \in \partial\Omega_j} [(j, \mathbf{x}) \times \Gamma^+(j, \mathbf{x})] \mapsto \mathcal{P}(\cup_{\mathbf{x} \in \partial\Omega_j} [(j, \mathbf{x}) \times \Gamma^-(j, \mathbf{x})])$ , each with target proportional to  $\tilde{\pi}(j, \cdot, \cdot)$ . When  $\mathbb{F}$  is finite, a sufficient condition for ergodicity of the global process  $(m_t, \mathbf{x}_t, \mathbf{v}_t)_{t \geq 0}$  is that  $\{(m_t^j, \mathbf{x}_t^j, \mathbf{v}_t^j)_{t \geq 0}\}_{j \in \mathbb{F}}$  are all ergodic, and that

$$\int_{\mathbf{x} \in \partial\Omega_m} \sum_{\mathbf{v} \in \Gamma^+(m, \mathbf{x})} \int_{\mathbf{y} \in \partial\Omega_j} \sum_{\mathbf{w} \in \Gamma^-(j, \mathbf{y})} Q(m, \mathbf{x}, \mathbf{v}; j, \mathbf{y}, \mathbf{w}) \tilde{\pi}(m, \mathbf{x}, \mathbf{v}) d\mathbf{x} > 0, \quad (8)$$

for ordered pairs  $(m, j) \in \mathbb{F}^2$  which form a cycle spanning the support of  $\tilde{\pi}$ . [BRZ19] provide conditions for ergodicity of single-domain zig-zag processes.

We conclude this section with a pseudocode specification of our zig-zag algorithm.

---

**Algorithm 1** Simulation the zig-zag process targeting  $\tilde{\pi}$

---

**Require:** Initial condition  $(m, \mathbf{x}_0, \mathbf{v}_0)$ , target  $\tilde{\pi}$ , jump kernel  $Q$ , terminal time  $t_{\text{end}}$

- 1: Set  $t \leftarrow 0, m_0 \leftarrow m, \mathbf{x} \leftarrow \mathbf{x}_0, \mathbf{v} \leftarrow \mathbf{v}_0$
  - 2: **while**  $t < t_{\text{end}}$  **do**
  - 3:     Set  $\tau \leftarrow T_{\Gamma^+(\partial\Omega)}(m, \mathbf{x}, \mathbf{v})$  and  $I \leftarrow 0$ .  $\triangleright I = 0 \Rightarrow$  boundary hit.
  - 4:     **for**  $i \in \{1, 2, \dots, d-1, d\}$  **do**
  - 5:         Sample  $U \sim \text{Exp}(1)$ .
  - 6:         Set  $\rho$  to be the solution of  $\int_0^\rho \lambda_i(m, \mathbf{x} + s\mathbf{v}, \mathbf{v}) ds = U$ .
  - 7:         **if**  $\rho < \tau$  **then**
  - 8:             Set  $\tau \leftarrow \rho$  and  $I \leftarrow i$
  - 9:     Set  $t \leftarrow t + \tau, \mathbf{x} \leftarrow \mathbf{x} + \tau\mathbf{v}$
  - 10:     **if**  $I > 0$  **then**
  - 11:         Set  $v_I \leftarrow -v_I$
  - 12:     **else**
  - 13:         Set  $(m, \mathbf{x}, \mathbf{v}) \leftarrow (j, \mathbf{y}, \mathbf{w}) \sim Q(m, \mathbf{x}, \mathbf{v}; \cdot, \cdot, \cdot)$
  - 14: **end while**
- 

### 3 The coalescent and a geometric embedding

An ultrametric binary tree with  $n$  labeled leaves is a rooted binary tree in which each leaf is an equal graph distance away from the root. We follow [GD16] and encode such a tree with leaf labels  $\{1, \dots, n\}$  as the pair  $(E_n, \mathbf{t}_n)$ , where  $E_n$  is a ranked topology and  $\mathbf{t}_n \in (0, \infty)^{n-1}$ . The continuous variables  $\mathbf{t}_n$  encode times between mergers. The time from the leaves to the first merger is  $t_1$ , and subsequent  $t_i$  variables are times between successive mergers. The ranked topology  $E_n$  is an  $(n-1)$ -tuple of pairs of labels, where the  $i$ th pair specifies the two child nodes

of the  $i$ th merger. Non-leaf nodes are labeled by the leaves they subtend. For example, the ranked topology

$$E_4 = (E_{4,1}, E_{4,2}, E_{4,3}) := (\{1, 2\}, \{\{1, 2\}, 3\}, \{\{1, 2, 3\}, 4\})$$

encodes the four leaf caterpillar tree depicted in Figures 3 and 7 with nodes labeled left to right. We order the two entries of  $E_{n,i}$  by their least element for definiteness.

The coalescent [Kin82] is a seminal model for the genetic ancestry of samples from large populations. Under the coalescent, a tree  $(E_n, \mathbf{t}_n)$  has probability density

$$\pi(E_n, \mathbf{t}_n) d\mathbf{t}_n := \exp\left(-\sum_{i=1}^{n-1} \binom{n+1-i}{2} t_i\right) d\mathbf{t}_n, \quad (9)$$

which arises as the law of a tree constructed by starting a lineage from each leaf, and merging each pair of lineages at rate 1 until the most recent common ancestor (MRCA) is reached. The success of the coalescent is due to robustness: distributions of ancestries of a large class of individual-based models converge to the coalescent in the infinite population limit under suitable rescalings of time. For details, see e.g. [Wak09].

We define swap operators  $s_i$  for  $i \in \{2, \dots, n-2\} : E_{n,i-1} \notin E_{n,i}$  via

$$s_i(E_n) := (E'_{n,1}, \dots, E'_{n,n-1}) \text{ where } E'_{n,k} := \begin{cases} E_{n,k} & \text{for } k < i-1 \text{ or } k > i, \\ E_{n,i} & \text{for } k = i-1, \\ E_{n,i-1} & \text{for } k = i. \end{cases}$$

In words,  $s_i$  swaps the order of the  $(i-1)$ th and  $i$ th mergers. We also define pivot operators  $p_i^\downarrow$  and  $p_i^\uparrow$  for  $i \in \{2, \dots, n-1\} : E_{n,i-1} \in E_{n,i}$  as

$$p_i^\downarrow(E_n) := (E'_{n,1}, \dots, E'_{n,n-1}) \text{ with } E'_{n,k} := \begin{cases} E_{n,k} & \text{for } k \notin \{i-1, i\}, \\ \{E_{n,i-1}^\downarrow, E_{n,i}^s\} & \text{for } k = i-1, \\ \{E'_{n,i-1}, E_{n,i-1}^\uparrow\} & \text{for } k = i, \end{cases} \quad (10)$$

where  $E_{n,i}^\uparrow$  (resp.  $E_{n,i}^\downarrow$ ) is the entry of  $E_{n,i}$  with the higher (resp. lower) least element, and  $E_{n,i}^s$  is the *sibling*: the entry of  $E_{n,i}$  that is not  $E_{n,i-1}$ . Pivot  $p_i^\uparrow$  is defined by interchanging  $\downarrow$  and  $\uparrow$  in (10). The pivots are the two nearest neighbor interchanges between the  $i$ th merger and the merger involving its nearest child. Figure 1 illustrates all three operators.

Next we describe  $\tau$ -space, which gives a geometric structure to the set of pairs  $(E_n, \mathbf{t}_n)$ . For fixed  $E_n$ , the space of  $\mathbf{t}_n$ -vectors is the orthant  $[0, \infty)^{n-1}$ . Each boundary point with  $t_i = 0$  corresponds to a degenerate tree in which one of three things happens:

1. The two leaves of  $E_{n,1}$  merge at time 0 if  $i = 1$ .
2. There are two simultaneous mergers if  $E_{n,i-1} \notin E_{n,i}$ .
3. There is a simultaneous merger of three lineages if  $E_{n,i-1} \in E_{n,i}$ .

Type 1 boundaries are boundaries of the whole  $\tau$ -space. Type 2 boundaries separate orthants corresponding to two ranked topologies separated by an  $s_i$ -step. Trajectories crossing the boundary move from one ranked topology to the other. Type 3 boundaries separate the three orthants which resolve the triple merger into two binary mergers, which differ by a  $p_i^\downarrow$  or  $p_i^\uparrow$ -step. A trajectory that crosses the boundary visits a tree with the triple merger. Figure 2 depicts the  $\tau$ -space with three leaves, and a type 2 boundary in a general  $\tau$ -space. An example with five leaves is depicted in [GD16, Figure 2], and the  $\mathbf{t}_n$  variables are illustrated in Figures 3 and 7 in Sections 4 and 5.




Figure 1: Operators  $p_2^\uparrow$ ,  $p_2^\downarrow$ , and  $s_2$ . The horizontal arrangement of leaves is arbitrary throughout this paper; only vertical distance is meaningful.

We will use  $\tau$ -space to construct zig-zag processes whose state spaces consist of tree topologies, branch lengths, and a scalar parameter introduced in the next section. The discrete variables  $\mathbb{F}$  will be the ranked topologies, with the boundary crossings described above defining the boundary jump kernel  $Q$ . For more on  $\tau$ -space, e.g. existence and uniqueness of geodesics and Fréchet means, we refer to [GD16].

