A NOTE ON THE EXACT SIMULATION OF SPHERICAL BROWNIAN MOTION

ALEKSANDAR MIJATOVIĆ, VENO MRAMOR, AND GERÓNIMO URIBE BRAVO

ABSTRACT. We describe an exact simulation algorithm for the increments of Brownian motion on a sphere of arbitrary dimension, based on the skew-product decomposition of the process with respect to the standard geodesic distance. The radial process is closely related to a Wright-Fisher diffusion, increments of which can be simulated exactly using the recent work of Jenkins & Spanò (2017) [JS17]. The rapid spinning phenomenon of the skew-product decomposition then yields the algorithm for the increments of the process on the sphere.

KEYWORDS. exact simulation; skew-product decomposition; spherical Brownian motion; Wright-Fisher diffusion

1. INTRODUCTION

Brownian motion (BM) can be viewed as a continuous time random walk with symmetric increments and is as such of fundamental importance in physics and other natural sciences. In applications one is often interested in the BM on curved surfaces and other manifolds, see e.g. [KDPN00] and [LTT08] for the modelling of the fluorescent marker molecules in cell membranes and the motion of bacteria or any other diffusing particles, respectively. Often Monte Carlo simulation algorithms for such models on curved spaces are constructed using approximate tangent plane methods, which are accurate only for very small time steps, making the algorithms computationally expensive. Algorithms allowing simulation over longer time steps are hence of particular interest in the physics literature. For example, [NEE03] gives a simulation algorithm for the BM on the three-dimensional sphere, based on its quaternionic structure. This algorithm is neither exact nor does it generalise easily to other dimensions. In [CEE10] it is applied to design an approximate simulation algorithm for the BM on the two-dimensional sphere. A further approximate algorithm for the simulation of BM on the two-sphere is given in [GSS12]. All the aforementioned algorithms are based on an approximation of the transition density of the BM on the relevant sphere.

In contrast, for any dimension $d \geq 3$, Algorithm 1 simulates exactly the increments of the BM Z on the sphere $\mathbb{S}^{d-1} := \{z \in \mathbb{R}^d \; ; \; |z| = 1\}$. It is based on two facts established in Section 2: the radial part of Z (with respect to the standard metric on $\mathbb{S}^{d-1}$) can be transformed to a Wright-Fisher diffusions and, due to the rapid spinning of the skew-product decomposition of Z, its angular component is uniform on $\mathbb{S}^{d-2}$.

Algorithm 1 Simulation of the increment of Brownian motion Z on $\mathbb{S}^{d-1}$ over any time interval

---

Require: Starting point $z \in \mathbb{S}^{d-1}$ and time horizon $t > 0$
1: Simulate the radial component: $X \sim \mathrm{WF}_{0,t}\left(\frac{d-1}{2}, \frac{d-1}{2}\right)$ ▷ Algorithm 2 in Appendix A
2: Simulate the angular component: Y uniform on $\mathbb{S}^{d-2}$
3: Set $u := (\mathrm{e}_d - z)/|\mathrm{e}_d - z|$ and $O(z) := I - 2uu^\top$
4: **return** $O(z)(2\sqrt{X(1-X)}Y^\top, 1 - 2X)^\top$

The vectors in Algorithm 1 are column vectors of appropriate dimension, $\mathrm{e}_d := (0, \ldots, 0, 1)^\top \in \mathbb{S}^{d-1}$ denotes the north pole of the sphere and $O(z)$ is the reflection of $\mathbb{R}^d$ across the hyperplane through the origin

with the normal u. Algorithm 2 is the exact simulation algorithm for the increment of the Wright-Fisher diffusion given in [JS17, Alg. 1] (see Eq. (A.1) below for a definition of Wright-Fisher diffusions). Step 2 in Algorithm 1 consists of simulating a vector N in $\mathbb{R}^{d-1}$ with independent standard normal components and setting $Y = N/|N|$ (we denote the standard Euclidean norm by $|\cdot|$ throughout).

The key property of the orthogonal matrix $O(z) \in \mathbb{R}^d \otimes \mathbb{R}^d$ in Algorithm 1 is $O(z)\mathrm{e}_d = z$. In fact, any orthogonal matrix in $\mathbb{R}^d \otimes \mathbb{R}^d$ with this property would lead to an exact sample from the increment of BM on $\mathbb{S}^{d-1}$. Indeed, if $O_1(z), O_2(z) \in \mathbb{R}^d \otimes \mathbb{R}^d$ are two such orthogonal matrices, then the product $O_1^{-1}(z)O_2(z)$ fixes e_d and is given by an orthogonal transformation $\widetilde{O}(z) \in \mathbb{R}^{d-1} \otimes \mathbb{R}^{d-1}$ on the orthogonal complement $\{\mathrm{e}_d\}^\perp$ in $\mathbb{R}^d$. Hence $O_2(z)(2\sqrt{X(1-X)}Y^\top, 1-2X)^\top = O_1(z)(2\sqrt{X(1-X)}(\widetilde{O}(z)Y)^\top, 1-2X)^\top$ implying

$$O_2(z)(2\sqrt{X(1-X)}Y^\top, 1-2X)^\top \overset{d}{=} O_1(z)(2\sqrt{X(1-X)}Y^\top, 1-2X)^\top,$$

where $\overset{d}{=}$ denotes equality in law. The formula for $O(z)$ in Algorithm 1 is chosen due to its simplicity.

