HEAT KERNEL ESTIMATES ON HARNACK MANIFOLDS AND BEYOND

LAURENT SALOFF-COSTE

ABSTRACT

On a Riemannian manifold, M, the heat kernel is a smooth function on $(0, +\infty) \times M \times M$, $(t, x, y) \mapsto p(t, x, y)$, and the shape of this function depends on the properties of M. This article pays particular attention to the long-time, large-scale behavior of the heat kernel and its relation to the global geometry of M. When does the heat kernel look like a bell curve? If it does not, what does it look like and why? To answer such questions, one needs tools to obtain sharp two-sided estimates for the heat kernel in terms of the time variable $t > 0$ and basic geometric quantities depending on $x, y \in M$. Under what assumptions on M, can one hope to obtain such bounds?

MATHEMATICS SUBJECT CLASSIFICATION 2020

Primary 58J35; Secondary 60J60, 58J65, 60J46

KEYWORDS

Heat kernel, manifolds, Brownian motion

ICM INTERNATIONAL CONGRESS OF MATHEMATICIANS 2022 JULY 6—14

Proc. Int. Cong. Math. 2022, Vol. 6, pp. 4452–4473
DOI 10.4171/ICM2022/11

1. INTRODUCTION

Over the last 50 years, the heat kernel has become the subject of many studies in many different settings and for many purposes. Several fields of mathematics have obvious good reasons to pay particular attention to the heat kernel. In partial differential equations, it is the fundamental solution of the most basic parabolic equation which is the model for all evolution equations. In probability theory, it is the density of the distribution of Brownian motion at a given time. In mathematical physics, beyond its original role in the theory of heat, it leads to the notion of "abstract Wiener space," a building block in quantum field theory. But interest in the heat kernel goes well beyond these natural areas. It has been called ubiquitous and a universal gadget by mathematicians interested in topology (index theorems) or number theory (trace formulae). On a Riemannian manifold M, the heat kernel is a smooth function p on $(0, +\infty) \times M \times M$, $(t, x, y) \mapsto p(t, x, y)$. When M is the real line, $y \mapsto p(t, x, y)$ is a scaled version of the bell curve. See Figure 1.

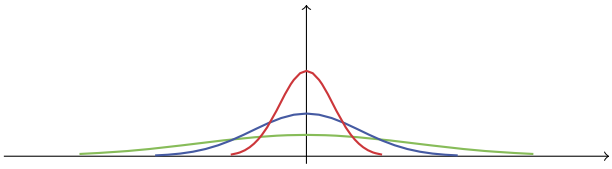

FIGURE 1

The bell curve as a model for the heat kernel: the heat kernel on the real line for $x = 0$ and three values of t: $t = 1/16, 1/4, 1$; $p(t, 0, y) = (4\pi t)^{-1/2} \exp(-|y|^2/4t)$.

This article pays particular attention to the long-time, large-scale behavior of the heat kernel and its meaning in global geometry. The goal is to develop tools to obtain sharp two-sided estimates for the heat kernel in terms of time, $t > 0$, and basic geometric quantities depending on $x, y \in M$. Because the heat kernel is the density function of the probability distribution of Brownian motion on M started at x at time t, if one can obtain sharp two-sided bounds on p valid for all $t > 0$, $x, y \in M$ and uniform over all manifolds M in a certain class $\mathcal{M}$, then one can say that, in that class $\mathcal{M}$, the geometry controls the behavior of Brownian motion in a precise sense. Such bounds have many further implications concerning spectral theory, potential theory, and global analysis.