## 4 The infinite sites model

The infinite sites model [Wat75] connects the coalescent tree to DNA sequence data by associating the MRCA with the unit interval  $(0, 1)$ . Mutations with independent,  $U(0, 1)$ -distributed locations accrue along branches of the tree (with branch lengths as specified by  $\mathbf{t}_n$ ) at rate  $\theta/2$ . The type of a leaf consist of the mutations along branches separating it from the MRCA. We denote the resulting list of types of the leaves by  $D_n$ . A realization of a coalescent tree and associated  $D_n$  is shown in Figure 3. The typical task is to sample from the conditional law  $(E_n, \mathbf{t}_n, \theta) | D_n$  corresponding to observing  $D_n$ , but not the tree which gave rise to it.

To handle mutations, we define  $F_n$  as the rooted graphical tree with  $2n - 1$  nodes, the first  $n$  of which are leaves labeled  $1, \dots, n$ , while the remaining  $n - 1$  are labeled as in  $E_n$ . Edges connect children to their parents, and edge lengths are determined by  $\mathbf{t}_n$ . For an edge  $\gamma \in F_n$ , we denote by  $c_\gamma$  and  $p_\gamma$  the respective labels of the child and parent nodes of  $\gamma$ , by  $m_\gamma$  the number of mutations on  $\gamma$ , and by  $l_\gamma := \sum_{t_j \in \gamma} t_j$  the edge length, where we write  $t_i \in \gamma$  if  $t_i$  contributes to the length of  $\gamma$  in which case we say  $\gamma$  spans  $t_i$ .

Given a prior density  $\pi_0(\theta)$  for  $\theta$ , the posterior distribution of  $(E_n, \mathbf{t}_n, \theta) | D_n$  follows from (9), the fact that the number of mutations on branch  $\gamma \in F_n$  is  $\text{Poisson}(\theta l_\gamma / 2)$ -distributed given  $l_\gamma$ ,




Figure 2: (Left)  $\tau$ -space with  $n = 3$  embedded into  $\mathbb{R}^3$ . Each square is a copy of  $[0, \infty)^2$  associated with the given topology. The coordinates  $(t_1, t_2)$  are the respective time of the first merger, and the time between the first and second merger. The dot is the origin, and the line on which all three orthants intersect is a type 3 boundary consisting of trees in which all three leaves merge simultaneously at time  $t_1$ . The dashed lines are boundaries at  $\infty$ . (Right) A segment of  $\tau$ -space depicting a type 2 boundary, in which each square represents  $[0, \infty)^{n-1}$ . Only the two orthants adjacent to the boundary are shown.




Figure 3: A realization of the infinite sites model with  $n = 4$ , two mutations, three types, and  $D_n = (\{0.7\}, \{0.7\}, \{\}, \{0.2\})$ . The holding times  $t_3$  are shown on the left.

and that mutations on distinct branches are independent. In particular,

$$\pi(E_n, \mathbf{t}_n, \theta | D_n) \propto \left\{ \prod_{\gamma \in F_n} \left( \frac{\theta l_\gamma}{2} \right)^{m_\gamma} \right\} \exp \left( - \sum_{i=1}^{n-1} \frac{(n+1-i)(n+\theta-i)}{2} t_i \right) \pi_0(\theta) \quad (11)$$

provided  $E_n$  is consistent  $D_n$ , and  $\pi = 0$  otherwise. This distribution can be sampled using a zig-zag algorithm by taking  $\mathbb{F}$  to be the set of ranked topologies on  $n$  leaves which are consistent with  $D_n$ , as well as  $\Omega_{E_n}^0 := \{(\mathbf{t}_n, \theta) \in (0, \infty)^n\}$  and  $\partial\Omega_{E_n} := \{(\mathbf{t}_n, \theta) \in [0, \infty)^n : t_i = 0 \text{ for one } i \in \{1, \dots, n-1\} \text{ or } \theta = 0\}$  for each  $E_n \in \mathbb{F}$ . In the boundary classification of Section 3,  $\theta = 0$  is another type 1 boundary. For  $(\mathbf{t}_n, \theta) \in \partial\Omega$ , we define  $k(\mathbf{t}_n, \theta)$  as the index of the zero variable, taken to be  $n$  in the case of  $\theta$ .

We augment the state space with  $n$  zig-zag velocities  $\mathbf{v}_n$ , of which  $(v_1, \dots, v_{n-1})$  drive  $\mathbf{t}_n$  and  $v_n$  drives  $\theta$ . For  $\gamma \in F_n$ , we also define  $v_\gamma := \sum_{j: t_j \in \gamma} v_j$  as the rate of change of  $l_\gamma$ . The boundary kernel  $Q$  is defined separately on each boundary type:

$$Q(E_n, (\mathbf{t}_n, \theta), \mathbf{v}_n; \cdot, \cdot, \cdot) := \begin{cases} \delta_{\{E_n\}}(\cdot) \otimes \delta_{\{(\mathbf{t}_n, \theta)\}}(\cdot) \otimes \delta_{\{F_{k(\mathbf{t}_n, \theta)}(\mathbf{v}_n)\}}(\cdot) & \text{for type 1,} \\ \delta_{\{s_{k(\mathbf{t}_n, \theta)}(E_n)\}}(\cdot) \otimes \delta_{\{(\mathbf{t}_n, \theta)\}}(\cdot) \otimes \delta_{\{F_{k(\mathbf{t}_n, \theta)}(\mathbf{v}_n)\}}(\cdot) & \text{for type 2,} \\ \left( \frac{\delta_{\{p_k^\uparrow(\mathbf{t}_n, \theta)(E_n)\}}(\cdot)}{2} + \frac{\delta_{\{p_k^\downarrow(\mathbf{t}_n, \theta)(E_n)\}}(\cdot)}{2} \right) \otimes \delta_{\{(\mathbf{t}_n, \theta)\}}(\cdot) \otimes \delta_{\{F_{k(\mathbf{t}_n, \theta)}(\mathbf{v}_n)\}}(\cdot) & \text{for type 3.} \end{cases} \quad (12)$$



At a type 1 boundary the process reflects back into  $\Omega_{E_n}^o$  via a velocity flip. At a type 2 boundary it undergoes an  $s_i$ -step and a velocity flip to pass through the boundary. For type 3 boundaries it chooses an adjacent orthant uniformly at random.

In the interiors of orthants, velocity flip rates are

$$\lambda_i(E_n, \mathbf{t}_n, \theta; \mathbf{v}_n) := \left[ v_i \left( \frac{(n+1-i)(n+\theta-i)}{2} - \sum_{\gamma \in F_n: t_i \in \gamma} \frac{m_\gamma}{l_\gamma} \right) \right]^+, \quad (13)$$

$$\lambda_\theta(E_n, \mathbf{t}_n, \theta; \mathbf{v}_n) := \left[ v_n \left( \sum_{i=1}^{n-1} \frac{n+1-i}{2} t_i - \frac{1}{\theta} \sum_{\gamma \in F_n} m_\gamma - \partial_\theta \log(\pi_0(\theta)) \right) \right]^+. \quad (14)$$

Simulating holding times with these rates is difficult due to the time intervals during which they vanish. One strategy is Poisson thinning via dominating rates consisting of only those terms in the round brackets in (13) and (14) whose sign matches that of the corresponding velocity  $v_i$ , but these can result in loose bounds and inefficient algorithms. Instead, we define  $t_n := \theta$  for brevity, and for  $i \in \{1, \dots, n-1\}$  define  $\gamma(E_n, i) := \operatorname{argmin}\{l_\gamma : p_\gamma = E_{n,i}\}$  as the shorter child branch from parent node  $E_{n,i}$ . For  $i \in \{1, n\}$  we adopt the short-hands

$$m_{\gamma(E_n, n)} := \sum_{\gamma \in F_n} m_\gamma \quad \text{and} \quad m_{\gamma(E_n, 1)} := \sum_{\gamma \in F_n: p_\gamma = E_{n,1}} m_\gamma,$$

and finally define the time localization  $T \equiv T(E_n, \mathbf{t}_n, \theta; \mathbf{v}_n)$  as

$$T := \min \left\{ \min_{i \in \{1, \dots, n\}: v_i < 0} \left\{ \frac{-t_i}{\{1 + c[\mathbb{1}_{E_{n,i}}(E_{n,i-1}) + \mathbb{1}_{\{1, n\}}(i)] \mathbb{1}_{\mathbb{Z}_+}(m_{\gamma(E_n, i)})\}} v_i \right\}, K \right\}, \quad (15)$$

where  $\min_\emptyset = \infty$ , and  $K \gg 0$  is a maximum increment for the case when all velocities are positive. The indicator functions in the denominator pick out boundaries where (11) vanishes: type 1 or 3 boundaries corresponding to length 0 branches which carry at least one mutation, and the  $\theta = 0$  boundary if there is at least one mutation in total. The parameter  $c > 0$  ensures that, when the current process time is  $t$ , such boundaries cannot be hit on  $[t, t+T]$ , and at most one other boundary can be reached. We found  $c = 4$  gave good performance across our tests in Sections 4 and 5. A larger value results in tighter bounds on (13) and (14), but wastes more computation as  $t+T$  is hit more often before an accepted velocity flip.