Algorithm 1 exploits the symmetry of both the geometry of $\mathbb{S}^{d-1}$ and the law of spherical BM to reduce the simulation problem in any dimension to the simulation of an increment of a one-dimensional diffusion. Unlike the simulation methods in [NEE03, CEE10, GSS12], Algorithm 1 depends on the dimension d only through the value of the mutation parameters in the Wright-Fisher diffusion. As discussed in [JS17, Sec. 4], Algorithm 2, which simulates exactly from the law of the increment of the Wright-Fisher diffusion, is numerically stable for the time intervals of length $t \geq 0.05$. This is essentially because, when t is very small, the modulus of the summands in the alternating series [JS17, Eq. (5)] is increasing as a function of the index for large values of the index, before it starts to decrease monotonically. Conversely, Algorithm 2 (and hence Algorithm 1) becomes more efficient with increasing time horizon t (for small t, a normal approximation can be used to obtain a fast approximate algorithm substituting Algorithm 2, see [JS17, Thm 1] and the comments therein). Dimension d does not affect the performance of Algorithm 1, cf. Appendix A.

Algorithm 1 above and the Markov property of the BM Z on the sphere $\mathbb{S}^{d-1}$ yield an exact sample from any finite-dimensional marginal of Z. Moreover, as the BM on $\mathbb{S}^1 \subseteq \mathbb{C}$ can be represented as $Z = \exp(\mathrm{i}W)$, where $\mathrm{i}^2 = -1$ and W denotes a standard BM on $\mathbb{R}$, the assumption $d \geq 3$ in Algorithm 1 is not restrictive.

Our main technical result (Proposition 2.1 below) states that the spherical BM Z enjoys a skew-product decomposition in which the last component of Z is independent of the normalized and suitably time-changed remaining components. Its main application in the present paper consists of establishing the validity of Algorithm 1. The proof of Proposition 2.1 is analogous to that of [MMU18, Thm 1.5]. However, we note that Proposition 2.1 is not a corollary of [MMU18, Thm 1.5], as the quoted theorem implies only the independence of the modulus of the last component (but not of its sign). This difference makes the claim of Proposition 2.1 more general than [MMU18, Thm 1.5] and (we hope) of independent interest.

The classical skew-product decomposition of Brownian motion ζ in $\mathbb{R}^d$, started away from the origin, features the BM Z on $\mathbb{S}^{d-1}$ as follows: $Z_t = \zeta_{C_t}/|\zeta_{C_t}|$, where the increasing process $(C_t)_{t\geq 0}$ is the inverse of the additive functional $\int_0^\cdot |\zeta_u|^{-2}\mathrm{d}u$ of ζ. Moreover, the process Z and $|\zeta|$ are independent of each other. It is tempting to try to use the above representation to obtain a sample from Z_t. One would first need to simulate C_t, then get a sample from the conditional law of ζ_{C_t}, given C_t, and finally normalise this sample to get an $\mathbb{S}^{d-1}$-valued random element identically distributed as Z_t. We stress however that this does not yield an approximate simulation algorithm for the increments of Z, let alone an exact one. Even if one is able to simulate C_t, which is in itself non-trivial and can to the best of our knowledge only be done approximately, the key obstacle is the fact that the process ζ is not independent of the time C_t. More precisely, since C_t

depends on the trajectory of ζ up to C_t, the conditional law of ζ_{C_t}, given C_t, is intractable, so the second step of the proposed sampling procedure cannot be achieved.

Finally, we note that Algorithm 1 yields an algorithm for the exact simulation of the increments of the BM on the real $\mathbb{RP}^n$, complex $\mathbb{CP}^n$ and quaternionic $\mathbb{HP}^n$ projective spaces (for any integer n). These are Riemannian manifolds of the (real) dimension n, $2n$ and $4n$, respectively, with canonical Riemannian metrics described in Appendix B. The Riemannian submersion π (mapping $\mathbb{S}^n \to \mathbb{RP}^n$, $\mathbb{S}^{2n+1} \to \mathbb{CP}^n$ and $\mathbb{S}^{4n+3} \to \mathbb{HP}^n$) by Lemma B.1 below projects the BM Z on the relevant sphere to the BM $\pi(Z)$ on the projective space. Since $\pi(z) = [z]$ only converts the standard coordinates of the point z on the sphere to the homogeneous coordinates $[z]$ in the projective space, the random element

$$\pi\left(O(z)(2\sqrt{X(1-X)}Y^\top, 1-2X)^\top\right),$$

where matrix $O(z)$ and random variables X, Y are the same as in Algorithm 1, gives the exact sample of the increment of the BM on the projective space, started at $[z]$, over the time horizon $t > 0$.

2. The skew-product decomposition and Algorithm 1

Brownian motion $Z = (Z_t)_{t\geq 0}$ on the sphere $\mathbb{S}^{d-1}$ is a Feller process generated by the Laplace-Beltrami operator corresponding to the Riemannian metric on $\mathbb{S}^{d-1}$ induced by the ambient Euclidean space $\mathbb{R}^d$, see [Hsu02]. There are a number of different ways of representing BM on a sphere. The most useful for our purposes is the Stroock representation of Z, given (in Itô form) by SDE (2.1) on $\mathbb{R}^d$,

$$(2.1) \qquad \mathrm{d}Z_t = (I - Z_t Z_t^\top)\mathrm{d}B_t - \frac{d-1}{2}Z_t\mathrm{d}t, \qquad Z_0 \in \mathbb{S}^{d-1},$$

which possesses the unique strong solution, where B is a BM on $\mathbb{R}^d$ and I denotes the identity matrix of appropriate dimension (cf. [Hsu02, Ch.3, §3, p.83]). However, as (2.1) and other representations of Z alluded to above are non-constructive, we cannot use them directly for the exact simulation of the increments of Z in dimension $d \geq 3$.