2. EXISTENCE

When does the heat kernel exist? What is needed to define it uniquely? Even in the basic setting of Riemannian manifolds, these questions require some attention, and the answers involve the use of some significant machinery. It is useful to proceed by stages: first, define and prove the existence of an abstract weaker version of the heat kernel; then extract from this weaker version a proper heat kernel. For instance, one could prove existence in the sense of distribution theory and then prove that the constructed distribution is, in fact, a smooth function. Instead, we appeal here to semigroup theory so that our first step is to

define the heat semigroup $(P_t)_{t>0}$, from which we later intend to extract the heat kernel itself. What is needed for this purpose is a reasonable underlying space M equipped with a measure μ, the Hilbert space $L^2(M,\mu)$, and a Dirichlet form, $(\mathcal{E},\mathcal{D}(\mathcal{E}))$, that is, a densely defined closed nonnegative bilinear form on $L^2(M,\mu)$ with one additional property, the Markovian property. Namely, one requires that, for any $u\in\mathcal{D}(\mathcal{E})$, it holds that $|u|\in\mathcal{D}(\mathcal{E})$ and $\mathcal{E}(|u|,|u|)\le\mathcal{E}(u,u)$. See [9] (the Markovian property is akin to restricting ourselves to a positivity-preserving semigroup, something related to versions of the maximum principle). Functional analysis associates to the data $(M,\mu,(\mathcal{E},\mathcal{D}(\mathcal{E})))$ a self-adjoint semigroup of operators

$$P_t : L^2(M,\mu)\to L^2(M,\mu),\quad f\mapsto P_t f,$$

which solves the initial value problem

$$\begin{cases}\partial_t u=\Delta u,\\ u(0,\cdot)=f,\end{cases}$$

in the sense that $u(t,x)=P_t f(x)$ is the only solution of this problem when $f\in L^2(M,\mu)$. Here, Δ is the operator extracted from $\mathcal{E}$ in the same way that a symmetric matrix can be associated with any given positive-definite bilinear form on a finite Euclidean space. Namely, $\Delta u\in L^2(M,\mu)$ is such that $\phi\mapsto\int_M(\Delta u)\phi\,d\mu=\mathcal{E}(u,\phi)$ for all $\phi\in\mathcal{D}(\mathcal{E})$ whenever $u\in\mathcal{D}(\mathcal{E})$ has the property that $|\mathcal{E}(u,\phi)|\le C_u\|\phi\|_{L^2(M,\mu)}$. This densely-defined linear operator is called the infinitesimal generator of the semigroup $(P_t)_{t>0}$, and it can also be obtained using the formula

$$\Delta v=\lim_{t\to 0_+} t^{-1}(P_t v-v)$$

when this limit exists in $L^2(M,\mu)$. Moreover, $P_t=e^{t\Delta}$ where one can think of the right-hand side as defined by using the spectral theory applied to the self-adjoint operator Δ. In the context of a Riemannian manifold M equipped with its Riemannian measure μ, the classical choice is

$$\mathcal{E}(f,f)=\int_M|\nabla f|^2 d\mu$$

with the domain equal to the closure of smooth compactly supported functions for the norm $(\|f\|_2^2+\mathcal{E}(f,f))^{1/2}$. Using integration by parts, the infinitesimal generator of the associated semigroup is, indeed, the Laplacian, $\Delta f=\operatorname{div}(\nabla f)$ (in the case $M=\mathbb{R}^n$, the Euclidian space, $\Delta=\sum_1^n\partial_i^2$ where ∂_i is the partial derivative in the direction of the ith basis unit vector).

This simple construction provides us with a *transition kernel* $p(t,x,dy)$, which, for each t and x, is a nonnegative measure of finite total mass at most 1 in y. This is sometimes referred to as the heat kernel measure (at time t and centered at $x\in M$), and its definition is simply that $P_t f(x)=\int_M f(y)p(t,x,dy)$ for all $f\in L^2(M,d\mu)$. The question of the existence of the heat kernel (as a function) becomes the question of the absolute continuity of the measure $p(t,x,dy)$ with respect to the base measure μ. If absolute continuity holds then, abusing notation somewhat, $p(t,x,y)$ is defined by