On time interval  $[t, t+T]$ , flip rates (13) and (14) are bounded above by constant rates

$$\lambda_i^* := \left[ v_i \left( \frac{(n+1-i)(n+\theta + (v_n T)^\pm - i)}{2} - \sum_{\gamma \in F_n: t_i \in \gamma} \frac{m_\gamma}{l_\gamma + (v_\gamma T)^\pm} \right) \right]^+,$$

$$\lambda_\theta^* := \left[ v_\theta \left( \sum_{i=1}^{n-1} \frac{n+1-i}{2} [t_i + (v_i T)^\pm] - \frac{1}{\theta + v_n T} \sum_{\gamma \in F_n} m_\gamma - \inf_{s \in [0, T]} \{\partial_\theta \log(\pi_0(\theta + v_\theta s))\} \right) \right]^+,$$

where, for each  $\lambda_i^*$ ,  $i \in \{1, \dots, n-1, \theta\}$ ,  $(x)^\pm := (x)^+$  if  $v_i > 0$  and  $(x)^\pm := (x)^- := \min\{x, 0\}$  if  $v_i < 0$ . Algorithms 2 and 3 in the appendix give pseudocode for simulating holding times, velocity flips, and boundary crossings as outlined above.

**Proposition 1.** *Suppose that the initial condition  $(E_n, \mathbf{t}_n, \theta)$  has a positive density on  $\mathbb{F} \times (0, \infty)^n$ , that  $\pi_0 > 0$ , and that  $\partial_\theta \log(\pi_0(\theta))$  is bounded on compact subsets of  $(0, \infty)$ . Then (11) is stationary under the dynamics simulated by Algorithm 2 and 3.*

*Proof.* Provided in the appendix. □

We compared the zig-zag process to a Metropolis–Hastings algorithm by reanalyzing the data of [WFDP91] with  $n = 55$ , 14 distinct types, and 18 mutations. We used the improper prior  $\pi_0(\theta) \propto \mathbf{1}_{(0,\infty)}(\theta)$ , set  $v_i = \pm 2/[(n+1-i)(n-i)]$ , and set  $v_n$  from trial runs to cross the  $\theta$ -mode in unit time. The appendix details the Metropolis–Hastings algorithm, and other tuning parameters in Table 3. We compared both methods to a hybrid combining zig-zag dynamics with continuous time Metropolis–Hastings moves at rate  $\kappa = 10$ . Performance was insensitive to  $\kappa$  provided it was not extreme: small values resemble a zig-zag process, while large values resemble Metropolis–Hastings.

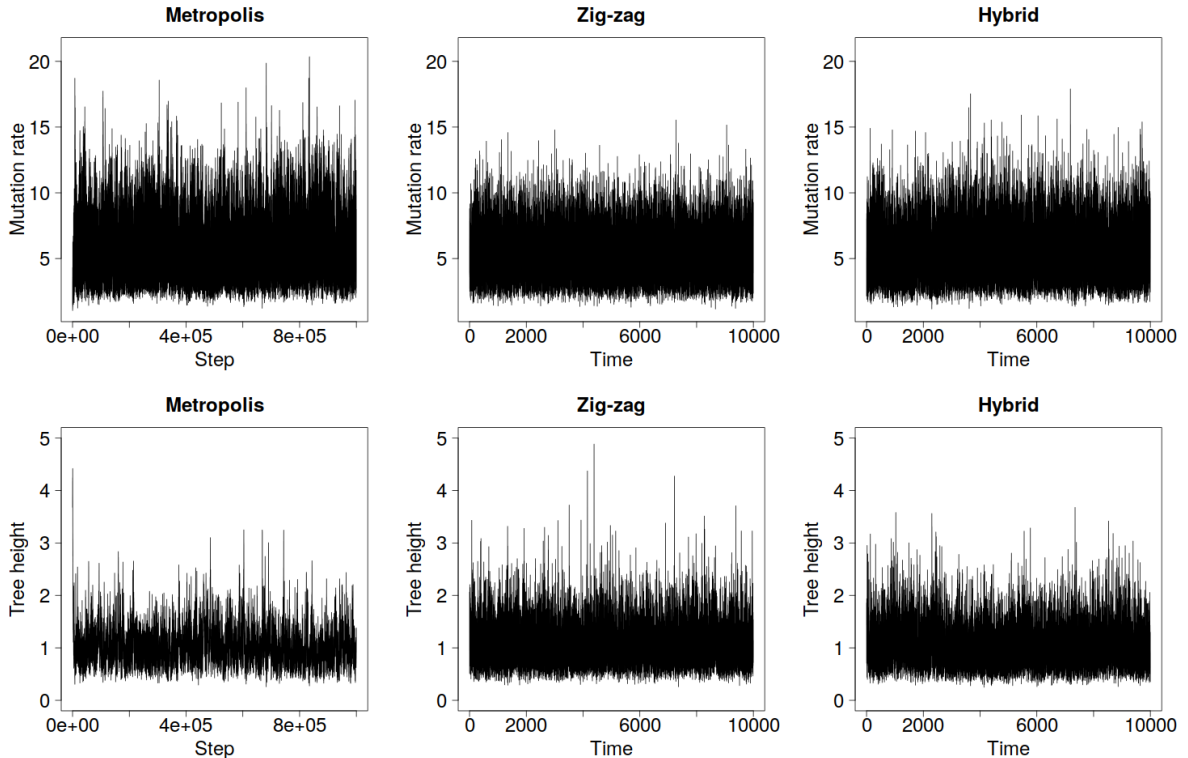


Figure 4: Trace plots under the infinite sites model and the data set of [WFDP91].

Figure 4 shows that the zig-zag and hybrid methods mix visibly better than Metropolis–Hastings over the latent tree, as measured by the tree height  $H_n := t_1 + \dots + t_{n-1}$ . However, they are not as effective at exploring the upper tail of the  $\theta$ -marginal, likely because they do not stay in regions of short trees for long enough for  $\theta$  to increase into the tail.

To assess scaling, we simulated two data sets: one of size  $n = 550$  with mutation rate  $\theta = 5.5$  (the approximate posterior mean in Figure 4) and one with  $n = 55$  and  $\theta = 55$ , which models a segment of DNA 10 times longer. Figures 5 and 6 demonstrate that the zig-zag and hybrid processes scale far better than Metropolis–Hastings, particularly when  $\theta = 55$ . Estimates in Table 1 quantify the improvement to 1–3 orders of magnitude.

## 5 The finite sites model

The finite sites model [JC69] is more detailed than the infinite sites model, but has greater computational cost. Consider a finite set of sites  $S$  with a finite number of possible types  $H$  per site; for example  $H = \{0, 1\}$  or  $H = \{A, T, C, G\}$ . Mutations occur along branches of the coalescent tree at each site with rate  $\theta/(2|S|)$ , and the type of a mutant child is drawn from stochastic matrix  $P$  with unique stationary distribution  $\nu$ . We denote the transition matrix of the  $H$ -valued compound Poisson process with jump rate  $\theta/(2|S|)$  and jump transition matrix  $P$

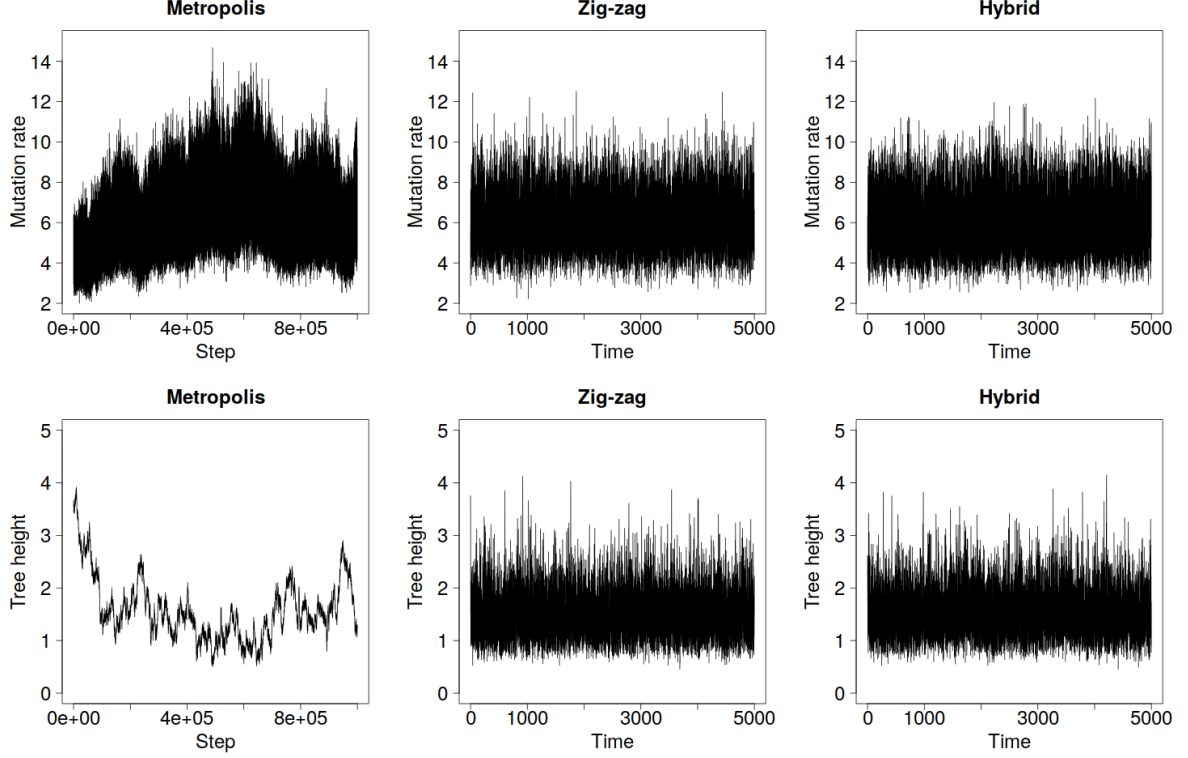


Figure 5: Trace plots for the infinite sites model and the data set with  $n = 550$ ,  $\theta = 5.5$ , 30 distinct types, and 38 mutations.

by  $(Q_{hg}^\theta(t))_{h,g \in H; t \geq 0}$ . Figure 7 depicts a realization of the finite sites coalescent.