As explained in the introduction, the key idea is to identify the skew-product decomposition of Z and exploit the rapid spinning phenomenon of the angular component at the starting point Z_0, together with the symmetries of the sphere $\mathbb{S}^{d-1}$ and the law of the spherical BM Z to obtain Algorithm 1. Let $D = (D_t)_{t\geq 0}$ be the geodesic distance $D_t := \mathrm{dis}(Z_t) \in [0, \pi]$ between Z_t and some fixed point $w \in \mathbb{S}^{d-1}$. In [IM96, p. 269] and [PR88], the authors show that D satisfies the SDE $\mathrm{d}D_t = \mathrm{d}\beta_t + \frac{d-2}{2}\cot(D_t)\mathrm{d}t$, where β is a scalar BM. Since $\mathrm{dis}(z) = \arccos(\langle z, w\rangle)$ for any $z, w \in \mathbb{S}^{d-1}$, the natural transformation $\widetilde{\mathcal{X}} := \cos(D)$ leads to the Jacobi diffusion (see e.g. [WY98] and the proof of Proposition 2.1 below) satisfying the SDE

$$(2.2) \qquad \mathrm{d}\widetilde{\mathcal{X}}_t = \sqrt{1 - \widetilde{\mathcal{X}}_t^2}\mathrm{d}\widetilde{\beta}_t - \frac{d-1}{2}\widetilde{\mathcal{X}}_t\mathrm{d}t, \qquad \widetilde{\mathcal{X}}_0 = \langle Z_0, w\rangle,$$

where $\widetilde{\beta} = -\beta$ is a standard scalar BM. This simple (and in our context crucial) observation has to the best of our knowledge not been made in the literature so far. Note that $\widetilde{\mathcal{X}}_t = \langle Z_t, w\rangle$. We henceforth fix $w = \mathrm{e}_d$ and obtain $\widetilde{\mathcal{X}}_t = Z_t^d$, making $\widetilde{\mathcal{X}}$ the process considered in [MMU18, Proposition 1.1] for $n = 1$ and $\ell = d-1$. A linear transformation $\mathcal{X}_t := \frac{1-\widetilde{\mathcal{X}}_t}{2}$ yields the SDE

$$(2.3) \qquad \mathrm{d}\mathcal{X}_t = \sqrt{\mathcal{X}_t(1-\mathcal{X}_t)}\mathrm{d}\beta_t + \left(\frac{d-1}{4}(1-\mathcal{X}_t) - \frac{d-1}{4}\mathcal{X}_t\right)\mathrm{d}t, \qquad \mathcal{X}_0 = \frac{1-Z_0^d}{2},$$

making $\mathcal{X}$ a Wright-Fisher diffusion with mutation parameters $\theta_1 = \frac{d-1}{2} = \theta_2$, see (A.1) below.

A weak form of the skew-product decomposition (on the level of generators) of the BM on the sphere has been established in [PR88], see also [IM96, p. 269]. In order to give a path-wise skew-product representation of Z, note first that for $d \geq 3$ the Wright-Fisher diffusion $\mathcal{X}$ visits neither 0 nor 1. We may thus introduce the following time-changes: for $0 \leq s \leq t$ define $S_s(t) := \int_s^t \frac{1}{4\mathcal{X}_u(1-\mathcal{X}_u)}du$, satisfying $\lim_{t\to\infty} S_s(t) = \infty$ (see proof of Proposition 2.1 below), and its inverse $T_s \colon [0,\infty) \to [s,\infty)$. We now state our main result.

Propositon 2.1 (Skew-product decomposition of the BM on $\mathbb{S}^{d-1}$). *Let $d \geq 3$ and Z be a solution of SDE* (2.1). *Pick $s \in [0,\infty)$ and assume that either $s > 0$ or $s = 0$ and $Z_0 \neq e_d$. Let $U = (Z^1,\ldots,Z^{d-1})^\top$ denote the first $d - 1$ components of Z. Then $|U| = 2\sqrt{\mathcal{X}(1-\mathcal{X})}$ and the process $\widehat{V} = (\widehat{V}_t)_{t\geq 0}$, given by $\widehat{V}_t := U_{T_s(t)}/|U_{T_s(t)}|$ is a BM on the sphere $\mathbb{S}^{d-2}$, started at $\widehat{V}_0 = U_s/|U_s|$, independent of $\mathcal{X}$. Hence we obtain the skew-product decomposition $Z_t = (2\sqrt{\mathcal{X}_t(1-\mathcal{X}_t)}\widehat{V}_{S_s(t)}^\top, 1 - 2\mathcal{X}_t)^\top$ for $t \geq s$. Furthermore, if $Z_0 = e_d$ (i.e. $U_0 = 0$), then $\widehat{V}_t$ is uniformly distributed on $\mathbb{S}^{d-2}$ for any $t > 0$ and subsequently evolves as a stationary BM on the sphere.*

Since time changing a stationary process by an independent time-change does not affect the marginal distributions of the process, the second part of the theorem immediately yields Corollary 2.2, which in turn justifies Algorithm 1 for the exact simulation of the increments of BM on the sphere $\mathbb{S}^{d-1}$.

