Electron. Commun. Probab. **25** (2020), no. 29, 1–12.
https://doi.org/10.1214/20-ECP299
ISSN: 1083-589X

ELECTRONIC COMMUNICATIONS in PROBABILITY

Extensions of Brownian motion to a family of Grushin-type singularities*

Ugo Boscain[†] Robert W. Neel[‡]

Abstract

We consider a one-parameter family of Grushin-type singularities on surfaces, and discuss the possible diffusions that extend Brownian motion to the singularity. This gives a quick proof and clear intuition for the fact that heat can only cross the singularity for an intermediate range of the parameter. When crossing is possible and the singularity consists of one point, we give a complete description of these diffusions, and we describe a "best" extension, which respects the isometry group of the surface and also realizes the unique symmetric one-point extension of the Brownian motion, in the sense of Chen-Fukushima. This extension, however, does not correspond to the bridging extension, which was introduced by Boscain-Prandi, when they previously considered self-adjoint extensions of the Laplace-Beltrami operator on the Riemannian part for these surfaces. We clarify that several of the extensions they considered induce diffusions that are carried by the Martin compactification at the singularity, which is much larger than the (one-point) metric completion. In the case when the singularity is more than one-point, a complete classification of diffusions extending Brownian motion would be unwieldy. Nonetheless, we again describe a "best" extension which respects the isometry group, and in this case, this diffusion corresponds to the bridging extension. A prominent role is played by Bessel processes (of every real dimension) and the classical theory of one-dimensional diffusions and their boundary conditions.

Keywords: Grushin plane; Bessel processes; self-adjointness of the Laplace-Beltrami operator.
AMS MSC 2010: 58J65; 53C17.
Submitted to ECP on October 5, 2019, final version accepted on February 15, 2020.
Supersedes arXiv:1910.02256.
Supersedes HAL:hal-02394958.

1 Introduction

Consider the (open) Riemannian manifold (M, g) where $M = (\mathbf{R} \setminus \{0\}) \times \mathbb{T}$ (here $\mathbb{T}$ is the one-dimensional torus), and

$$g = dx^2 + |x|^{-2\alpha} d\theta^2, \text{ that is, in matrix notation } g = \begin{pmatrix} 1 & 0 \\ 0 & |x|^{-2\alpha} \end{pmatrix}. \tag{1.1}$$

*The first author was supported by the ANR project SRGI ANR-15-CE40-0018 and by the ANR project Quaco ANR-17-CE40-0007-01. The second author was partially supported by grant #524713 from the Simons Foundation and by the National Security Agency under Grant Number H98230-15-1-0171.

[†]CNRS, Sorbonne Université, Inria, Université de Paris, Laboratoire Jacques-Louis Lions, Paris, France.
E-mail: ugo.boscain@upmc.fr

[‡]Department of Mathematics, Lehigh University, Bethlehem, PA, USA.
E-mail: robert.neel@lehigh.edu

Here $x \in \mathbf{R} \setminus \{0\}$, $\theta \in \mathbb{T}$ and $\alpha \in \mathbf{R}$ is a parameter. An orthonormal frame for the metric (1.1) is given by the pair of vector fields

$$X = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 \\ |x|^\alpha \end{pmatrix}.$$

Define

$$M_{\text{cylinder}} = \mathbf{R} \times \mathbb{T}, \qquad M_{\text{cone}} = M_{\text{cylinder}} / \sim,$$

where $(x_1, \theta_1) \sim (x_2, \theta_2)$ if and only if $x_1 = x_2 = 0$.

When $\alpha \geq 0$, extending the vector fields X and Y to M_{cylinder}, the natural control-theoretic notion of the length of a curve shows that there are paths of finite length between $M^+ = \{x > 0\} \times \mathbb{T}$ and $M^- = \{x < 0\} \times \mathbb{T}$ (which are the Riemannian geodesics for $x \neq 0$ and which are tangent to X when $x = 0$), and, as explained in [5], this extended distance makes M_{cylinder} into a metric space (and a length space) in a way that induces on M_{cylinder} its original topology. Similarly, when $\alpha < 0$ the distance induced by the metric (1.1) extend naturally to M_{cone}. Said differently, M_{cylinder} (for $\alpha \geq 0$) and M_{cone} (for $\alpha < 0$) give the metric compactifications of M with respect to this distance (at the singularity). We denote these metric spaces by M_α ($M_\alpha = M_{\text{cylinder}}$ if $\alpha \geq 0$ and $M_\alpha = M_{\text{cone}}$ if $\alpha < 0$) and note that M is the Riemannian subset, while we let $\mathcal{Z} = M_\alpha \setminus M = \{x = 0\}$ be the singular set (terminology which will be further justified in a moment).

This construction gives a one-parameter family of natural singularity models relevant to rank-varying sub-Riemannian/almost-Riemannian geometry that includes the well-known Grushin cylinder (the obvious quotient of the Grushin plane), when $\alpha = 1$. Moreover, the case $\alpha \geq 1$ corresponds to an almost-Riemannian structure in the sense of [2, 4] and of [1, Chapter 9].

However, even though the metric (and length minimizing curves) extends across the singularity, the Riemannian metric g (except for $\alpha = 0$, which gives a standard cylinder) is singular on $\mathcal{Z}$, and this is also the case for the Riemannian volume ω and for the Laplace-Beltrami operator, that take the form

$$\omega = \sqrt{\det g}\, dx\, d\theta = |x|^{-\alpha} dx\, d\theta, \text{ and}$$
$$\Delta = \frac{1}{\sqrt{\det g}} \sum_{j,k=1}^{2} \partial_j \left(\sqrt{\det g}\, g^{jk} \partial_k \right) = \partial_x^2 + |x|^{2\alpha} \partial_\theta^2 u - \frac{\alpha}{x} \partial_x.$$

Thus, if one wishes to consider the heat equation or Schrödinger equation on M, or to consider Brownian motion, one must consider the behavior at the boundary. From the perspective of functional analysis, this means considering self-adjoint extensions of the Laplacian on M. Indeed, in [5], the following two basic results were proven.