As in Section 4, we denote the configuration of types at the leaves by  $D_n$ , and seek to sample from the posterior  $\pi(E_n, \mathbf{t}_n, \theta | D_n)$ , which can be written as a sum over the types of internal nodes:

$$\pi(E_n, \mathbf{t}_n, \theta | D_n) \propto \left\{ \prod_{s \in S} \sum_{\substack{h(s, c_\gamma) \in H \\ \text{for } \gamma \in F_n: |c_\gamma| > 1}} \prod_{\gamma \in F_n} Q_{h(s; p_\gamma)h(s; c_\gamma)}^\theta(l_\gamma) \right\} \times \exp \left( - \sum_{i=1}^{n-1} \binom{n+1-i}{2} t_i \right) \pi_0(\theta), \quad (16)$$

where  $h(s; \eta) \in H$  is the type at site  $s \in S$  on the node with label  $\eta \in E_n$ , and  $\gamma \in F_n : |c_\gamma| > 1$  denotes edges which do not end in a leaf. The target (16) can be evaluated efficiently using the pruning algorithm of [Fel81].

The posterior (16) can be sampled using zig-zag dynamics with the same construction as in



Data set	Method	ESS( $\theta$ )/sec	ESS( $H_n$ )/sec	Run time (min)
[WFDP91]	Metropolis	5	3	3.5
	Zig-zag	160	179	0.5
	Hybrid	128	67	0.5
$n = 550, \theta = 5.5$	Metropolis	0.1*	0.002*	70
	Zig-zag	1.7	1.2	45
	Hybrid	0.5	0.3	115
$n = 55, \theta = 55$	Metropolis	0.1*	0.1*	50
	Zig-zag	36	37	1
	Hybrid	22	22	1.5

Table 1: Effective sample sizes and run times for all three methods and data sets for the infinite sites model. Estimates were computed with the `ess` method [FHVD17] under default settings for Metropolis–Hastings, and as in [BFR19a, Section 2] for the other two. Stars highlight poorly mixing chains with unreliable ESS estimates.

which can be evaluated using the linear-cost method of [JZH<sup>+</sup>20].

We show that events with rates (17) and (18) can be simulated using the example of [GT94, Section 7.4], in which  $S = \{1, \dots, 20\}$ ,  $H = \{0, 1\}$  and  $P$  is the  $2 \times 2$  matrix which always changes state, corresponding to  $\nu = (1/2, 1/2)$  and

$$Q_{hg}^\theta(t) := \frac{1}{2} + \left( \mathbb{1}_{\{h\}}(g) - \frac{1}{2} \right) e^{-\theta t}. \quad (19)$$

As (19) is not bounded away from 0 when  $h \neq g$ , (17) and (18) lack simple bounds for Poisson thinning. As in (15), bounds can be obtained by time localization using

$$T \equiv T(E_n, \mathbf{t}_n, \theta; \mathbf{v}_n) := \min \left\{ \min_{i \in \{1, \dots, n\}: v_i < 0} \left\{ \frac{-t_i}{[1 + c \mathbb{1}_{\{1, \theta\}}(i) \mathbb{1}_{\mathbb{Z}_+}(m_\gamma(E_n, i))] v_i} \right\}, K \right\},$$

where  $K \gg 0$  is a default increment in case all velocities are positive. The variable  $T$  localizes the next zig-zag time step beginning at time  $t$  so that at most one branch can shrink to length zero on  $[t, t + T]$ ,  $\theta$  can fall by at most  $1/(1 + c)$  of its present value, and  $t_1$  can fall by at most  $1/(1 + c)$  of its present value if the first two leaves to merge are of distinct types. This treatment of  $\theta$  and  $t_1$  is needed as (17) and (18) diverge in these cases, rendering the  $\theta = 0$  and  $t_1 = 0$  boundaries inaccessible.

Given  $T \in (0, \infty)$ , we have the following bounds on (19) on the time interval  $[t, t + T]$ :

$$\begin{aligned} Q_{hh}^\theta(l_\gamma) &\leq \frac{1}{2} \{1 + \exp(-[\theta + (v_n T)^-][l_\gamma + (v_\gamma T)^-])\}, \\ Q_{hg}^\theta(l_\gamma) &\leq \frac{1}{2} \{1 - \exp(-[\theta + (v_n T)^+][l_\gamma + (v_\gamma T)^+])\}, \\ Q_{hh}^\theta(l_\gamma) &\geq \frac{1}{2} \{1 + \exp(-[\theta + (v_n T)^+][l_\gamma + (v_\gamma T)^+])\}, \\ Q_{hg}^\theta(l_\gamma) &\geq \frac{1}{2} \{1 - \exp(-[\theta + (v_n T)^-][l_\gamma + (v_\gamma T)^-])\}, \end{aligned}$$

where  $h \neq g$ . Substituting these bounds into [JZH<sup>+</sup>20, Equation (9)] provides bounds on flip rates that can be evaluated with  $O(|S|n)$  cost.

Figure 8 demonstrates that the zig-zag process mixes over latent trees faster than Metropolis–Hastings again, but struggles to explore the upper tail of the  $\theta$ -marginal. The hybrid method was run with  $\kappa = 100$  to compensate for shorter run lengths than in Section 4, and thus resembles Metropolis–Hastings rather than the zig-zag process.

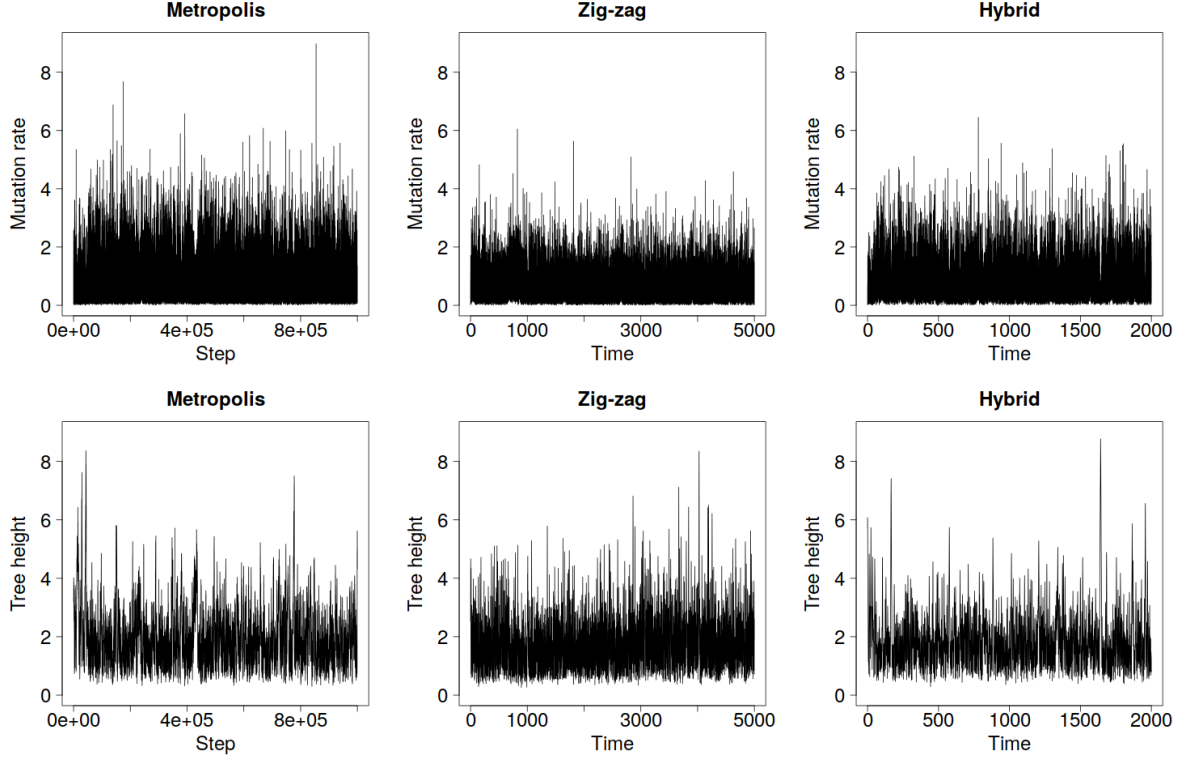


Figure 8: Trace plots for the finite sites model and data from [GT94].