Corollary 2.2. *Let $d \geq 3$ and Z be the BM on $\mathbb{S}^{d-1}$ started at $Z_0 = e_d$. Additionally, let $\mathcal{W}$ be the unique strong solution of SDE* (A.1) *with mutation parameters $\theta_1 = \frac{d-1}{2} = \theta_2$, started at $\mathcal{W}_0 = 0$, and Y a uniformly distributed random vector on $\mathbb{S}^{d-2}$, independent of $\mathcal{W}$. Then for each $t \geq 0$, the random vectors Z_t and $(2\sqrt{\mathcal{W}_t(1-\mathcal{W}_t)}Y^\top, 1 - 2\mathcal{W}_t)^\top$ are identically distributed.*

It may appear that the skew-product decomposition in Proposition 2.1 can be applied directly to simulate the increments of the spherical BM started at any point in $\mathbb{S}^{d-1}$. Since the increment of the Wright-Fisher diffusion can be simulated exactly using Algorithm 2 below, Proposition 2.1 appears to reduce the problem to the simulation of the increment of the BM on $\mathbb{S}^{d-2}$. Recursively, this would reduce the problem to the simulation of the BM on $\mathbb{S}^1$, where the algorithm is trivial. Unfortunately, for this direct approach to work we would need to obtain a sample from law of the pair $(\mathcal{X}_t, \int_0^t \frac{1}{4\mathcal{X}_u(1-\mathcal{X}_u)}du)$, which at the time of writing we do not know how to do. As stated in Corollary 2.2, this problem disappears when the BM Z starts at the north pole e_d. Since for any orthogonal matrix $A \in \mathbb{R}^d \otimes \mathbb{R}^d$ the process AZ solves the SDE in (2.1) started at AZ_0, Corollary 2.2 and the orthogonal matrix $O(z)$ in Algorithm 1 circumvent the need to simulate from the law of the pair $(\mathcal{X}_t, \int_0^t \frac{1}{4\mathcal{X}_u(1-\mathcal{X}_u)}du)$ for an arbitrary starting point $z \in \mathbb{S}^{d-1}$.

Remark 2.3. It is not difficult to see that the Brownian motion $\widetilde{Z} = (\widetilde{Z}_t)_{t\geq 0}$ on the sphere of radius $R > 0$ in $\mathbb{R}^d$ satisfies the following SDE with a unique strong solution (see e.g. [GMW18, Lem. 3.6(e)]):

$$d\widetilde{Z}_t = (I - \widetilde{Z}_t\widetilde{Z}_t^\top/|\widetilde{Z}_t|^2)dB_t - \frac{d-1}{2}\widetilde{Z}_t/|\widetilde{Z}_t|^2 dt, \qquad |\widetilde{Z}_0| = R.$$

Thus we may define the BM $\widetilde{Z}$ by $\widetilde{Z}_t = RZ_{t/R^2}$, where Z is the BM on $\mathbb{S}^{d-1}$ satisfying the SDE in (2.1), and Algorithm 1 can be applied to produce the exact sample of the increments of $\widetilde{Z}$.

Proof of Proposition 2.1. In order to apply [MMU18, Thm 1.5], let $n = d - 1$, $\ell = 1$ and set $\gamma \equiv 1$, $g \equiv \frac{d-1}{2}$. Then the process U in Proposition 2.1 equals the process X considered in [MMU18, Thm 1.5]. In particular, the processes $|U|$ and R are equal implying that the time-changes $S_s(t)$ and $T_s(t)$ used in both

Proposition 2.1 and [MMU18, Thm 1.5] coincide. Moreover, since [MMU18, Lem. 2.3] holds in the current setting the time-changes are well-defined. We finally note that $|U| = 2\sqrt{\mathcal{X}(1-\mathcal{X})}$.

As mentioned in the introduction, the only thing left to prove is the independence of $\widehat{V}$ and $\mathcal{X} = (1-Z^d)/2$. The argument is analogous to the one that yielded the independence of $\widehat{V}$ and R in the proof of [MMU18, Thm 1.5], but the statement does not follow directly from [MMU18, Thm 1.5]. Let $V := U/|U|$ and note that, as in the proof of [MMU18, Thm 1.5], Itô's formula applied to V yields the following $\mathrm{d}V_t = |U_t|^{-1}(I - V_tV_t^\top)\mathrm{d}\widetilde{B}_t - |U_t|^{-2}\frac{d-2}{2}V_t\mathrm{d}t$, where $\widetilde{B} = (B^1,\ldots,B^{d-1})^\top$. By [MMU18, Prop. 1.1] (with $n = 1, \ell = d-1$) we have that $\widetilde{\mathcal{X}} = Z^d$ satisfies SDE (2.2) above with $\widetilde{\mathcal{X}}_0 = Z_0^d$. Moreover, as in the proof of [MMU18, Prop. 1.1], the scalar BM $\widetilde{\beta}$ in (2.2) is given by $\mathrm{d}\widetilde{\beta}_t = -(1-2\mathcal{X}_t)V_t^\top\mathrm{d}\widetilde{B}_t + |U_t|\,\mathrm{d}B_t^d$. Consequently, $\mathcal{X}$ satisfies SDE (2.3) with $\mathcal{X}_0 = \frac{1-Z_0^d}{2}$ and the scalar BM $\beta = -\widetilde{\beta}$. Since $\widehat{V} = V_{T_s(\cdot)}$, the change-of-time formulae [RY99, Ch. V, §1] for the Itô and the Lebesgue-Stieltjes integrals imply that $\widehat{V}_t - \widehat{V}_0 = \int_0^t (I - \widehat{V}_u\widehat{V}_u^\top)\mathrm{d}\widetilde{W}_u - \int_0^t \frac{d-2}{2}\widehat{V}_u\mathrm{d}u$, where $\widetilde{W}_t = \int_s^{T_s(t)} |U_u|^{-1}\,\mathrm{d}\widetilde{B}_u$ is a standard BM on $\mathbb{R}^{d-1}$ by Levy's characterization. This implies that $\widehat{V}$ solves the SDE in (2.1) in dimension $d-1$, making it a BM on $\mathbb{S}^{d-2}$.