Theorem 1.1 ([5]). *The operator $\Delta|_{C_c^\infty(M)}$ is essentially self-adjoint in $L^2(M, \omega)$ if and only if $\alpha \in (-\infty, -3] \cup [1, \infty)$.*

An immediate consequence of Theorem 1.1 is that for $\alpha \in (-\infty, -3] \cup [1, \infty)$ the only self-adjoint extension of Δ is the Friedrich extension Δ_F.

If $\alpha \notin (-\infty, -3] \cup [1, \infty)$ the next theorem gives some additional information.

Theorem 1.2 ([5]). *Let $\widehat{\Delta}$ be the Fourier transform of Δ in the variable θ. We have the following*

- *if $\alpha \in (-3, -1]$, only the first Fourier component of $\widehat{\Delta}$ is not essentially self-adjoint.*
- *if $\alpha \in (-1, 1)$, all the Fourier components of $\widehat{\Delta}$ are not essentially self-adjoint.*

Notice that the essential self-adjointness of a Fourier component of $\hat{\Delta}$ means that such an operator can be naturally and uniquely extended (by taking its closure) to a self-adjoint operator without adding any additional boundary conditions. Moreover, one extension that is always possible is to take the Friedrichs extensions for the operators defined on M^+ and M^- and to "concatenate" them. This extension is the one which separates the dynamics. Hence, the essential self-adjointness of a Fourier component of $\hat{\Delta}$ means that its unique self-adjoint extension is the one that separates the dynamics between M^+ and M^-. On the contrary, if a Fourier component of $\hat{\Delta}$ is not essentially self-adjoint, one can construct a self-adjoint extension which connects the boundary conditions, and thus the dynamics, on the two sides.

We now focus on the heat equation $\partial_t u = \Delta u$ on $L^2(M, \omega)$. When only the first Fourier component of $\hat{\Delta}$ is not essentially self-adjoint, a self-adjoint extension of Δ permits at most the first Fourier component of u (i.e. the average in the variable θ) to flow from M^+ to M^-. When all Fourier components of $\hat{\Delta}$ are not essentially self-adjoint, we have "full communication" between M^+ and M^- in the sense that one can construct self-adjoint extensions permitting all Fourier components of u to flow from M^+ to M^-. We talk about "maximal communication" when we choose a self-adjoint extension that, on each Fourier component, put the least possible constraints on the boundary conditions and that connects them on the the two sides (see below for an explicit formula).

For the heat equation, a consequence of Theorem 1.2 is that

- when $\alpha \in (-3, -1]$ there are self-adjoint extensions of Δ that permit only the average over $\mathbb{T}$ of u to flow through $\mathcal{Z}$. However, as explained in [5], the only Markovian extension of Δ is Δ_F, which does not permit any communication between M^+ and M^-.
- when $\alpha \in (-1, 1)$ there are self-adjoint extensions of Δ that permit full communication between M^+ and M^- and that are Markovian. In particular, there is a self-adjoint extension called the *bridging extension* realizing the maximal communication between the two sides, the domain of which is

$$\left\{ H^2\left(M_\alpha, \omega\right) \,\middle|\, u(0^+, \cdot) = u(0^-, \cdot),\ \lim_{x \to 0^+} |x|^{-\alpha} \partial_x u(x, \cdot) = \lim_{x \to 0^-} |x|^{-\alpha} \partial_x u(x, \cdot) \right\},$$

where $H^2\left(M_\alpha, \omega\right) = \{u \in L^2(M, \omega),\ \ |\nabla u|, \Delta u \in L^2(M, \omega)\}$.

The purpose of the present note is to consider diffusions on M_α that extend Brownian motion on M. (By a diffusion, we mean a strong Markov process with continuous paths.) One aspect of this is to give the path properties that correspond to many of the above results. For example, the fact that there are Markov extensions allowing communication between M^+ and M^- exactly when $\alpha \in (-1, 1)$ corresponds to the fact that for $\alpha < -1$, Brownian motion on M never hits the singularity, and thus cannot cross it, while for $\alpha > 1$, is an exit-only boundary for M (essentially in the sense of the classical Feller classification), so the process must be absorbed at the singularity (assuming it is conservative and cannot be killed) and thus also cannot cross it. In particular, the x-marginal of Brownian motion on M is given by a Bessel process of dimension $d = 1 - \alpha$, so both the behavior of the process at the singularity, as well as the stochastic completeness near infinity, follow, and thus stochastic methods provide elementary proofs and intuition for these results. From the other side, this one-parameter family of geometries provides examples "in nature" of Bessel processes of all real dimensions (such examples also arise in SLE, but not naturally in Riemannian geometry). This material is treated in Section 2. When considering the possible extensions for $\alpha \in (-1, 1)$, it is important to note that the Martin compactification of M at the singularity (in what follows, we are concerned with the behavior at the singularity, and thus all of our compactifications are done there, ignoring what happens near infinity, since there the structure is Riemannian) is larger than $\mathcal{Z}$, that is, larger than the metric compactification (at the singularity). Various self-adjoint extensions of Δ mentioned above are carried by the Martin boundary (at the singularity). For example, Neumann boundary conditions make the process undergo instantaneous normal reflection at the singularity, back into the component of M it came from. But such an extension clearly cannot descend to a strong Markov process on M_α. Thus, here we treat extensions that are carried by M_α itself, so that our results differ from, and complement, those of [5]. We also do not restrict our attention to symmetric extensions. (In this connection, it is worth mentioning that we do not treat non-Markov extensions, so we have no contribution to the above results for $\alpha \in (-3,-1]$.)