Figures 9 and 10 show results for two further simulated data sets: one with  $n = 500$  and  $S = 20$ , and one with  $n = 50$  and  $S = 200$ . The superior mixing of the zig-zag process over the latent tree is clear. The lack of mixing in the upper tail of the  $\theta$ -marginal is also stark, particularly in Figure 9 where zig-zag significantly underestimates posterior variance. The estimated posterior means of all three methods coincide in all cases (results not shown).

Data set	Method	ESS( $\theta$ )/sec	ESS( $H_n$ )/sec	Run time (min)
[GT94]	Metropolis	1.8	0.2	29
	Zig-zag	2.3	1.0	18
	Hybrid	2.1	0.6	16
$n = 500, S = 20$	Metropolis	0.09	0.0006*	1567
	Zig-zag	0.01*	0.004	1749
	Hybrid	0.01	0.003	1800
$n = 50, S = 200$	Metropolis	0.04	0.03	257
	Zig-zag	0.05	0.13	175
	Hybrid	0.08	0.07	183

Table 2: Effective sample sizes and run times for all three methods and data sets. Estimates were computed as in Table 1. Stars highlight unmixed chains with unreliable ESS estimates.

## 6 Discussion

We have presented a general method for using zig-zag processes to sample targets defined on hybrid spaces consisting of discrete and continuous variables. This was done by introducing boundaries into the state space of continuous variables and updating discrete components via a Markov jump kernel  $Q$  whenever a boundary was hit. The resulting algorithm remains a piecewise-deterministic Markov processes in the sense of [Dav93, Section 24], and generalizes

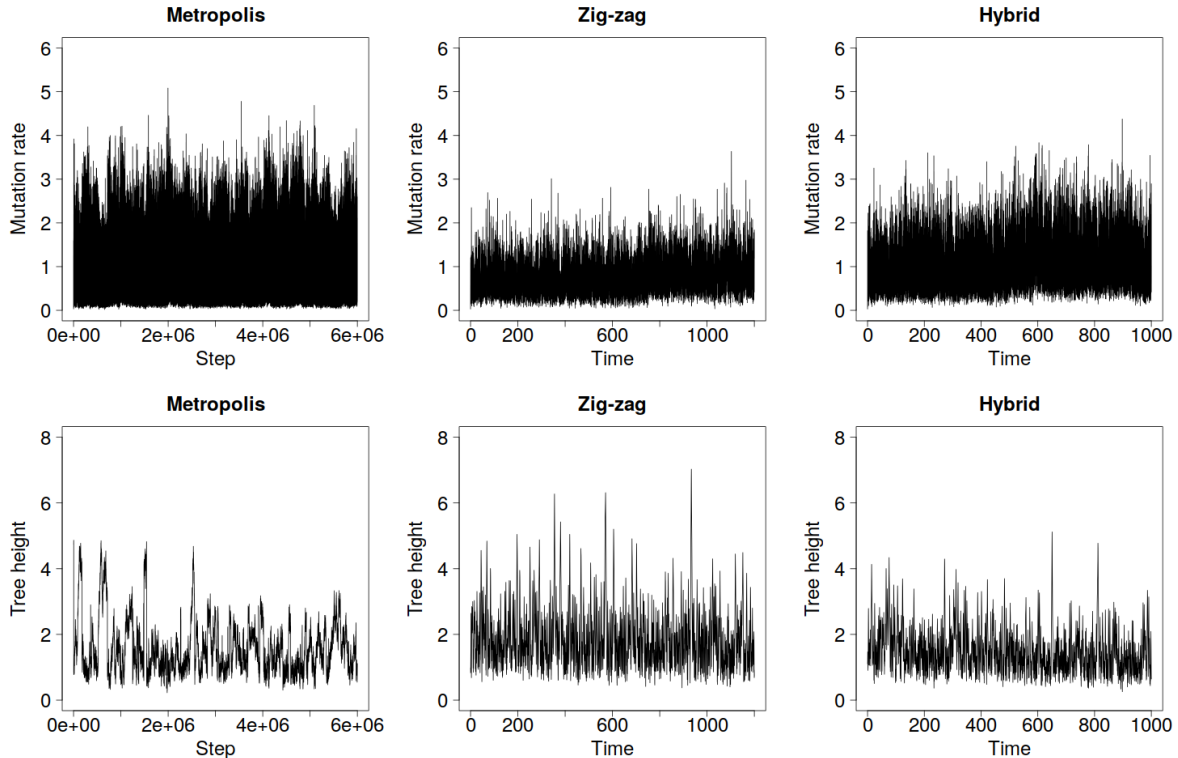


Figure 9: Trace plots for the finite sites model and data set with  $n = 500$  and  $S = 20$  consisting of five distinct sequences.

existing zig-zag processes for restricted domains [BBCD<sup>+</sup>18]. Crucially, no assumptions of structure among the discrete variables are required. The key conditions on  $Q$  are the skew-detailed balance (2), which is local, and (3), which involves an integral with respect to  $Q$  but not the target  $\tilde{\pi}$ . Both are verifiable in applications, and do not require  $\tilde{\pi}$  to be normalized. Our method is reminiscent of discrete Hamiltonian Monte Carlo [DBZM17, NDL20], but the lack of time discretization simplifies boundary crossings (though see [NDL20, Section S6.4]).

We have demonstrated the method on two examples involving the coalescent, which is a gold-standard model in phylogenetics. It is defined on the space of binary trees consisting of discrete tree topologies and continuous branch lengths, which lacks a simple geometric structure, e.g. a partial order or a tractable norm. We have also shown that the zig-zag process can improve mixing over trees relative to Metropolis-Hastings, particularly under the infinite sites model. This model is widely used to analyze ever larger data sets, and the zig-zag process shows promise for expanding the scope of feasible MCMC computations.

The zig-zag process was more efficient than Metropolis-Hastings under the infinite sites model in terms of effective sample size, but struggled to explore the tails of the  $\theta$ -marginal. A likely reason is correlation in the target: high mutation rates are only be attainable when branch lengths are short. A Metropolis-Hastings algorithm can jump to a high mutation rate as soon as the latent tree has short branches, while the zig-zag process must traverse all intervening mutation rates before branch lengths grow. The speed up zig-zag method of [VR21] has state-dependent velocities, and could provide further improvement. The hybrid method with both zig-zag motion and Metropolis-Hastings updates interpolated between the two algorithms.

All three algorithms exhibited much longer run times under the finite sites model than under infinite sites. For the zig-zag and hybrid methods, that is due to the  $O(|S|n)$  cost per evaluation of (17) and (18), of which there are  $O(n)$ . However, flip times for different velocities are conditionally independent given the current state and can be generated in parallel, unlike

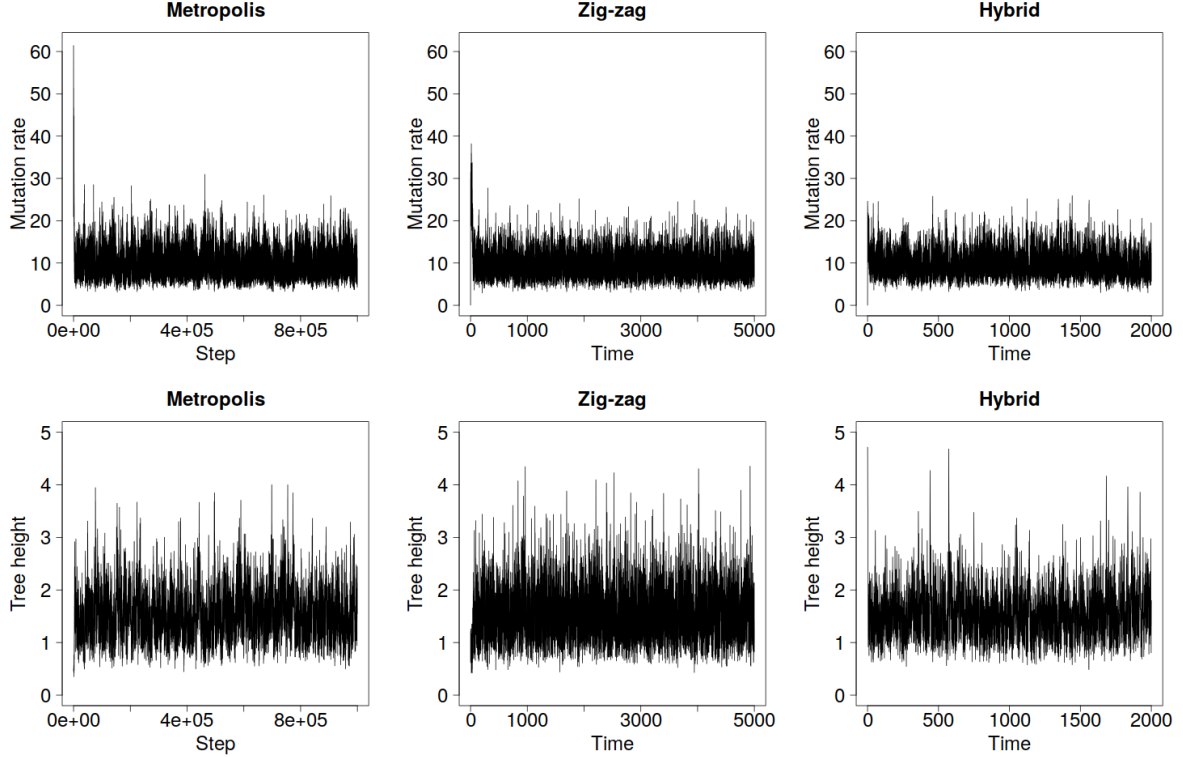


Figure 10: Trace plots for the finite sites model and data set with  $n = 50$  and  $S = 200$  consisting of 18 distinct sequences.

steps of the Metropolis–Hastings algorithm. Hence the zig-zag process is well suited to parallel architectures.