$$p(t,x,dy)=p(t,x,y)d\mu(y).$$

Hence, to prove the existence of the heat kernel as a measurable locally bounded function, it suffices to prove bounds of the type $\sup_{x\in U}\{|P_t f(x)|\} \le C(t,U,V)\|f\|_1$, $f \in \mathcal{C}_c(V)$, for pairs (U,V) of open relatively compact sets that cover $M\times M$. Here $\mathcal{C}_c(V)$ is the space of continuous functions with compact support in V. On a smooth Riemannian manifold, the parabolic nature of the heat equation and local PDE theory provide such bounds, as well as the smoothness of the heat kernel. Properties of this type are known under the name of "ultracontractivity," and they can often be proved via the use of functional inequalities such as Sobolev or Nash inequalities. For instance, on a Riemannian manifold M, for any fixed $\nu > 0$, the Nash inequality ([24])

$$\|f\|_2^{2+4/\nu} \le C_1\|\nabla f\|_2^2\|f\|_1^{4/\nu}, \quad f \in \mathcal{C}_c^\infty(M), \tag{2.1}$$

is equivalent to the ultracontractivity inequality

$$\|P_t f\|_\infty \le C_2 t^{-\nu/2}\|f\|_1, \quad t > 0, f \in L^1(M,\mu),$$

which, in turn, is equivalent to

$$\sup_{x,y\in M}\{p(t,x,y)\} \le C_2 t^{-\nu/2}, \quad t > 0. \tag{2.2}$$

Even in the case of Riemannian manifolds where, thanks to local PDE theory, the existence and smoothness of the heat kernel are not in question, this Nash inequality technique gives access to a more quantitative control of the heat kernel. Because $\sup_{x,y}\{p_t(x,y)\} = \sup_x\{p(t,x,x)\}$, bounds of type (2.2), possibly with different functions of t on the right-hand side, are often called "on-diagonal heat kernel upper-bounds." They capture the decay of the heat kernel as time tends to infinity; see, e.g., [5]. In fact, with a little more work, the same set of ideas leads to the fact that (2.1) implies a Gaussian upper-bound involving the Riemannian distance between two points x, y, that is,

$$\forall\, t > 0,\ x, y \in M, \quad p(t,x,y) \le C_\varepsilon t^{-\nu/2}\exp\left(-\frac{d(x,y)^2}{4(1+\varepsilon)t}\right) \tag{2.3}$$

for any small $\varepsilon > 0$ (see, e.g., [27, CHAPTER 4, SECTION 2] and the references therein). For comparison, in our notation, the heat kernel in $\mathbb{R}^n$ is

$$\frac{1}{(4\pi t)^{n/2}}\exp\left(-\frac{|x-y|^2}{4t}\right).$$

The essentially universal nature of the Gaussian factor, $\exp(-d^2/4t)$, is somewhat surprising and very useful in practice: it makes the heat kernel behave almost like a compactly supported function, and this facilitates many manipulations.

3. THE GEOMETRY OF NICE DIRICHLET SPACES

The Nash inequality easily makes sense on a Dirichlet space $(M,\mu,(\mathcal{E},\mathcal{D}(\mathcal{E})))$ simply by interpreting $\|\nabla f\|_2^2$ as $\mathcal{E}(f,f)$ for $f \in \mathcal{D}(\mathcal{E}) \cap L^1(M,\mu)$. It then implies (2.2) and, in particular, the existence of a bounded heat kernel for all $t > 0$. Dirichlet forms can be local or nonlocal (the associated Markov process has continuous paths in the first case

and includes jumps in the second). Under some additional relatively mild assumptions, it is possible to extract from a (strictly) local Dirichlet form $(\mathcal{E}, \mathcal{D}(\mathcal{E}))$ a measure-valued bilinear form defined on $\mathcal{D}(\mathcal{E}) \times \mathcal{D}(\mathcal{E})$, $(f, g) \mapsto d\Gamma(f, g)$, such that $\mathcal{E}(f, g) = \int_M d\Gamma(f, g)$. For a given function $f \in \mathcal{D}(\mathcal{E})$, $d\Gamma(f, f)$ may or may not be absolutely continuous with respect to $d\mu$. When it is, we write $d\Gamma(f, f) = \Gamma(f, f)d\mu$. This is called the "carré du champ," and it is a substitute for the classical $|\nabla f|^2$. It extends naturally to a local version $\mathcal{D}_{\text{loc}}(\mathcal{E})$ of $\mathcal{D}(\mathcal{E})$. For arbitrary points $x, y \in M$, set