In the case $\alpha \in (-1,0)$ when M_α is a topological cone, we are able to give a complete description of (conservative) diffusions on M_α the extend Brownian motion on M. It is worth noting that these correspond to one-point extensions in the sense of Chen and Fukushima (see [6, 7]), and we identify the unique symmetric extension spending zero time at the singularity. For this extension, we see that only the average over $\mathbb{T}$ of a function flows through the singularity under the corresponding semigroup (note that in this case, the bridging extension does not correspond to a diffusion on M_α). Thus the same phenomenon observed in [5] for non-Markov self-adjoint extensions of Δ for $\alpha \in (-3,-1]$ is replicated here for Markov processes that respect the topology of M_α when $\alpha \in (-1,0)$. This is carried out in Section 3.

In the case $\alpha \in [0,1)$ when M is a topological cylinder, the larger singular set makes a complete classification of extensions complicated. However, one can describe the basic features, and we also construct the unique diffusion spending time 0 at $\mathcal{Z}$ and respecting the symmetries of M_α, and show that in this case it corresponds to the bridging extension. This comprises Section 4.

2 Bessel processes, stochastic completeness, and boundary conditions

From the above description of the metric and the induced Laplacian on M, we see that, in the (x,θ) coordinates, Brownian motion evolves by the system of SDEs

$$\begin{aligned} dx_t &= dW_t^1 - \frac{\alpha}{2x}\,dt \\ d\theta_t &= |x|^\alpha \; dW_t^2, \end{aligned} \tag{2.1}$$

at least until T_0, the first hitting time of $\{x=0\}=\mathcal{Z}$, where W_t^1 and W_t^2 are independent one-dimensional Brownian motions. (The SDE for x_t should be understood as giving a local semi-martingale on $(0,\infty)$ for general real α, but the extension until T_0 is standard, say, by squared-Bessel processes as mentioned below, or by explicit construction of the transition density, etc.) It is the x_t process that mainly interests us in this section. Note that its evolution does not depend on θ_t (except possibly on the singular set), so that the situation reduces to a one-dimensional problem. Moreover, observe that, for $x_t > 0$ (equivalently, on M^+), the SDE satisfied by x_t is just that of a Bessel process of dimension $d = 1-\alpha$ (and for $x_t < 0$, it is just -1 times such a process).

Theorem 2.1. *Let (x_t,θ_t) be a diffusion on M_α extending Brownian motion on M, with no killing (at $\mathcal{Z}$, or anywhere else). Then the diffusion a.s. does not explode in finite time (making M_α stochastically complete). Moreover,*

1. *If $\alpha \le -1$, then (x_t,θ_t) a.s. does not hit $\mathcal{Z}$ if $(x_0,\theta_0) \notin \mathcal{Z}$, but it is possible for the diffusion to enter M from $\mathcal{Z}$. Thus (x_t,θ_t) is uniquely determined, except when $(x_0,\theta_0) \in \mathcal{Z}$, when we must specify the entrance law.*

2. *If $\alpha \in (-1,1)$, then (x_t,θ_t) a.s. hits $\mathcal{Z}$ in finite time, but is then able to re-enter M. Thus (x_t,θ_t) is determined by the choice of behavior at $\mathcal{Z}$.*
3. *If $\alpha \geq 1$, then (x_t,θ_t) a.s. hits $\mathcal{Z}$ in finite time, and is necessarily absorbed by $\mathcal{Z}$.*

Proof. Taking advantage of the reflection symmetry in x, consider the process $z_t = x_t^2$ on $[0,\infty)$, so that $\mathcal{Z}$ corresponds to an included boundary point. Then z_t is a squared Bessel process of dimension $1-\alpha$, which is a true semi-martingale for any value of α and which satisfies the SDE $dz_t = 2\sqrt{z_t}\,dW_t + (1-\alpha)\,dt$ until T_0. Moreover, the behavior of (squared) Bessel processes is well-understood. In particular, the process doesn't explode to infinity (in finite time) for any value of α. For $\alpha \leq -1$, 0 is an entrance-only boundary (in the standard Feller classification for one-dimensional diffusions). For $\alpha \in (-1,1)$, 0 is a regular boundary, and thus one needs to specify boundary conditions. For $\alpha \geq 1$, 0 is an exit-only boundary, and thus the process must be absorbed at 0, since we don't allow killing. Since both T_0 and any potential explosion of the process depend only on x_t, the theorem is a restatement of the above. □

(Note that the phrasing of this theorem contrasts slightly with [5], since they consider only self-adjoint extensions and thus kill the process at $\mathcal{Z}$ when it is an exit-only boundary, instead of letting it be absorbed. But the real point is that the process never explodes to infinity in finite time.) This not only recovers the stochastic completeness and difference in (Markov) extensions of Δ depending on whether $\alpha \in (-1,1)$ or not from [5], but also allows other properties of the heat flow on M_α to deduced from known properties of Bessel processes. For example, for the Grushin cylinder (or Grushin plane), which corresponds to $\alpha = 1$, the rate at which heat is absorbed at $\mathcal{Z}$ is given by the transition measure for a 0-dimensional Bessel process (see Section A.2 of [8], for example).

The above makes it clear that the only case which requires further exploration is $\alpha \in (-1,1)$ (apart from the entrance law when $\alpha \leq -1$ and $(x_0,\theta_0) \in \mathcal{Z}$ which amounts to a simpler version of the $\alpha \in (-1,0)$ case— see Section 3.3).

3 The case $-1 < \alpha < 0$

When $-1 < \alpha < 0$, M_α has a cone structure at the singularity, and the singularity reduces to a single point. As we saw above, the singularity is a regular boundary for the $|x_t|$ process, so that any diffusion a.s. hits the singular point in finite time, and it is possible for the diffusion to leave the singular point. In this situation, there are many possible diffusions extending Brownian motion on M to all of M_α, but the singularity is simple enough that we can describe them all.

3.1 Classification of diffusions on M_α

Because the singularity is a single point, the behavior of the diffusion at the singularity doesn't depend on the θ_t process, which means that x_t is a one-dimensional diffusion. As mentioned, the theory of one-dimensional diffusions is completely understood (see, for example, [11]). Thus we can give all possible (conservative) diffusion extending x_t in this case, and this is the first step in determining the possible diffusions on M_α. Essentially, the possible x_t-diffusions depend on two parameters, the degree of "stickiness" at 0 and the skewness at 0. We also note that x_t is a Bessel process of dimension between 1 and 2 (at least until it hits 0, at which point we don't necessarily instantaneously reflect it), and thus x_t is a semi-martingale.