In order to conclude that $\mathcal{X}$ and $\widehat{V}$ are independent, we assume without loss of generality that the underlying probability space with filtration $(\mathcal{F}_t)_{t\geq 0}$ supports a further scalar $(\mathcal{F}_t)$-BM ξ, independent of B. Define the continuous local martingale $W = (W_t)_{t\geq 0}$ in $\mathbb{R}^{d-1}$ by

$$W_t = \int_s^{T_s(t)} \frac{1}{|U_u|}(I - V_uV_u^\top)\mathrm{d}\widetilde{B}_u + \int_s^{T_s(t)} \frac{V_u}{|U_u|}\mathrm{d}\xi_u.$$

Since $(I - V_uV_u^\top)(I - V_uV_u^\top) = I - V_uV_u^\top$, it follows easily from the independence of B and ξ that the components of its quadratic covariation W are given by $\langle W^i,W^j\rangle_t = \int_s^{T_s(t)} \delta_{ij}|U_u|^{-2}\mathrm{d}u = \delta_{ij}t$ for any $i,j \in \{1,\ldots,d-1\}$. Hence, by Levy's characterisation, the process W is a standard $(\mathcal{G}_t)$-BM, where the filtration $(\mathcal{G}_t)_{t\geq 0}$ is defined by $\mathcal{G}_t := \mathcal{F}_{T_s(t)}$ for all $t \geq 0$. Moreover, since

$$(I - V_uV_u^\top)\cdot\left[\tfrac{1}{|U_u|}(I - V_uV_u^\top),\quad \tfrac{V_u}{|U_u|}\right] = \left[\tfrac{1}{|U_u|}(I - V_uV_u^\top),\quad 0\right]$$

we can apply the change-of-time formulae for the stochastic and Lebesgue-Stieltjes integral [RY99, Ch. V, §1] to obtain $\widehat{V}_t - \widehat{V}_0 = \int_0^t (I - \widehat{V}_u\widehat{V}_u^\top)\mathrm{d}W_u - \int_0^t \frac{d-2}{2}\widehat{V}_u\mathrm{d}u$.

Since SDE in (2.1) in dimension $d-1$ has a unique strong solution, the independence of $\widehat{V}$ and $\mathcal{X}$ follows from the independence of the driving BMs W and β. Since W and β run on different time scales, we first note that the Markov property of W implies that W depends on $\mathcal{G}_0 = \mathcal{F}_s$ only through its position at time zero, i.e. $W_0 = 0$, making it independent of $\mathcal{F}_s$. In particular, W is independent of $(\beta_t)_{t\in[0,s]}$. Therefore it is sufficient to prove that W is independent of $(\beta_{u+s} - \beta_s)_{u\geq 0}$. For any $t \in [0,\infty)$, define $\eta_t := \beta_{T_s(t)} - \beta_s = \int_s^{T_s(t)}(1 - 2\mathcal{X}_u)V_u^\top\mathrm{d}\widetilde{B}_u - \int_s^{T_s(t)} |U_u|\,\mathrm{d}B_u^d$ so that η is a $(\mathcal{G}_t)$-local martingale. Simple calculations show that $\langle W^i,\eta\rangle_t = 0$ for any component W^i of W and the quadratic variation $\langle\eta\rangle_t = T_s(t) - s$ satisfies $\langle\eta\rangle_{S_s(u+s)} = u$ for all $u \geq 0$. By Knight's Theorem (also known as the multidimensional Dambis-Dubins-Schwarz Theorem [RY99, Ch. V, Thm 1.9]) we have that W and $(\eta_{S_s(t+s)})_{t\geq 0} = (\beta_{s+t} - \beta_s)_{t\geq 0}$ are independent BMs. This concludes the proof of the proposition. $\qquad\square$

APPENDIX A. ALGORITHMS FOR SIMULATION OF WRIGHT-FISHER DIFFUSION

We recall the algorithm for the exact simulation of the marginals of Wright-Fisher diffusions. This appendix is taken from [JS17, Section 2]. Fix parameters $\theta_1,\theta_2 \geq 0$ and consider the SDE

$$(A.1) \qquad \mathrm{d}\mathcal{W}_t = \sqrt{\mathcal{W}_t(1-\mathcal{W}_t)}\mathrm{d}B_t + \frac{1}{2}(\theta_1(1-\mathcal{W}_t) - \theta_2\mathcal{W}_t)\mathrm{d}t, \qquad \mathcal{W}_0 = x \in [0,1].$$

Its unique strong solution is know as the Wright-Fisher diffusion with mutation parameters θ_1 and θ_2.[1]

Algorithm 2 Exact simulation from the law $\mathrm{WF}_{x,t}(\theta_1,\theta_2)$

Require: Mutation parameters θ_1 and θ_2, starting point $x \in [0,1]$ and time horizon $t > 0$
1: Simulate $M \overset{d}{=} A_\infty^\theta(t)$ ▷ Use [JS17, Alg. 2]
2: Simulate $L \sim \mathrm{Binomial}(M,x)$
3: Simulate $Y \sim \mathrm{Beta}(\theta_1 + L,\theta_2 + M - L)$
4: **return** Y