$$d(x, y) = \sup\{|f(x) - f(y)| : f \in \mathcal{C}(M) \cap \mathcal{D}_{\text{loc}}(\mathcal{E}), d\Gamma(f, f) \le d\mu\}.$$

This symmetric function of x and y can vanish or take the value $+\infty$. When it is finite, continuous, and defines the topology of M, it provides a good notion of distance called the "intrinsic distance" of the given Dirichlet form; see, e.g., [22, 28]. In the Riemannian setting, the intrinsic distance is simply the Riemannian distance. To see a different (non-Riemannian), yet classical set of examples, let $M = G$ be a unimodular Lie group equipped with its Haar measure μ, and a family $\{X_1, \dots, X_k\}$ of left-invariant vector fields which generates the Lie algebra $\mathfrak{g}$ of G (i.e., these fields together with all their iterated Lie brackets span $\mathfrak{g}$, linearly). Set

$$\mathcal{E}(f, f) = \int_G \sum_1^k |X_i f|^2 d\mu$$

for f in the closure of $\mathcal{C}_c^\infty(G)$ for the norm $(\int_G |f|^2 d\mu + \int_G \sum_1^k |X_i f|^2 d\mu)^{1/2}$. In this case, the distance d is the sub-Riemannian distance associated with the family $\{X_1, \dots, X_k\}$ and this example serves as a model for the development of sub-Riemannian geometry and the analysis of the related subelliptic Laplacians (here, $\Delta = \sum_1^k X_i^2$, because G is unimodular).

In general, when the intrinsic distance d is continuous and defines the topology of M, it is possible again to obtain (2.3) from (2.1) (e.g., [28]). Recently, based in part on the notions and techniques described here, Carron and Tewodrose proved the following rigidity result [4]. Consider a σ-compact complete metric space (M, d), equipped with a positive Radon measure μ and with a Dirichlet form $(\mathcal{E}, \mathcal{D}(\mathcal{E}))$. Assume that the associated heat semigroup admits a heat kernel p which satisfies, for all $(t, x, y) \in (0, +\infty) \times M \times M$,

$$p(t, x, y) = \frac{1}{(4\pi t)^{\alpha/2}} \exp\left(-\frac{d(x, y)^2}{4t}\right).$$

Then α is an integer, M is $\mathbb{R}^\alpha$, d is the Euclidean metric on $\mathbb{R}^\alpha$, and μ is the α-dimensional Hausdorff measure. The Dirichlet form $(\mathcal{E}, \mathcal{D}(\mathcal{E}))$ is the usual Euclidean Dirichlet form.

On the other hand, there are many interesting examples that are definitively not Euclidean and whose heat kernel satisfies

$$\frac{c_1}{t^{\alpha/2}} \exp\left(-C_1 \frac{d(x, y)^2}{t}\right) \le p(t, x, y) \le \frac{C_2}{t^{\alpha/2}} \exp\left(-c_2 \frac{d(x, y)^2}{t}\right).$$