As the above suggests, here we take the perspective that $x=0$ gives an interior singular point of the diffusion. To make the connection to the approach of the previous section clear, we note that (the law of) z_t determines (the law of) x_t up to the sign of each excursion of x_t away from 0, that is, away from $\mathcal{Z}$.

We begin by recalling some basics of one-dimensional diffusions, in order to describe the possible diffusions that extend x_t from $\mathbb{R} \setminus \{0\}$ to $\mathbb{R}$. In the classification scheme of Itô and McKean [10], 0 can be a regular point, a left or right shunt, or a trap. Most interesting for us is when 0 is a regular point, so that the process can cross 0 in either direction. In this case, the x_t-diffusion is determined by its scale function $s(x)$ and its speed measure m. Further, the diffusion must agree with the appropriate Bessel process on $\mathbb{R} \setminus \{0\}$. Hence the scale function is determined up to affine transformations on each of $\{x < 0\}$ and $\{x > 0\}$, subject to the additional constraint that it is continuous. We normalize s by translation so that $s(0) = 0$. Starting from the "standard" scale function $s(x) = x^{\alpha+1}$ for a Bessel process of dimension $1 - \alpha \in (1, 2)$ (see Section 11.1 of [14], also for the speed measure of a Bessel process which we are about to use), we see that the most general normalized scale function for x_t with $s(0) = 0$ is

$$s(x) = \begin{cases} -a(-x)^{1+\alpha} & \text{for } x < 0 \\ (1-a)x^{1+\alpha} & \text{for } x \geq 0 \end{cases}$$

for $0 < a < 1$. Here we see that a gives the skewness of x_t at 0, in the sense that, for any $y > 0$,

$$\mathbb{P}\left(x_t \text{ starting from 0 hits } y \text{ before } -y\right) = a.$$

Continuing, the speed measure is uniquely determined on $\mathbb{R} \setminus \{0\}$ by the speed measure of a Bessel process and the above choice of scaling function (as having density $2/s'$ with respect to Lebesgue measure), so that the most general speed measure for x_t is

$$m = \frac{2}{a(1+\alpha)(-x)^{\alpha}} \mathbf{1}_{\{x<0\}}\, dx + \gamma\delta_0 + \frac{2}{(1-a)(1+\alpha)x^{\alpha}} \mathbf{1}_{\{x>0\}}\, dx$$

for some $\gamma \in [0, \infty)$, where dx denotes Lebesgue measure and δ_0 a point mass at $x = 0$. Here γ gives the degree of "stickiness" at 0, in the sense that if $\gamma = 0$, the set $\{t > 0 : x_t = 0\}$ almost surely has Lebesgue measure 0, whereas if $\gamma > 0$, this set has positive measure. Note that for $\gamma = 0$ and $0 < a < 1$, x_t will be a time-changed skew Brownian motion (see the survey [12] for a detailed introduction to skew Brownian motion).

A perhaps more appealing (and slightly more general) way to describe x_t is as follows. While 0 is a regular point for $0 < a < 1$, 0 is a left shunt if $a = 0$, and a right shunt if $a = 1$. (If $\gamma = \infty$, then 0 is a trap, viewed as an interior point of M_α rather than as a boundary point of $[0, \infty)$ as above.)

Next, we consider the θ-process. Recall that this is an $\mathbb{T}$-valued process that satisfies the SDE $d\theta_t = |x|^\alpha dW_t^2$ for $x \neq 0$. Suppose $x_0 > 0$, and make the change of variables $y(x) = \frac{1}{\alpha+1} x^{\alpha+1}$, which puts x_t on its natural scale. Recall that T_0 is the first hitting time of 0 for x_t (and thus the first time the process on M_α hits the singularity). Then the process satisfies the system of SDEs

$$\begin{aligned} dy_t &= (1+\alpha)^{\alpha/(1+\alpha)} y^{\alpha/(1+\alpha)}\, dW_t^1 \\ d\theta_t &= (1+\alpha)^{\alpha/(1+\alpha)} y^{\alpha/(1+\alpha)}\, dW_t^2, \end{aligned} \tag{3.1}$$

on the time interval $[0, T_0]$. Note that (y_t, θ_t) is a time-changed Brownian motion on $(\mathbb{R} \setminus \{0\}) \times \mathbb{T}$, and that the case when $x_0 < 0$ is given by reflection.

With these preliminaries, we can now state and prove the following classification of diffusions on M_α (in the present case of $-1 < \alpha < 0$) that extend Brownian motion on M.

Theorem 3.1. *Let M_α be as above for $-1 < \alpha < 0$, and let (x_t, θ_t) be a (conservative) diffusion on M_α extending Brownian motion on M, written in the standard coordinates. Then (x_t, θ_t) is determined by its behavior starting from $\mathcal{Z}$, which is given by the following* *parameters: $\gamma \in [0,\infty]$, $a \in [0,1]$, and (Borel) probability measures μ^+ and μ^- on $\mathbb{T}$. More concretely, x_t^2 is a squared-Bessel process of dimension $1-\alpha$ on $[0,\infty)$ with reflecting boundary condition at 0 determined by γ (making 0 instantaneously reflecting, slowly reflecting, or absorbing, as above), x_t is recovered from x_t^2 by assigning each excursion away from 0 a positive sign with probability a (and thus a negative sign with probability $1-a$), and on any excursion $t \in (t_1,t_2)$ of x_t away from 0, θ_t is given as the solution of the SDE $d\theta_t = |x_t|^\alpha \, dW_t^2$ with the initial condition θ_{t_1} distributed as μ^+ if x_t is positive on (t_1,t_2) and μ^- if x_t is negative on (t_1,t_2). Moreover, if $\gamma=\infty$, $\{x=0\}$ is absorbing and none of the other parameters are relevant, if $\gamma<\infty$ and $a=0$, all excursions of x_t are negative and μ^+ is irrelevant, and if $\gamma<\infty$ and $a=1$, all excursions of x_t are positive and μ^- is irrelevant, but aside from these exceptions, there is a one-to-one correspondence between the choice of diffusion and the choice of parameters.*

Proof. Let (x_t,θ_t) be such a diffusion. Then the process is uniquely determined until T_0 (the first hitting time of the singularity), so by the strong Markov property, such diffusions are determined by their behavior starting from the singularity.