The random variable $A_\infty^\theta(t)$ in step 1 is integer-valued with the mass function $\{q_m^\theta(t); m = 0,1,\ldots\}$, where $\theta = \theta_1 + \theta_2$, that can be described as follows. Let $\{A_n^\theta(t); t \geq 0\}$ be a pure death process on the non-negative integers, started at $A_n^\theta(0) = n$, where the only transitions are of the form $m \mapsto m - 1$ and occur at rate $m(m + \theta - 1)/2$ for each $m \in \{1, \ldots, n\}$. Then $q_m^\theta(t) = \mathbb{P}[A_\infty^\theta(t) = m]$ can be expressed as the limit $q_m^\theta(t) = \lim_{n\to\infty} \mathbb{P}[A_n^\theta(t) = m]$. The coefficients $q_m^\theta(t)$ are known in terms of an alternating series [JS17, Eq. (5)], with summands whose absolute values are ultimately monotonically decreasing. This property is exploited in [JS17, Alg. 2] for the simulation of $A_\infty^\theta(t)$. In theory, [JS17, Alg. 2] is exact for all values of parameters. In practice, when the time horizon t is very small, the algorithm runs into the running time and floating number precision problems, as it requires the multiplication of very large and very small numbers. It is noted in [JS17] that [JS17, Alg. 2] works well for $t \geq 0.05$ (the values of θ_1,θ_2 do not seem to affect the performance). For $t < 0.05$, a normal approximation for $A_\infty^\theta(t)$ can be used, see [JS17, Thm 1]. This yields an efficient approximate simulation algorithm for the law $\mathrm{WF}_{x,t}(\theta_1,\theta_2)$.

Appendix B. Brownian motion on projective spaces

Pick a field $\mathbb{F} \in \{\mathbb{R},\mathbb{C},\mathbb{H}\}$, where $\mathbb{C}$ are the complex numbers and $\mathbb{H}$ denotes the quaternions, consider $\mathbb{F}^{n+1}$ as an $(n+1)$-dimensional vector space over $\mathbb{F}$ and recall that the n-dimensional projective space $\mathbb{FP}^n$ is defined as a space of all 1-dimensional subspaces of $\mathbb{F}^{n+1}$. More precisely, define $\mathbb{FP}^n$ to be the set of all equivalence classes $[x_0 : \cdots : x_n]$, where the $(n+1)$-tuple $(x_0,\ldots,x_n)^\top$ is in $\mathbb{F}^{n+1}\backslash\{0\}$ and $[x_0 : \cdots : x_n] = [y_0 : \cdots : y_n]$ if and only if there exists a scalar $\lambda \in \mathbb{F}\backslash\{0\}$ such that $x_i = \lambda y_i$ for each $i \in \{0,\ldots,n\}$. For any $x \in \mathbb{F}^{n+1}\backslash\{0\}$, $[x] := [x_0 : \cdots : x_n]$ denotes the *homogeneous coordinates* of the corresponding point in $\mathbb{FP}^n$.

Let $\pi \colon \mathbb{F}^{n+1}\backslash\{0\} \to \mathbb{FP}^n$ be the quotient map given by $\pi(x) := [x]$. $\mathbb{FP}^n$ is topologised by the quotient topology i.e. a subset $U \subseteq \mathbb{FP}^n$ is open if and only if $\pi^{-1}(U)$ is open in $\mathbb{F}^{n+1}\backslash\{0\}$. It is easy to see that $\mathbb{FP}^n$ is Hausdorff, second-countable and compact since the restriction of π to the sphere $\mathbb{S}^n$ (resp. $\mathbb{S}^{2n+1}$, $\mathbb{S}^{4n+3}$) if $\mathbb{F} = \mathbb{R}$ (resp. $\mathbb{C}$, $\mathbb{H}$), is surjective. Moreover, $\mathbb{FP}^n$ is a smooth manifold. Indeed, for any $i \in \{0,\ldots,n\}$, let $\widetilde{U}_i = \{x \in \mathbb{F}^{n+1} \; ; \; x_i \neq 0\}$ and note that $\widetilde{U}_i = \pi^{-1}(\pi(\widetilde{U}_i))$, making the set $U_i = \pi(\widetilde{U}_i)$ open in $\mathbb{FP}^n$. Define the chart $\varphi_i \colon U_i \to \mathbb{F}^n$ by $\varphi_i([x]) = (\frac{x_0}{x_i},\ldots,\frac{x_{i-1}}{x_i},\frac{x_{i+1}}{x_i},\ldots,\frac{x_n}{x_i})^\top$. The map φ_i is well-defined (since it assigns the same value to each $(n + 1)$-tuple in a given equivalence class in U_i) and a homeomorphism with the inverse $\varphi_i^{-1}(y_1,\ldots,y_n) = [y_1 : \cdots : y_{i-1} : 1 : y_i : \cdots : y_n]$. The charts (U_i,φ_i), $i \in \{0,\ldots,n\}$ are smoothly compatible. Thus $\mathbb{FP}^n$ is a compact smooth manifold of dimension n (resp. $2n$, $4n$) if $\mathbb{F} = \mathbb{R}$ (resp. $\mathbb{C}$, $\mathbb{H}$).