These examples include uniformly elliptic operators in $\mathbb{R}^n$, for which

$$\mathcal{E}(f, g) = \int_{\mathbb{R}^n} \sum_{i,j=1}^n a_{ij}(x) \nabla f(x) \cdot \nabla g(x) dx,$$

where the coefficients a_{ij} are bounded, measurable, and satisfy $\sum_{ij} a_{ij}(x)\xi_i\xi_j \geq \varepsilon\|\xi\|_2^2$ for all $x, \xi \in \mathbb{R}^n$ and some $\varepsilon > 0$. In this case $\alpha = n$ (see [1, 2, 25]). Another example is the Heisenberg group H of 3×3 matrices

$$\begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix}, \quad x, y, z \in \mathbb{R},$$

equipped with the Dirichlet form $\mathcal{E}(f, f) = \int_H (|Xf|^2 + |Yf|^2)d\mu$ where X (resp. Y) is the left-invariant vector field equal to $\partial/\partial x$ (resp. $\partial/\partial y$) at the identity. In this case $\alpha = 4$. In all these examples the correct interpretation of the on-diagonal factor, $t^{-\alpha/2}$, is that it is $1/V(x, \sqrt{t})$ where $V(x, r) = \mu(\{z \in M : d(x, z) < r\})$, the volume of the ball of radius r and center x. This leads us to consider the following hypothetical two-sided Gaussian bound:

$$\frac{c_1}{V(x, \sqrt{t})} \exp\left(-C_1 \frac{d(x, y)^2}{t}\right) \leq p(t, x, y) \leq \frac{C_2}{V(x, \sqrt{t})} \exp\left(-c_2 \frac{d(x, y)^2}{t}\right). \tag{3.1}$$

When (3.1) holds on a Riemannian manifold, one can answer many questions. For instance, Brownian motion on such a manifold is transient if and only if $\int^{+\infty} \frac{ds}{V(x,\sqrt{s})} < +\infty$. If this integral is finite then the Green function $G(x, y)$, i.e., the function such that

$$\Delta^{-1} f(x) = \int_M G(x, y) f(y) d\mu(y), \quad f \in \mathcal{C}_c^\infty(M),$$

satisfies

$$c \int_{d(x,y)^2}^{+\infty} \frac{ds}{V(x, \sqrt{s})} \leq G(x, y) \leq C \int_{d(x,y)^2}^{+\infty} \frac{ds}{V(x, \sqrt{s})}.$$

By integrating (3.1) with respect to y over the ball $B(x, 2\sqrt{t})$ and noting that the integral of the heat kernel is at most 1, one easily sees that (3.1) implies that the manifold M must be doubling in the following sense.

Definition 3.1. A metric measure space is called doubling if there exists a constant D such that, for all $x \in M$ and $r > 0$, $V(x, 2r) = \mu(B(x, 2r)) \leq DV(x, r) = D\mu(B(x, r))$.

In the next section, we answer the following question: Which Riemannian manifolds satisfy (3.1)?

4. HARNACK MANIFOLDS AND DIRICHLET SPACES

On a smooth manifold with boundary, a function u is harmonic in an open ball B if it is smooth in B, satisfies $\Delta u = 0$, and has vanishing normal derivative on $\delta_M \cap B$. Similarly, a function u is a solution of the heat equation in a time-space cylinder $Q = (a, b) \times B$ if it is smooth there, satisfies $(\partial_t - \Delta)u = 0$ in Q, and has vanishing normal derivative on $\delta_M \cap B$. When dealing with more general contexts, including local (regular) Dirichlet spaces, the appropriate notion of a *weak solution* must be used instead; see, e.g., [20, 29] for details.

The elliptic Harnack inequality is one of the most well-known inequalities in analysis and goes back to the nineteenth century. In $\mathbb{R}^n$, it states that there is a constant C_n

such that any nonnegative function u, harmonic in a ball $B = B(x,r)$, satisfies $\sup_{\frac{1}{2}B}\{u\} \leq C_n \inf_{\frac{1}{2}B}\{u\}$ where $\frac{1}{2}B = B(x,r/2)$. A Riemannian manifold (or Dirichlet space as above which admits a good intrinsic distance) satisfies the elliptic Harnack inequality when there is a constant C_M such that any nonnegative function u which is harmonic in a ball $B = B(x,r)$ satisfies

$$\sup_{\frac{1}{2}B}\{u\} \leq C_M \inf_{\frac{1}{2}B}\{u\}.$$

Until recently, there was no clear characterization of this property in geometric terms, but this problem is resolved beautifully in [3], to which the reader is referred.