From (3.1), we see that y_t is a time-changed Brownian motion that a.s. hits 0 in finite time, and thus it a.s. accumulates finite quadratic variation $\int_0^{T_0} (1+\alpha)^{2\alpha/(1+\alpha)} y^{2\alpha/(1+\alpha)} \, dt$ over $[0,T_0]$. Since the quadratic variation of θ_t on $[0,T_0]$ is equal to that of y_t, it is also a.s. finite. Further, θ_t is a martingale and thus a time-changed Brownian motion, and it follows that θ_t a.s. has a limit as $t \nearrow T_0$. (In particular, the existence of an exit angle from M implies that the invariant sigma-algebra of this stopped process, and thus also the Martin boundary of M, is non-trivial, even though the process converges to a single point on M_α as $t \nearrow T_0$. We discuss this further below.)

We can now consider how the process leaves the singularity. We have already seen that x_t is determined by γ and a, independent of θ_t (and the description in terms of x_t^2 and the sign of the excursions of x_t is equivalent), so we need only consider θ_t. Because the law of the Brownian excursion is preserved by time-reversal, it follows from the above that y_t a.s. accumulates finite quadratic variation on each excursion from 0, and thus θ_t does as well. Hence, starting from the singularity, θ_t must have a limit as $t \searrow 0$, that is, the process leaves the singularity, and enters $M = M \setminus \{x=0\}$, with an entrance angle. Thus, let μ^+ and μ^- be two probability measures on $\mathbb{T}$. If an excursion of x_t has positive sign, the entrance angle of θ_t is distributed according to μ^+, and similarly for negative excursions and μ^-. (Equivalently, a, μ^+, and μ^- determine a probability measure on $\{-1,1\}\times\mathbb{T}$ that gives the entrance behavior of the diffusion from the singularity.) Moreover, since $\theta_t = \theta_0 + \int_0^t |x_s|^\alpha \; dW_s^2$ on the excursion $[0,T_0]$, we see that the θ-process is completely determined by this entrance behavior. Hence any such diffusion is uniquely specified by the data described in the theorem.

Conversely, for any data as in the theorem, there is a corresponding diffusion. If $\gamma=\infty$, this is immediate, so assume $\gamma<\infty$. Then the construction of a diffusion matching these parameters is a straightforward variant of the argument in [3] for Walsh's Brownian motion (also see the Remark just before Section 3). In particular, suppose we show that the simplified process (y_t',θ_t') with $\gamma=0$ and time-changed so that (y_t',θ_t') is a Brownian motion on $(\mathbb{R}\setminus\{0\})\times\mathbb{T}$ is a diffusion. Then the desired process with general $\gamma\in[0,\infty)$ can be constructed from the simplified process via time-change by a continuous additive functional, and such time-changes preserve diffusions. Hence it is enough to show that, for any a, μ^+, and μ^-, the associated simplified process (y_t',θ_t') is a diffusion.

But the semigroup P_T for such a process (or equivalently, the transition density) can be explicitly given in terms of well-known objects. In particular, let $p_T^0(y_0,\theta_0;y,\theta)$ be the density of ($\mathbb{R}^2$-) Brownian motion from (y_0,θ_0) killed at $\{y=0\}$ with respect to $dy\,d\theta$ (which is given by the method of images), let $T_{0,y_0}(t)$ be the distribution of T_0 for Brownian motion started from (y_0, θ_0) (which has a density with respect to dt given by the classical first passage time), let $p_s^{\mathbb{T}}(\varphi;\theta)$ be the density (with respect to $d\theta$) of Brownian motion on $\mathbb{T}$ started from φ, let $G_t(s)$ be the distribution of the last zero of a 1-dimensional Brownian motion starting from the origin on the time interval $[0,t]$ (which has density with respect to ds given by Levy's arcsine law, see Section 7 of [13]), and let $p_s^{\mathrm{me}}(y)$ be the density (with respect to dy) of the position of a Brownian meander in time s (again see Section 7 of [13]). We also let $\mu(\varphi,\sigma)$ be the measure on $\mathbb{T}^1 \times \{-1,1\}$ induced by a, μ^+ and μ^-, which encodes the entrance law from $\mathcal{Z}$. Then for $f \in C_0(M_\alpha)$,

$$
\begin{aligned}
P_T f(y_0,\theta_0) &= \int_{y\in\mathbb{R}}\int_{\theta\in\mathbb{T}} p_T^0(y_0,\theta_0;y,\theta) f(y,\theta)\,d\theta\,dy + \\
&\int_{t=0}^{T}\int_{s=0}^{T-t}\int_{\substack{\varphi\in\mathbb{T}\\ \sigma\in\{-1,1\}}}\int_{y=0}^{\infty}\int_{\theta\in\mathbb{T}} p_s^{\mathrm{me}}(y) p_s^{\mathbb{T}}(\varphi;\theta) f(\sigma y,\theta)\,d\theta\,dy\,d\mu(\varphi,\sigma)\,dG_{T-t}(s)\,dT_{0,y_0}(t).
\end{aligned}
$$

Then it is an exercise to verify that this semigroup is Feller, so that there is a corresponding strong Markov process. To see that the paths can be taken to be continuous, we can either verify Kolmogorov's condition holds (say, by taking advantage of the fact that y_t is a skew Brownian motion and M_α has the cone topology), or we can observe that it follows from verifying that the process has the desired decomposition in terms of excursions from $\mathcal{Z}$. Either way, with the strong Markov property in hand, we can see that (y'_t, θ'_t) has excursions determined by a, μ^+, and μ^- as desired, which completes the construction. □

Note that while we gave an efficient argument using some general process theory, we do not attempt to develop the corresponding stochastic analysis, as has been pursued for other generalizations of Walsh's Brownian motion, in [9], for example.