An equivalent way of defining the smooth structure on $\mathbb{FP}^n$ is via the action on $\mathbb{S}^n$ (resp. $\mathbb{S}^{2n+1}$, $\mathbb{S}^{4n+3}$) if $\mathbb{F} = \mathbb{R}$ (resp. $\mathbb{C}$, $\mathbb{H}$) of the multiplicative group of unit elements in $\mathbb{F}$, see proof of Lemma B.1 below. This action is free, proper and smooth, so by the Quotient Manifold Theorem [Lee12, Thm 21.10] there exists

---

[1]This notation coincides with the one used in [JS17] but differs from the one in [MMU18] by a factor of 2.

a unique smooth structure on the projective spaces $\mathbb{RP}^n = \mathbb{S}^n/\mathbb{S}^0$, $\mathbb{CP}^n = \mathbb{S}^{2n+1}/\mathbb{S}^1$ and $\mathbb{HP}^n = \mathbb{S}^{4n+3}/\mathbb{S}^3$ that makes the quotient projection π from a sphere onto $\mathbb{FP}^n$ a smooth submersion i.e. its differential $d\pi_x$ is surjective at each point x in the sphere.

The spheres $\mathbb{S}^n$, $\mathbb{S}^{2n+1}$ and $\mathbb{S}^{4n+3}$ are Riemannian manifolds with the respective metrics induced by the ambient Euclidean spaces and the groups of unit elements $\mathbb{S}^0$, $\mathbb{S}^1$ and $\mathbb{S}^3$ act via isometries. Thus, the pushforward Riemannian metric on $\mathbb{FP}^n$ is well-defined and turns the projection π into a Riemannian submersion. Put differently, the sphere $\mathbb{S}^n$ (resp. $\mathbb{S}^{2n+1}$, $\mathbb{S}^{4n+3}$) is a smooth fibre bundle over $\mathbb{RP}^n$ (resp. $\mathbb{CP}^n$, $\mathbb{HP}^n$) with the fibre diffeomorphic to $\mathbb{S}^0$ (resp. $\mathbb{S}^1$, $\mathbb{S}^3$) and the differential $d\pi_x$ is an isometry when restricted to the space of horizontal tangent vectors at any x in the sphere and the tangent space at $\pi(x)$ in the projective space. This statement is trivial when $\mathbb{F} = \mathbb{R}$, since in that case we have $\mathbb{S}^0 = \{-1,1\}$ and π is a local diffeomorphism. The other cases are given in [Lee97, Problem 3-8]. If $\mathbb{F} = \mathbb{C}$ this construction yields the well-known Fubini-Study metric on $\mathbb{CP}^n$ [Jos17, Ch. 7.1].

The BM on $\mathbb{FP}^n$ can be defined as a strong Markov process with the generator equal to $\frac{1}{2}\Delta_{\mathbb{FP}^n}$, where $\Delta_{\mathbb{FP}^n}$ is the Laplace-Beltrami operator corresponding to the Riemannian metric on $\mathbb{FP}^n$ [Hsu02, p.74]. Finally, we establish a connection between the spherical BM and the BM on projective spaces.

Lemma B.1. *Let $n \geq 1$ and let Z be BM on sphere $\mathbb{S}^n$ (resp. $\mathbb{S}^{2n+1}$, $\mathbb{S}^{4n+3}$) started at z. Then the process $\pi(Z)$ is the BM on $\mathbb{RP}^n$ (resp. $\mathbb{CP}^n$, $\mathbb{HP}^n$) started at $\pi(z)$.*

Proof. Fibres of the Riemannian submersion π are orbits in $\mathbb{S}^n \subseteq \mathbb{R}^{n+1}$ (resp. $\mathbb{S}^{2n+1} \subseteq \mathbb{C}^{n+1}$, $\mathbb{S}^{4n+3} \subseteq \mathbb{H}^{n+1}$) of the (right) group action $((z_0, \ldots, z_n), \lambda) \mapsto (z_0\lambda, \ldots, z_n\lambda)$, where $(z_0, \ldots, z_n)^\top \in \mathbb{S}^n$ (resp. $\mathbb{S}^{2n+1}$, $\mathbb{S}^{4n+3}$) and λ is an element of multiplicative group of unit elements in $\mathbb{R}$ (resp. $\mathbb{C}$, $\mathbb{H}$). This group is isomorphic to $\mathbb{S}^0$ (resp. $\mathbb{S}^1$, $\mathbb{S}^3$) and isometric to the submanifold $\{(z_0, \ldots, z_n)^\top \in \mathbb{S}^n$ (resp. $\mathbb{S}^{2n+1}$, $\mathbb{S}^{4n+3}$); $z_1 = \cdots = z_n = 0\}$ in $\mathbb{S}^n$ (resp. $\mathbb{S}^{2n+1}$, $\mathbb{S}^{4n+3}$). This submanifold is easily seen to be totally geodesic i.e. any geodesic in the submanifold is also a geodesic in the ambient manifold. Moreover, each fibre (i.e. orbit) can be isometrically mapped onto this submanifold by multiplication from the left by suitable element of special orthogonal group (for any element x in the fibre pick any matrix $O(x)$ in the special orthogonal group such that $O(x)x = (1,0,\ldots,0)^\top$). Hence all fibres are totally geodesic submanifolds.