The importance and usefulness of the parabolic Harnack inequality only became apparent in the second half of the twentieth century in the work of Nash [24], Moser [23], and many others after them. Consider a time-space cylinder $Q = (s-r^2, s) \times B(x,r)$ and the smaller separated subcylinders $Q_- = (s-3r^2/4, s-r^2/2) \times B(x,r/2)$ and $Q_+ = (s-r^2/4, s] \times B(x,r/2)$. We say that M satisfies the parabolic Harnack inequality (at all scales and locations) if there is a constant C_M such that any nonnegative solution u of the heat equation in Q satisfies

$$\sup_{Q_-}\{u\} \leq C_M \inf_{Q_+}\{u\}. \tag{4.1}$$

In what follows we will make constant use of the following definition.

Definition 4.1. We say that a Riemannian manifold M is Harnack with constant C if it satisfies the parabolic Harnack inequality (4.1), at all scales and locations, with a constant $C_M \leq C$.

Given a precompact open subset $\Omega \subseteq M$, set

$$\lambda(\Omega) = \inf\left\{\frac{\int_\Omega |\nabla f|^2 d\mu}{\int_\Omega |f - f_\Omega|^2 d\mu} : f \in W^1(\Omega), f - f_\Omega \neq 0\right\}.$$

Here, f_Ω is the average value of f on Ω and $W^1(\Omega)$ is the set of all L^2-functions in Ω whose gradient in Ω (in the sense of distributions) can be represented as an L^2-vector field in Ω. The language of Dirichlet spaces allows us to view $\lambda(\Omega)$ as the lowest *positive* eigenvalue of the Neumann-Laplacian in Ω (even if Ω does not have a smooth boundary).

Definition 4.2. We say that a Riemannian manifold M satisfies the Poincaré inequality at all scales and locations, with a constant at most P, if

$$\forall x \in M,\ r > 0, \quad \lambda\big(B(x,r)\big) \geq 1/\big(Pr^2\big). \tag{4.2}$$

Definition 4.3. Fix $\kappa \geq 1$. We say that a Riemannian manifold M satisfies the weak Poincaré inequality with parameter κ at all scales and locations, with a constant at most P, if

$$\forall x \in M, r > 0, \forall f \in \mathcal{C}_b^\infty\big(B(x,\kappa r)\big), \quad \int_{B(x,r)} |f - f_{B(x,r)}|^2 d\mu \leq P r^2 \int_{B(x,\kappa r)} |\nabla f|^2 d\mu.$$

With these definitions we can answer the question posed at the end of the previous section: Which Riemannian manifolds satisfy the two-sided Gaussian heat kernel estimate (3.1)?

Theorem 4.4. *A complete Riemannian manifold M is Harnack if and only if it satisfies* (3.1). *These properties are also equivalent to the fact that M is doubling and satisfies the Poincaré inequality* (4.2) *at all scales and locations. Finally, doubling and the weak Poincaré inequality with a fixed parameter $\kappa \geq 1$ are enough to imply that M is Harnack.*

This theorem is essentially taken from the independent works [10] and [26], which both used ideas developed in the earlier works by various other authors. The proof in [26,27] uses the well-know techniques of Nash and Moser, as well as ideas developed by D. Jerison and S. Kusuoka and D. Stroock. K. T. Sturm extended this theorem in an important and useful way to the context of local Dirichlet spaces admitting a good intrinsic distance [28,29]. In this abstract context, the well-known fact that (4.1) implies the Hölder continuity of the (weak) solutions of the heat equation provides an important method to prove the continuity of the heat kernel. Each property used in Theorem 4.4, the parabolic Harnack inequality (4.1), the two-sided-Gaussian bound (3.1), and the conjunction of doubling and the Poincaré inequality, comes with a small set of fundamental constants, the constant C_M in (4.1), the constants c_1, C_1, c_2, C_2 in (3.1), and the doubling constant D and Poincaré constants κ, P. In each case, the constants of a given property can be controlled solely in terms of the constants of one of the other equivalent properties. For instance, fix large, positive, reals D and P. There is a constant $C = C(D, P)$ such that any complete manifold satisfying doubling with constant at most D and the Poincaré inequality (4.2) with constant at most P also satisfies the parabolic Harnack inequality (4.1) with constant at most C.