3.2 The state space and symmetric extensions

The previous classification was restricted to diffusions on M_α. To clarify, if we start with M, since Brownian motion on M explodes toward $\mathcal{Z} = \{x = 0\}$ in finite time, to continue the process for all time requires enlarging the state space. In order to make the terminology and relationship to other literature clearer, we call a compactification of $M \cap \{-1 \le x \le 1\}$ an interior compactification of M. The idea is that we want to compactify M at the singularity, but this doesn't, in fact, give a compactification, since M has two more ends, corresponding to $x \to \pm\infty$. However, we've seen that Brownian motion on M never escapes out of these ends. Thus we want to restrict our attention to a neighborhood of the singularity, and this is what looking at interior compactifications accomplishes.

One way of enlarging the state space is to add a single point for $\{x = 0\}$, and this gives the metric space M_α that we have been working with, based on the geodesic distance. However, this is not the only possible extension of M. Indeed, starting, more functional analytically, from either the Laplacian or the associated Dirichlet form on M, one can consider extending the domain of the operator beyond smooth functions compactly supported on M. Such extensions are naturally carried by an interior compactification of M, but the compactification will depend on the extension and won't necessarily coincide with the one-point compactification that gives M. This is the approach followed in [5], and we now briefly explain the relationship between their results and the above.

To understand (interior) compactifications of M, it is useful to observe that the coordinates (y, θ) on M, where

$$y = \mathrm{sign}(x)\frac{1}{\alpha+1}|x|^{\alpha+1}$$

and θ, of course, is the same θ from the standard coordinates, give a conformal diffeomorphism from M to the subset of the (Euclidean) cylinder

$$D = ((-\infty, 0) \cup (0, \infty)) \times \mathbb{S}^1 \subset \mathbb{R} \times \mathbb{S}^1.$$

That these coordinates are conformal is contained in (3.1), since it shows that Brownian motion on M is a time-change of (Euclidean) Brownian motion on D.

The maximal extension of the domain of Δ, as described in [5], corresponds to Neumann boundary conditions at the singularity. Unsurprisingly, this corresponds to the "maximal" (interior) compactification (from a potential-theoretic viewpoint) of M, which is the Martin boundary $\partial_M M$ of $M \cap \{-1 \le x \le 1\}$, or equivalently, of $D \cap \{|y| \le 1/(\alpha+1)\}$. (Here we put Neumann boundary conditions on $\{x = \pm 1\}$ for convenience, in order to restrict the process to $M \cap \{-1 \le x \le 1\}$.) More concretely, the conformal equivalence with D shows that $\partial_M M$ can be identified with the "doubled" Euclidean boundary of D, $\{y = 0\}$, which is the disjoint union of two copies of $\mathbb{T}$. Thus $\partial_M M$ records both the exit angle $\lim_{t \nearrow T_0} \theta_t$ and exit "side" $\lim_{t \nearrow T_0} \mathrm{sign}(x_t) = \lim_{t \nearrow T_0} \mathrm{sign}(y_t)$ of the diffusion, where T_0 is the first hitting time of $\{y = 0\}$ (or the first exit time of M) for the diffusion started from a point in M. Then the process associated with the Neumann boundary conditions is determined by instantaneous normal reflection of (y_t, θ_t) back into the component of D the process started in. It's clear how to construct this diffusion analogously to what was done in the previous section, with M_α replaced by $M \cup \partial_M M$, and it's also clear that this process does not descend to a strong Markov process on M_α.

A second extension of Δ considered in [5] is what they call the bridging extension. The corresponding interior compactification is given by identifying pairs of points in $\partial_M M$ with the same θ-coordinate. Since this boundary is also the Euclidean boundary of D, we denote it by $\partial_E M$. The corresponding process can be constructed by starting with the diffusion $(|y_t|, \theta_t)$ with Neumann boundary conditions (as just discussed) but assigning signs to the excursions of $(|y_t|, \theta_t)$ randomly with equal probabilities. To see this, note that the condition $\lim_{x \to 0^+} |x|^{-\alpha} \partial_x u(x, \cdot) = \lim_{x \to 0^-} |x|^{-\alpha} \partial_x u(x, \cdot)$ in the domain of the Laplacian for the bridging extension becomes $\lim_{y \to 0^+} \partial_y u(y, \cdot) = \lim_{y \to 0^-} \partial_y u(y, \cdot)$ after changing coordinates. Then if we consider functions that are even in the y (or x) variable, we see that the boundary condition for the $(|y_t|, \theta_t)$-process is just that the normal derivative vanishes, which corresponds to instantaneous normal reflection. Thus the y_t-process is independent of the θ_t-process, and we see that the domain of the operator is exactly that of skew-Brownian motion in the trivial case when the skewness vanishes; see Equation (7.6.10) of [6] and the surrounding discussion (that is, the process is just Brownian motion, realized in a slightly non-standard way). This justifies the above claim; alternatively, the process can be thought of as Euclidean Brownian motion on $\overline{D}$ time-changed to spend Lebesgue measure 0 time on $\{y = 0\}$ and to solve (3.1) on D. Again, the natural state space for this process is $M \cup \partial_E M$, and it does not induce a strong Markov process on M.

The third explicit extension of Δ considered in [5] is the Friedrich extension, which corresponds to the diffusion killed at the singularity and gives the minimal extension of the domain of Δ. This has M plus a graveyard state for when the particle is killed as its natural state space, which means this case is not covered by Theorem 3.1, so the interior compactification of M is "minimal" as well, but the diffusion is not conservative. (The case when $\mathcal{Z}$ is absorbing has $\mathcal{Z}$ as a stationary state, so it is not symmetric, and thus not considered in [5].)