## Acknowledgements

The author was supported by ESPRC grant EP/R044732/1, and is grateful to Jure Vogrinc and Andi Wang for productive conversations on MCMC for coalescent processes, and non-reversible MCMC in general.

## Appendix

*Proof of Theorem 1.* Since the initial point  $(m, \mathbf{x})$  has a density, the Poisson construction of velocity flips and (4) ensure that corners (where the process is undefined) are hit with probability 0.

Next, we will verify the four conditions of [Dav93, Section (24.8)]. Firstly, the deterministic zig-zag motion is at a constant, bounded velocity, and hence forms a non-explosive, Lipschitz vector field. The flip rates  $\lambda_i(m, \mathbf{x}, \mathbf{v})$  are bounded, and hence integrable, on compact subsets of  $[0, T_{\Gamma^+(\partial\Omega)}(m, \mathbf{x}, \mathbf{v})]$  because  $\tilde{\pi}$  is continuously differentiable in  $\mathbf{x}$  away from boundaries. Velocity flips and boundary transitions both change the state  $(m, \mathbf{x}, \mathbf{v})$  with probability 1, the former by explicit construction and the latter because  $\Gamma^+(\partial\Omega) \cap \Gamma^-(\partial\Omega) = \emptyset$ . The final standard condition is (7). In addition, the first three conditions of [CPWF21, Theorem 1] hold by construction, and hence it suffices for stationarity of  $\tilde{\pi}$  to verify their fourth condition by checking that for all



$f \in D(L)$ ,

$$\begin{aligned} \mathbb{E}_{\tilde{\pi}}[Lf(m, \mathbf{x}, \mathbf{v})] &= \sum_{i=1}^d \int_{(m, \mathbf{x}) \in \Omega^o} \sum_{\mathbf{v} \in \{-1, 1\}^d} \{v_i \partial_i f(m, \mathbf{x}, \mathbf{v}) \\ &\quad + \lambda_i(m, \mathbf{x}, \mathbf{v})[f(m, \mathbf{x}, F_i \mathbf{v}) - f(m, \mathbf{x}, \mathbf{v})]\} \tilde{\pi}(m, \mathbf{x}, \mathbf{v}) d\mathbf{x} = 0. \end{aligned}$$

The change of variable  $\mathbf{v} \mapsto F_i \mathbf{v}$  combined with the fact that  $\tilde{\pi}(m, \mathbf{x}, \cdot)$  is uniform, then (1), and then an integration by parts yield

$$\begin{aligned} &\sum_{i=1}^d \int_{(m, \mathbf{x}) \in \Omega^o} \sum_{\mathbf{v} \in \{-1, 1\}^d} \{v_i \partial_i f(m, \mathbf{x}, \mathbf{v}) + \lambda_i(m, \mathbf{x}, \mathbf{v})[f(m, \mathbf{x}, F_i \mathbf{v}) - f(m, \mathbf{x}, \mathbf{v})]\} \tilde{\pi}(m, \mathbf{x}, \mathbf{v}) d\mathbf{x} \\ &= \sum_{i=1}^d \int_{(m, \mathbf{x}) \in \Omega^o} \sum_{\mathbf{v} \in \{-1, 1\}^d} \{v_i \partial_i f(m, \mathbf{x}, \mathbf{v}) + f(m, \mathbf{x}, \mathbf{v})[\lambda_i(m, \mathbf{x}, F_i \mathbf{v}) - \lambda_i(m, \mathbf{x}, \mathbf{v})]\} \tilde{\pi}(m, \mathbf{x}, \mathbf{v}) d\mathbf{x} \\ &= \int_{(m, \mathbf{x}) \in \Omega^o} \sum_{\mathbf{v} \in \{-1, 1\}^d} \mathbf{v} \cdot \{\nabla_{\mathbf{x}} f(m, \mathbf{x}, \mathbf{v}) - f(m, \mathbf{x}, \mathbf{v}) \nabla_{\mathbf{x}} \log(\tilde{\pi}(m, \mathbf{x}, \mathbf{v}))\} \tilde{\pi}(m, \mathbf{x}, \mathbf{v}) d\mathbf{x} \\ &= - \int_{(m, \mathbf{x}) \in \partial\Omega} \sum_{\mathbf{v} \in \Gamma^+(m, \mathbf{x})} f(m, \mathbf{x}, \mathbf{v}) (\mathbf{v} \cdot \mathbf{n}(m, \mathbf{x})) \tilde{\pi}(m, \mathbf{x}, \mathbf{v}) d\mathbf{x} \\ &\quad - \int_{(m, \mathbf{x}) \in \partial\Omega} \sum_{\mathbf{v} \in \Gamma^-(m, \mathbf{x})} f(m, \mathbf{x}, \mathbf{v}) (\mathbf{v} \cdot \mathbf{n}(m, \mathbf{x})) \tilde{\pi}(m, \mathbf{x}, \mathbf{v}) d\mathbf{x}, \end{aligned} \tag{20}$$

where  $\nabla_{\mathbf{x}}$  is the gradient operator with respect to  $\mathbf{x}$ . To conclude, we will show that the first term on the right hand side of (20) cancels with the second. Using (6), then (2), and then  $\{\mathbf{v} \in \Gamma^+(m, \mathbf{x})\} = \{-\mathbf{v} : \mathbf{v} \in \Gamma^-(m, \mathbf{x})\}$ ,

$$\begin{aligned} &- \int_{(m, \mathbf{x}) \in \partial\Omega} \sum_{\mathbf{v} \in \Gamma^+(m, \mathbf{x})} f(m, \mathbf{x}, \mathbf{v}) (\mathbf{v} \cdot \mathbf{n}(m, \mathbf{x})) \tilde{\pi}(m, \mathbf{x}, \mathbf{v}) d\mathbf{x} \\ &= - \int_{(m, \mathbf{x}) \in \partial\Omega} \sum_{\mathbf{v} \in \Gamma^+(m, \mathbf{x})} \left[ \int_{(j, \mathbf{y}) \in \partial\Omega} \sum_{\mathbf{w} \in \Gamma^-(j, \mathbf{y})} f(j, \mathbf{y}, \mathbf{w}) Q(m, \mathbf{x}, \mathbf{v}; j, d\mathbf{y}, \mathbf{w}) \right] \\ &\quad \times (\mathbf{v} \cdot \mathbf{n}(m, \mathbf{x})) \tilde{\pi}(m, \mathbf{x}, \mathbf{v}) d\mathbf{x} \\ &= - \int_{(m, \mathbf{x}) \in \partial\Omega} \sum_{\mathbf{v} \in \Gamma^+(m, \mathbf{x})} \int_{(j, \mathbf{y}) \in \partial\Omega} \sum_{\mathbf{w} \in \Gamma^-(j, \mathbf{y})} f(m, \mathbf{y}, \mathbf{w}) (\mathbf{v} \cdot \mathbf{n}(m, \mathbf{x})) Q(j, \mathbf{y}, -\mathbf{w}; m, d\mathbf{x}, -\mathbf{v}) \\ &\quad \times \tilde{\pi}(j, \mathbf{y}, -\mathbf{w}) d\mathbf{y} \\ &= - \int_{(m, \mathbf{x}) \in \partial\Omega} \sum_{\mathbf{v} \in \Gamma^-(m, \mathbf{x})} \int_{(j, \mathbf{y}) \in \partial\Omega} \sum_{\mathbf{w} \in \Gamma^-(j, \mathbf{y})} f(j, \mathbf{y}, \mathbf{w}) (\mathbf{v} \cdot \mathbf{n}(m, \mathbf{x})) Q(j, \mathbf{y}, -\mathbf{w}; m, d\mathbf{x}, \mathbf{v}) \\ &\quad \times \tilde{\pi}(j, \mathbf{y}, -\mathbf{w}) d\mathbf{y}. \end{aligned}$$

We now exchange order of integration (justified by the fact that  $f$  is bounded), then apply (3)

and the fact that  $\tilde{\pi}(\mathbf{x}, \cdot)$  is uniform to obtain

$$\begin{aligned}
& - \int_{\partial\Omega} \sum_{\mathbf{v} \in \Gamma^+(\mathbf{x})} f(\mathbf{x}, \mathbf{v}) (\mathbf{v} \cdot \mathbf{n}(\mathbf{x})) \tilde{\pi}(\mathbf{x}, \mathbf{v}) d\mathbf{x} \\
&= - \int_{\mathbf{y} \in \partial\Omega} \sum_{\mathbf{w} \in \Gamma^-(\mathbf{y})} f(\mathbf{y}, \mathbf{w}) \left[ \int_{\mathbf{x} \in \partial\Omega} \sum_{\mathbf{v} \in \Gamma^-(\mathbf{x})} (\mathbf{v} \cdot \mathbf{n}(\mathbf{x})) Q(\mathbf{y}, -\mathbf{w}; d\mathbf{x}, \mathbf{v}) \right] \tilde{\pi}(\mathbf{y}, -\mathbf{w}) d\mathbf{y} \\
&= \int_{\partial\Omega} \sum_{\mathbf{w} \in \Gamma^-(\mathbf{y})} f(\mathbf{y}, \mathbf{w}) (\mathbf{w} \cdot \mathbf{n}(\mathbf{y})) \tilde{\pi}(\mathbf{y}, \mathbf{w}) d\mathbf{y}. \tag{21}
\end{aligned}$$

Substituting (21) into (20) concludes the proof.  $\square$

*Proof of Proposition 1.* The proof consists of checking the hypotheses of Theorem 1. The set of ranked topologies of binary trees with  $n$  leaves is finite, and hence so is its restriction to topologies consistent with  $D_n$ . The initial condition has a density by assumption. The set of points from which deterministic motion at a constant velocity first hits a boundary at a corner consists of finitely many Lebesgue-null lines. Hence, since the initial condition has a density and the boundary kernel  $Q$  does not change  $(\mathbf{t}_n, \theta)$ , corners are hit with probability 0.