Let (M,g) be a k-dimensional Riemannian submanifold of a Riemannian manifold $(\widetilde{M}, \widetilde{g})$, i.e. the metric g is a restriction of $\widetilde{g}$ to the tangent bundle of M. For any two vector fields X,Y on M, the second fundamental form II is given by $II(X,Y) = \widetilde{\Delta}_{\widetilde{X}}\widetilde{Y} - \Delta_X Y$, where Δ (resp. $\widetilde{\Delta}$) represent the metric connections on M (resp. $\widetilde{M}$) and the vector fields $\widetilde{X}, \widetilde{Y}$ are arbitrary extensions of the vector fields X,Y to $\widetilde{M}$. The second fundamental form is well-defined and takes values in the normal bundle of the submanifold M (see [Lee97, Thm 8.2]). By [Lee97, Exercise 8.4], the submanifold M is totally geodesic if and only if $II(X,Y) = 0$ for any vector fields X,Y on M. The mean curvature H_x at any $x \in M$ is proportional to the metric trace of II, i.e. for any orthonormal frame $e_1, \ldots, e_k$ in the neighbourhood of x we have $H_x = \frac{1}{k}\sum_{i=1}^k II(e_i, e_i)_x$. Clearly, H_x is equal to 0 for each $x \in M$ if M is totally geodesic in $\widetilde{M}$.

By [Pau90, Thm 1], the projection $\pi(Z)$ in $\mathbb{RP}^n$ (resp. $\mathbb{CP}^n$, $\mathbb{HP}^n$) of the BM Z on $\mathbb{S}^n$ (resp. $\mathbb{S}^{2n+1}$, $\mathbb{S}^{4n+3}$) is a BM with the drift given by $V_x = -\frac{n}{2}d\pi_x(H_x)$ (resp. $-\frac{2n}{2}d\pi_x(H_x)$, $-\frac{4n}{2}d\pi_x(H_x)$), where H_x is the mean curvature of the fibre of x in the ambient sphere. Since all the fibres are totally geodesic, the drift vanishes and the lemma follows. $\qquad\square$

Acknowledgements

References

[CEE10] Tobias Carlsson, Tobias Ekholm, and Christer Elvingson, *Algorithm for generating a Brownian motion on a sphere*, Journal of Physics A: Mathematical and Theoretical **43** (2010), no. 50, 505001.

[GMW18] N. Georgiou, A. Mijatović, and A. R. Wade, *Invariance principle for non-homogeneous random walks*, ArXiv e-prints, 1801.07882 (2018).

[GSS12] A. Ghosh, J. Samuel, and S. Sinha, *A "Gaussian" for diffusion on the sphere*, EPL (Europhysics Letters) **98** (2012), no. 3, 30003.

[Hsu02] E.P. Hsu, *Stochastic analysis on manifolds*, Graduate Studies in Mathematics 38, vol. 38, American Mathematical Society, 2002.

[IM96] K. Ito and H.P. McKean, *Diffusion Processes and their Sample Paths*, Classics in Mathematics, Springer, 1996.

[Jos17] J. Jost, *Riemannian Geometry and Geometric Analysis*, 7 ed., Universitext, Springer International Publishing, 2017.

[JS17] P.A. Jenkins and D. Spanò, *Exact simulation of the Wright–Fisher diffusion*, Ann. Appl. Probab. **27** (2017), no. 3, 1478–1509.

[KDPN00] M. M. G. Krishna, Ranjan Das, N. Periasamy, and Rajaram Nityananda, *Translational diffusion of fluorescent probes on a sphere: Monte Carlo simulations, theory, and fluorescence anisotropy experiment*, The Journal of Chemical Physics **112** (2000), no. 19, 8502–8514.

[Lee97] J.M. Lee, *Riemannian Manifolds: An Introduction to Curvature*, 1st ed., Graduate Texts in Mathematics, Springer, 1997.

[Lee12] _____, *Introduction to Smooth Manifolds*, 2nd ed., Graduate Texts in Mathematics 218, Springer-Verlag New York, 2012.

[LTT08] Guanglai Li, Lick-Kong Tam, and Jay X. Tang, *Amplified effect of Brownian motion in bacterial near-surface swimming*, Proceedings of the National Academy of Sciences **105** (2008), no. 47, 18355–18359.

[MMU18] A. Mijatović, V. Mramor, and G. Uribe Bravo, *Projections of spherical Brownian motion*, Electron. Commun. Probab. **23** (2018), 12 pp.

[NEE03] J. Nissfolk, T. Ekholm, and C. Elvingson, *Brownian dynamics simulations on a hypersphere in 4-space*, The Journal of Chemical Physics **119** (2003), no. 13, 6423–6432.

[Pau90] J.E Pauwels, *Riemannian submersions of Brownian motions*, Stochastics and Stochastic Reports **29** (1990), no. 4, 425–436.

[PR88] E.J. Pauwels and L.C.G. Rogers, *Skew-Product Decompositions of Brownian motions*, Contemporary Math. **73** (1988), 237–262.

[RY99] D. Revuz and M. Yor, *Continuous martingales and Brownian motion*, 3rd ed., Springer, 1999.

[WY98] J. Warren and M. Yor, *The Brownian Burglar: conditioning Brownian motion by its local time process*, Séminaire de Probabilités XXXII (Berlin, Heidelberg) (Jacques Azéma, Marc Yor, Michel Émery, and Michel Ledoux, eds.), Springer Berlin Heidelberg, 1998, pp. 328–342.

Department of Statistics, University of Warwick, & The Alan Turing Institute, UK
Email address: a.mijatovic@warwick.ac.uk

Department of Statistics, University of Warwick, & The Alan Turing Institute, UK
Email address: veno.mramor@warwick.ac.uk

Instituto de Matematicas, Universidad Nacional Autónoma de México, México
Email address: geronimo@matem.unam.mx