Recall that ω is the Riemannian volume measure on M, and let $\overline{\omega}$ be the extension of ω to M_α given by assigning measure 0 to the singularity. Since Brownian motion on a Riemannian manifold is symmetric with respect to the Riemannian volume, it is natural to ask for a diffusion on M_α that is symmetric with respect to $\overline{\omega}$. Indeed, both the Neumann and bridging extensions of [5] are symmetric with respect to the extension of ω given by assigning measure 0 to $\partial_M M$ or $\partial_E M$. On the other hand, the diffusions of Theorem 3.1 aren't, in general, symmetric with respect to $\overline{\omega}$. Indeed, this is also a consequence of the following.

Theorem 3.2. *Let M and $\overline{\omega}$ be as above, for $-1 < \alpha < 0$. Then the unique (conservative) diffusion on M_α that extends Brownian motion on M, spends time 0 at $\mathcal{Z}$, and is $\overline{\omega}$-symmetric is given by taking $a = 1/2$, $\gamma = 0$, and μ^+ and μ^- both to be the uniform probability measure on $\mathbb{S}^1$ in Theorem 3.1. This diffusion is also the unique extension of Brownian motion that spends time 0 at $\mathcal{Z}$ and is invariant under the isometry group of M_α.*

Further, let P_t be the semigroup associated to this diffusion, and let f and g be functions in $L^\infty(M_\alpha)$ such that $f = g$ on M^+ and for almost every $u < 0$, $\int_{\mathbb{T}} f(u,\theta)\,d\theta = \int_{\mathbb{T}} g(u,\theta)\,d\theta$. Then $P_t f(x) = P_t g(x)$ for every $x > 0$ and $t > 0$.

Proof. The interior compactification taking M into M_α gives a one-point compactification of $M \cap \{-1 \leq x \leq 1\}$, in the terminology of Chapter 7 of [6] (although what we write as ω and $\overline{\omega}$ correspond to ω_0 and ω, respectively, in their notation). Thus, according to Theorems 7.5.4 and 7.5.6 of [6], there is a unique diffusion on M_α that extends Brownian motion on M and is symmetric with respect to $\overline{\omega}$.

Note that the isometry group of M_α is generated by reflection in x and the action of $\mathrm{SO}(2)$ on θ. Then uniqueness of the $\overline{\omega}$-symmetric extension implies that such a diffusion, when started from the singularity, must be invariant under $\mathbb{Z}/2\mathbb{Z} \times \mathrm{SO}(2)$. Hence, in Theorem 3.1, we must have that $a = 1/2$ and both μ^+ and μ^- are the uniform probability measure on $\mathbb{S}^1$, simply because the action of $\mathbb{Z}/2\mathbb{Z} \times \mathrm{SO}(2)$ is transitive "on the entrance directions" of M. Further, for the process to spend 0 time at the singularity, we must have $\gamma = 0$. This establishes the uniqueness claims.

Finally, the $\mathrm{SO}(2)$ invariance of μ^- and the strong Markov property relative to T_0 imply that $P_t f(x)$ (for $x > 0$) is invariant under letting $\mathrm{SO}(2)$ act on M^-. Hence $P_t f(x)$ is unchanged by replacing f with its θ-integral for each $u < 0$. Applying this to both f and g gives the last result of the theorem. □

Referring to the final result of this theorem, note that, of course, an analogous result holds for $x < 0$ and $u > 0$. Moreover, we interpret this to mean that the diffusion loses information, and only certain "average" features of f are communicated across $\mathcal{Z}$. If μ^- or μ^+ is not uniform, then a similar result holds, except that the θ-averages for each u must be computed with respect to a non-uniform measure (depending on y). In any case, the fact that $\mathcal{Z}$ is a single point means that some information must be lost when the process crosses $\mathcal{Z}$.

3.3 The case $\alpha \leq -1$

In the case when $\alpha \leq -1$, the process can enter M from $\mathcal{Z}$, but then never returns (and never hits $\mathcal{Z}$ if it starts from M). Thus, if we want a diffusion starting from any point of M_α, we need only describe how it enters M from $\mathcal{Z}$ (which is a single point). One possibility is for the process to never leave $\mathcal{Z}$, which one can think of as a type of absorbing boundary condition. If the process leaves, it must do so immediately (by the strong Markov property), and just as above, how it enters M is determined by a choice of $a \in [0,1]$ and probability measures μ^+ and μ^-.

4 The case $0 \le \alpha < 1$

In this case, M_α has a cylinder structure at the singularity. In particular, the singularity is now a circle, naturally parametrized by the θ-coordinate, and this would make a complete description of all possible (conservative) diffusions extending Brownian motion on M rather complicated. For instance, such a description is connected to the boundary theory of multidimensional diffusions. More concretely, consider the process $(x_t^2, \theta_t) = (z_t, \theta_t)$ on $[0, \infty) \times \mathbb{T}$. Then z_t can undergo sticky, oblique reflection at the boundary, with the parameters determining this reflection depending on θ. Determining a solution without assuming (much) regularity of these parameters (or potentially of the boundary) is a longstanding topic of interest. For example, a construction of a process on a halfspace with general Wentzell boundary conditions was given fairly recently by Watanabe [15] by extending Itô's excursion theory, and one can see the references therein for other probabilistic approaches. To extend Brownian motion to M_α, one would expect a "two-sided" version of this type of construction, where the process is potentially sticky at the boundary (and perhaps even diffuses within the boundary) in a way that depends on θ, and when the process re-enters M, the distribution of sign of the excursion depends on θ as does the obliqueness of the "reflection." Constructing such a process, especially for low regularity of the parameters describing this behavior, is well beyond the scope of this note, and it is also in the opposite direction from the more geometrically natural question of determining a "good" or "best" extension.