The augmented posterior  $\tilde{\pi}$  corresponding to (11) is  $C^1(\Omega^*)$  by inspection, and only vanishes at a boundary when  $\theta = 0$  or a branch with mutations has length zero.

The unit outward normal is  $\mathbf{n}(\mathbf{t}_n, \theta) = (0, \dots, 0, -1, 0, \dots, 0)$ , where the  $-1$  is in the  $k(\mathbf{t}_n, \theta)$ th place. We also have  $\mathbf{v}_n = F_i(F_i(\mathbf{v}_n))$ ,  $E_n = s_i(s_i(E_n))$ , the third case of (12) is invariant under permutation of  $\{E_n, p_i^\uparrow(E_n), p_i^\downarrow(E_n)\}$ , and (11) is continuous even at the boundaries (though its gradient is not), so verifying (2) and (3) is a routine calculation.

It remains to verify (7), which we do by stochastic domination. By construction, the zig-zag process crosses boundaries only when a corresponding velocity is negative. The boundary crossing flips the velocity, and a boundary corresponding to the same coordinate cannot be crossed again until a further flip. Hence, the number of jumps in the zig-zag process started at  $(E_n, \mathbf{t}_n, \theta)$  by time  $t \in (0, \infty)$  is dominated by  $n + 2Y(\mathbf{t}_n, \theta, t)$ , where

$$Y(\mathbf{t}_n, \theta, t) \sim \text{Pois} \left( |\mathbb{F}| (\|\mathbf{v}_n\|_1 t)^n \sum_{i=1}^{n-1} \left\{ \frac{(n+1-i)}{2} [|v_i|(n+\theta+|v_n|t-i) + |v_n|(t_i+|v_i|t)] \right\} \right),$$

which arises as the volume of the set which is reachable from  $(E_n, \mathbf{t}_n, \theta)$  by time  $t$  multiplied by upper bounds for (13) and (14) for positive velocities on the reachable set. Here  $\|\cdot\|_1$  is the  $L^1$ -norm, which is invariant for all valid zig-zag velocities  $\mathbf{v}_n$ . In the construction of the dominating random variable, negative velocities are assumed to flip to positive instantaneously which accounts for the initial summand of  $n$  (the maximum number of initial negative velocities) and the factor of 2.  $\square$

The remainder of the appendix details the zig-zag and Metropolis–Hastings algorithms used in Sections 4 and 5. Algorithms 2 and 3 provide pseudocode implementations of the zig-zag methods. Line 7 of Algorithm 2 is given twice, with the first applicable to Section 4 and the second to Section 5. Otherwise, both algorithms apply to both sections.

The for-loop on lines 4–10 of Algorithm 1 implements the time localization and ensures that time steps cannot cause negative entries in  $\mathbf{t}_n$ . The distribution of the increment  $\tau$  thus has an atom at the value determined by lines 4–10, signifying either a boundary crossing or a need to recompute the time localization and rates  $\{\lambda_i^* : i \in \{1, \dots, n-1, \theta\}\}$ . However, the continuous-time motion of the trajectory is not interrupted in either case; see also [BBCD<sup>+</sup>18, Section 2.2].

---

**Algorithm 2** Simulation of zig-zag process targeting  $\pi$ 

---

**Require:**  $E_n, \mathbf{t}_n, \theta, \mathbf{v}_n, t_{\text{end}}, c, K$ 

```
1: Set  $t \leftarrow 0$ 
2: while  $t < t_{\text{end}}$  do
3:   Set  $\tau \leftarrow K$  and  $I \leftarrow 0$   $\triangleright K \in (0, \infty)$  is a default increment if all  $v_i > 0$ .
4:   for  $i \in \{1, \dots, n\}$  do  $\triangleright t_n := \theta$  for brevity.
5:     if  $v_i < 0$  then
6:       Set  $d \leftarrow 1$  and  $J \leftarrow i$ 
7:       if  $(i \in \{1, n\}$  or  $E_{n,i-1} \in E_{n,i})$  and  $m_{\gamma(E_n,i)} > 0$  then  $\triangleright$  For Section 4
7:       if  $i \in \{1, n\}$  and  $m_{\gamma(E_n,i)} > 0$  then  $\triangleright$  For Section 5
8:         Set  $d \leftarrow 1 + c$  and  $J \leftarrow 0$   $\triangleright I = 0 \Rightarrow$  localization refresh with no flip.
9:         if  $-t_i/(d \times v_i) < \tau$  then
10:          Set  $\tau \leftarrow -t_i/(d \times v_i)$  and  $I \leftarrow J$ 
11:       for  $i \in \{1, \dots, n\}$  do  $\triangleright t_n := \theta$  as above.
12:         Sample  $\rho \leftarrow \text{NextFlip}(E_n, \mathbf{t}_n, \theta, \mathbf{v}_n; i, \tau)$ 
13:         if  $\rho < \tau$  then
14:           Set  $\tau \leftarrow \rho$  and  $I \leftarrow i$ 
15:       Set  $t \leftarrow t + \tau$ 
16:       for  $i \in \{1, \dots, n\}$  do  $\triangleright t_n := \theta$  as above.
17:         Set  $t_i \leftarrow t_i + v_i \tau$ 
18:       if  $I \neq 0$  then
19:         Set  $v_I \leftarrow -v_I$ 
20:       if  $I \notin \{0, 1, n\}$  and  $t_I = 0$  then
21:         if  $E_{n,I-1} \in E_{n,I}$  then
22:           Sample  $Y \sim U(\{\uparrow, \downarrow\})$ 
23:           Set  $E_n \leftarrow p_I^Y(E_n)$ 
24:         else
25:           Set  $E_n \leftarrow s_I(E_n)$ 
26: end while
```

---

Iterations of the outer while-loop of Algorithm 1 flip a velocity for one of two reasons: a flip event occurs or a boundary is hit. In the former case the reason for the flip is clear by construction. In the latter, a trajectory arriving at a boundary must have  $v_i < 0$  for the corresponding velocity. After crossing, the trajectory moves into the interior of an adjacent orthant, which corresponds to  $v_i > 0$  in its local coordinates.

---

**Algorithm 3** NextFlip

---

**Require:**  $E_n, \mathbf{t}_n, \theta, \mathbf{v}_n, i, \tau$

```

1: Set  $\rho \leftarrow 0$ 
2: repeat
3:   Sample  $s \sim \text{Exp}(\lambda_i^*)$ 
4:   Set  $\rho \leftarrow \rho + s$ 
5:   if  $\rho < \tau$  then
6:     Set  $\alpha \leftarrow \lambda_i(E_n, \mathbf{t}_n + \mathbf{v}_n \rho, \theta + v_n \rho; \mathbf{v}_n) / \lambda_i^*$ 
7:   else ▷ Guaranteed to not be the shortest waiting time.
8:     Set  $\alpha \leftarrow 1$ 
9:   Sample  $U \sim U(0, 1)$ 
10: until  $U < \alpha$ 
11: return  $\rho$ 

```

---

The Metropolis–Hastings algorithms in Sections 4 and 5 both used the same proposal mechanisms consisting of three steps. An iteration of the algorithm is one scan through all three steps, with an accept/reject correction after each step. The hybrid method only used steps 1 and 3.

1. A Gaussian random walk update of  $\theta$  reflected at 0 with proposal variance  $\sigma_\theta^2$  tuned to obtain an average acceptance probability  $\alpha_\theta \approx 1/4$ .
2. An update of holding times  $\mathbf{t}_n$  under a fixed topology.
3. A subtree-prune-regraft step.

Type 2 updates first additively perturb the initial holding time  $t_1$  as

$$\xi_1 \sim \mathcal{N}(t_1, \sigma_{\mathbf{t}_n}^2 / [n(n-1)^2]) | (\xi_1 > 0).$$

Further holding times  $i \in \{2, \dots, n-1\}$  are conditionally independently perturbed as

$$\xi_i | (\xi_1, \dots, \xi_{i-1}) \sim \mathcal{N}(t_i, \sigma_{\mathbf{t}_n}^2 / [(n-1)(n-i+1)(n-i)]) | (\xi_i > \xi_{c_{i,1}} \vee \xi_{c_{i,2}}),$$

where  $\xi_{c_{i,1}}$  and  $\xi_{c_{i,2}}$  are the perturbed times of the child nodes of the node at time  $t_i$ . Again,  $\sigma_{\mathbf{t}_n}$  was tuned so that the average acceptance probability was  $\alpha_{\mathbf{t}_n} \approx 1/4$ .