Before doing this, motivated by the earlier emphasis on whether or not the process can cross the singularity, we give a simple example to illustrate that the way in which the process crosses the singularity can be unusual. Let $A \subset \mathbb{T}$ be a non-empty open subset of $\mathcal{Z}$ such that A^c has non-empty interior. Let $(|x_t|, \theta_t)$ (as a process on $[0, \infty) \times \mathbb{T}$) be given by instantaneous normal reflection at the boundary, and let the sign of each excursion of x_t be positive if it begins in A and negative if it begins in A^c. Then because the process hits both A and A^c with positive probability from either side of $\mathcal{Z}$, we see that the process will (almost surely) cross $\mathcal{Z}$ infinitely often. However, the crossing is "non-local," in the sense that when the process hits the interior of A from M^+, it is distance 0 from M^-, but cannot cross into M^- immediately. Instead it must "go around" A and cross at A^c, and similarly for the process hitting the interior of A^c from M^-.

Returning to the question of a "best" extension, we have the following.

Theorem 4.1. *For $0 \le \alpha < 1$, the only (conservative) diffusion on M_α extending Brownian on M that spends 0 time at $\mathcal{Z}$ and is invariant under the isometry group of M_α is given by letting $(|x_t|, \theta_t)$ be the diffusion on $[0, \infty) \times \mathbb{T}$ that undergoes instantaneous normal reflection at the boundary and letting x_t be constructed from $|x_t|$ by giving each excursion a positive or negative sign with probability $1/2$. Moreover, this is the diffusion associated to the bridging extension.*

Proof. As before, the isometry group of M_α is $\mathbb{Z}/2\mathbb{Z} \times \mathrm{SO}(2)$. Then we note that if the process spends time 0 at $\mathcal{Z}$ and is symmetric with respect to reflection in x, $(|x_t|, \theta_t)$ must be a diffusion on $[0, \infty) \times \mathbb{T}$ that reflects instantaneously at the boundary, and x_t can be recovered from $|x_t|$ by giving each excursion a positive or negative sign with probability $1/2$. Additionally, if $(|x_t|, \theta_t)$ is invariant with respect to the $\mathrm{SO}(2)$ action, the reflection must be normal. Further, the conformal map described in Section 3.2 and given by Equation (3.1) (combined with reflection in y) extends to the current case of $0 \le \alpha < 1$. Thus we see that the Martin boundary is the same as before, although now it is only "twice" the singularity, since $\mathcal{Z}$ is $\partial_E M$ in this case. So while the Neumann extension doesn't give a diffusion on M, now the bridging extension does. Also, we see that the construction of the associated diffusion in Section 3.2 (which remains valid here) agrees with the one we just gave under the assumption of invariance under the isometry group. □

References

[1] Andrei Agrachev, Davide Barilari, and Ugo Boscain, *A comprehensive introduction to sub-riemannian geometry*, Cambridge Studies in Advanced Mathematics, vol. 181, Cambridge University Press, Cambridge, 2020. MR-3971262

[2] Andrei Agrachev, Ugo Boscain, and Mario Sigalotti, *A Gauss-Bonnet-like formula on two-dimensional almost-Riemannian manifolds*, Discrete Contin. Dyn. Syst. **20** (2008), no. 4, 801–822. MR-2379474

[3] Martin Barlow, Jim Pitman, and Marc Yor, *On Walsh's Brownian motions*, Séminaire de Probabilités, XXIII, Lecture Notes in Math., vol. 1372, Springer, Berlin, 1989, pp. 275–293. MR-1022917

[4] Ugo Boscain and Camille Laurent, *The Laplace-Beltrami operator in almost-Riemannian geometry*, Ann. Inst. Fourier (Grenoble) **63** (2013), no. 5, 1739–1770. MR-3186507

[5] Ugo Boscain and Dario Prandi, *Self-adjoint extensions and stochastic completeness of the Laplace-Beltrami operator on conic and anticonic surfaces*, J. Differential Equations **260** (2016), no. 4, 3234–3269. MR-3434398

[6] Zhen-Qing Chen and Masatoshi Fukushima, *Symmetric Markov processes, time change, and boundary theory*, London Mathematical Society Monographs Series, vol. 35, Princeton University Press, Princeton, NJ, 2012. MR-2849840

[7] Zhen-Qing Chen and Masatoshi Fukushima, *One-point reflection*, Stochastic Process. Appl. **125** (2015), no. 4, 1368–1393. MR-3310351

[8] Anja Göing-Jaeschke and Marc Yor, *A survey and some generalizations of Bessel processes*, Bernoulli **9** (2003), no. 2, 313–349. MR-1997032

[9] Tomoyuki Ichiba, Ioannis Karatzas, Vilmos Prokaj, and Minghan Yan, *Stochastic integral equations for Walsh semimartingales*, Ann. Inst. Henri Poincaré Probab. Stat. **54** (2018), no. 2, 726–756. MR-3795064

[10] Kiyoshi Itô and Henry P. McKean, Jr., *Diffusion processes and their sample paths*, Die Grundlehren der Mathematischen Wissenschaften, Band 125, Academic Press, Inc., Publishers, New York; Springer-Verlag, Berlin-New York, 1965. MR-0199891

[11] Frank B. Knight, *Essentials of Brownian motion and diffusion*, Mathematical Surveys, vol. 18, American Mathematical Society, Providence, R.I., 1981. MR-0613983

[12] Antoine Lejay, *On the constructions of the skew Brownian motion*, Probab. Surv. **3** (2006), 413–466. MR-2280299

[13] Jim Pitman and Marc Yor, *A guide to Brownian motion and related stochastic processes*, arXiv:1802.09679 (2018). MR-3827967

[14] Daniel Revuz and Marc Yor, *Continuous martingales and Brownian motion*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 293, Springer-Verlag, Berlin, 1991. MR-1083357

[15] Shinzo Watanabe, *Itô's theory of excursion point processes and its developments*, Stochastic Process. Appl. **120** (2010), no. 5, 653–677. MR-2603058

Acknowledgments. We thank Dario Prandi, Masha Gordina, Ioannis Karatzas, and Nate Eldredge for helpful conversations, and the anonymous referee for suggestions that improved the exposition.