Type 3 updates sample an ordered pair of edges  $(\gamma, \gamma') \sim U(F_n \times [F_n \cup \gamma_{MRCA}])$ , where  $\gamma_{MRCA}$  is an edge connecting the MRCA to  $\infty$ . Edge  $\gamma$  is deleted, and its child  $c_\gamma$  is reattached into  $\gamma'$ . Letting  $t_\eta$  be the time of the node with label  $\eta \in E_n$  in an abuse of notation, the reattachment time  $t'$  has distribution

$$t' \sim U(t_{c_\gamma} \wedge t_{c_{\gamma'}}, t_{p_{\gamma'}}) \text{ if } \gamma' \neq \gamma_{MRCA},$$

$$t' - t_{c_{\gamma_{MRCA}}} \sim \text{Exp}(1) \text{ otherwise.}$$

Under the infinite sites model, moves into topologies incompatible with observed mutations were rejected essentially instantaneously, before costly likelihood evaluations.

Table 3 summarizes hyperparameters and quantities of interest used in the simulation.

Data set	Method	$v_\theta$	$\sigma_\theta$	$\sigma_{t_n}$	$\alpha_\theta$	$\alpha_{t_n}$	$\alpha_{\text{SPR}}$	$\kappa$
[WFDP91]	Metropolis	-	8	0.6	0.27	0.25	0.06	-
	Zig-zag	8	-	-	-	-	-	-
	Hybrid	8	10	-	0.24	-	0.06	10
$n = 550, \theta = 5.5$	Metropolis	-	6	0.25	0.23	0.24	0.03	-
	Zig-zag	6	-	-	-	-	-	-
	Hybrid	6	6	-	0.23	-	0.03	10
$n = 55, \theta = 55$	Metropolis	-	18	0.4	0.26	0.23	0.02	-
	Zig-zag	40	-	-	-	-	-	-
	Hybrid	40	18	-	0.24	-	0.01	10
[GT94]	Metropolis	-	4	0.7	0.23	0.25	0.12	-
	Zig-zag	4	-	-	-	-	-	-
	Hybrid	4	4	-	0.24	-	0.12	100
$n = 500, S = 20$	Metropolis	-	4	0.3	0.21	0.21	0.31	-
	Zig-zag	4	-	-	-	-	-	-
	Hybrid	4	3	-	0.29	-	0.32	100
$n = 50, S = 200$	Metropolis	-	14	0.6	0.20	0.27	0.04	-
	Zig-zag	20	-	-	-	-	-	-
	Hybrid	20	14	-	0.20	-	0.04	100

Table 3: Hyperparameters and acceptance probabilities for simulations in sections 4 and 5.

## References

- [ASR16] A. J. Aberer, A. Stamatakis, and F. Ronquist. An efficient independence sampler for updating branches in Bayesian Markov chain Monte Carlo sampling of phylogenetic trees. *Syst. Biol.*, 65(1):161–176, 2016.
- [BBCD<sup>+</sup>18] J. Bierkens, A. Bouchard-Côte, A. Doucet, A. B. Duncan, P. Fearnhead, T. Lienart, G. Roberts, and S. J. Vollmer. Piecewise deterministic Markov processes for scalable Monte Carlo on restricted domains. *Stat. Probab. Lett.*, 136:148–154, 2018.
- [BFR19a] J. Bierkens, P. Fearnhead, and G. Roberts. Supplement to the zig-zag process and super-efficient sampling for Bayesian analysis of big data. *Ann. Stat.*, 47, 2019.
- [BFR19b] J. Bierkens, P. Fearnhead, and G. Roberts. The zig-zag process and super-efficient sampling for Bayesian analysis of big data. *Ann. Stat.*, 47(3):1288–1320, 2019.
- [BHV01] L. J. Billera, S. P. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Adv. Appl. Math.*, 27:733–767, 2001.
- [BR17] J. Bierkens and G. Roberts. A piecewise deterministic scaling limit of lifted Metropolis-Hastings for the Curie-Weiss model. *Ann. Appl. Probab.*, 27(2):846–882, 2017.
- [BRZ19] J. Bierkens, G. Roberts, and P.-A. Zitt. Ergodicity of the zigzag process. *Ann. Appl. Probab.*, 29(4):2266–2301, 2019.
- [CFS20] A. Chevalier, P. Fearnhead, and M. Sutton. Reversible jump PDMP samplers for variable selection. Preprint, arXiv:2010.11771, 2020+.
- [CPWF21] A. Chevalier, S. Power, A. Wang, and P. Fearnhead. PDMP Monte Carlo methods for piecewise-smooth densities. Preprint, arXiv:2111.05859, 2021+.
- [Dav93] M. Davis. *Markov models and Optimization*. Chapman Hall, 1993.

- [DBZM17] V. Dinh, A. Bilge, C. Zhang, and F. A. Matsen IV. Probabilistic path Hamiltonian Monte Carlo. In *Proceedings of the 34th International Conference on Machine Learning*, Sydney, Australia, 2017.
- [Fel81] J. Felsenstein. Evolutionary trees from DNA sequences: a maximum likelihood approach. *J. Mol. Evol.*, 17:368–376, 1981.
- [FHVD17] J. M. Flegel, J. Hughes, D. Vats, and N. Dai. *mcmcse: Monte Carlo Standard Errors for MCMC*, 2017.
- [GD16] A. Gavryushkin and A. J. Drummond. The space of ultrametric phylogenetic trees. *J. Theor. Biol.*, 403:197–208, 2016.
- [GD21] P. Gagnon and A. Doucet. Nonreversible jump algorithms for Bayesian nested model selection. *J. Comput. Graph. Stat.*, 30(2):312–323, 2021.
- [Gey92] C. J. Geyer. Practical Markov chain Monte Carlo. *Statist. Sci.*, 7:473–483, 1992.
- [GS90] A. E. Gelfand and A. F. M. Smith. Sampling-based approaches to calculating marginal densities. *J. Am. Stat. Assoc.*, 85(410):398–409, 1990.
- [GT94] R. C. Griffiths and S. Tavaré. Simulating probability distributions in the coalescent. *Theor. Popln Biol.*, 46:131–159, 1994.
- [Has70] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57:97–109, 1970.
- [HD12] S. Höhna and A. J. Drummond. Guided tree topology proposals for Bayesian phylogenetic inference. *Syst. Biol.*, 61(1):1–11, 2012.
- [HDPD08] S. Höhna, M. Defoin-Platel, and A. J. Drummond. Clock-constrained tree proposal operators in Bayesian phylogenetic inference. In *BioInformatics and BioEngineering*, pages 1–7, 8th IEEE International Conference on IEEE, 2008.
- [JC69] T. H. Jukes and C. R. Cantor. Evolution of protein molecules. In H. N. Munro, editor, *Mammalian protein metabolism*, pages 21–132. Academic Press, New York, 1969.
- [JZH<sup>+</sup>20] X. Ji, Z. Zhang, A. Holbrook, A. Nishimura, G. Baele, A. Rambaut, P. Lemey, and M. A. Suchard. Gradients do grow on trees: a linear-time  $O(N)$ -dimensional gradient for statistical phylogenetics. *Mol. Biol. Evol.*, 37(10):3047–3060, 2020.
- [Kin82] J. F. C. Kingman. The coalescent. *Stochast. Process. Applic.*, 13(3):235–248, 1982.
- [LVH<sup>+</sup>08] C. Lakner, P. Van Der Mark, J. P. Huelsenbeck, B. Larget, and F. Ronquist. Efficiency of Markov chain Monte Carlo tree proposals in Bayesian phylogenetics. *Syst. Biol.*, 57:86–103, 2008.
- [MV05] E. Mossel and E. Vigoda. Phylogenetic MCMC algorithms are misleading on mixtures of trees. *Science*, 309:2207–2209, 2005.
- [NDL20] A. Nishimura, D. B. Dunson, and J. Lu. Discontinuous Hamiltonian Monte Carlo for discrete parameters and discontinuous likelihoods. *Biometrika*, 107(2):365–380, 2020.
- [Nea10] R. M. Neal. MCMC using Hamiltonian dynamics. In *Handbook of Markov chain Monte Carlo*. CRC Press, 2010.
- [VR21] G. Vasdekis and G. O. Roberts. Speed up zig-zag. Preprint, arXiv:2103.16620, 2021+.

- [Wak09] J. Wakeley. *Coalescent theory: an introduction*. Roberts & Co, 2009.
- [Wat75] G. A. Watterson. On the number of segregating sites in genetical models without recombination. *Theor. Popln Biol.*, 7:256–276, 1975.
- [WFDP91] R. H. Ward, B. L. Frazier, K. Dew, and S. Pääbo. Extensive mitochondrial diversity within a single Amerindian tribe. *Proc. Natl. Acad. Sci. USA*, 88:8720–8724, 1991.