This brings us to the following conjecture due to Maria Gordina, Nate Eldredge, and the author.

Conjecture 4.5 ([8]). *Given a compact Lie group G, there exists a constant $H(G)$ such that all left-invariant Riemannian metrics on G are Harnack with constant at most $H(G)$.*

Because of (3.1), what this would mean is that, on a given compact Lie group, all left-invariant diffusion processes are uniformly controlled by their own geometry.

A simple argument going back to the work of N. Varopoulos shows that a left-invariant Riemannian metric on a unimodular Lie group G which is doubling also satisfies the Poincaré inequality. It follows that the conjecture above can be stated in the following much simpler form.

Conjecture 4.6 ([8]). *Given a compact Lie group G, there exists a constant $D(G)$ such that all left-invariant Riemannian metrics on G are doubling with constant at most $D(G)$.*

The best evidence for these conjectures is that they hold for abelian Lie groups (compact or not), for nilpotent Lie groups (not compact, if not abelian), and for SU(2) (see [8]). The case of $U(2)$ and $\mathrm{SU}(2) \times A$ where A is an abelian Lie group is in preparation by the same authors (it is surprisingly more involved than the case of SU(2)).

These conjectures can be compared with the following theorem which follows from [21] (this theorem covers the case of abelian groups because their left-invariant Riemannian metrics have 0 curvature).

Theorem 4.7. *Fix a dimension n. There is a constant C_n such that all complete Riemannian manifolds with nonnegative Ricci curvature are Harnack with constant at most C_n.*

Repeating something said above, one of the consequence of this theorem is that, on manifolds with nonnegative Ricci curvature and dimension at most n, the behavior of Brownian motion is controlled by the geometry of the manifold, uniformly over all such manifolds.

For a typical compact Lie group G in Conjectures 4.5–4.6, there is no finite common lower bound on the Ricci curvature of all left-invariant Riemannian metrics. So, Theorem 4.7 does not help much in settling these conjectures. In fact, although we stated these conjectures for left-invariant Riemannian metrics, they automatically extend to (e.g., include) left-invariant sub-Riemannian geometries because the desired property is uniform over all Riemannian metrics on G. They further extend to left-invariant structures on analytic subgroups of the compact group G, showing that G cannot contain an analytic subgroup of exponential volume growth (the fact that a compact Lie group G cannot contain an analytic subgroup of exponential volume growth is indeed known; it can be viewed as supporting evidence for the conjectures).

Remark 4.8. Most results concerning Harnack inequalities (elliptic or parabolic) in the context of Riemannian geometry are based on Ricci curvature lower bounds in the spirit of the famous works of S. T. Yau, Cheng and Yau, and Li and Yau. This technique leads to "gradient Harnack inequalities" which imply inequalities of the type (4.1). One can strengthen Conjecture 4.5 by asking for a uniform *parabolic Harnack gradient inequality* over all left-invariant Riemannian metrics on G. This stronger conjecture is open even for SU(2).

5. NON-HARNACK MANIFOLDS

In the remaining sections, we will focus on examples of non-Harnack manifolds. Finding ways to obtain sharp heat kernel estimates for manifolds to which Theorem 4.4 DOES NOT apply is a major challenge. Techniques exist that provide good on-diagonal estimates in various particular situations, and it is known that the Gaussian factor of the type $\exp(-cd(x,y)^2/t)$ is somewhat universal (although not always sharp at large scales and large time). One can phrase this challenge more precisely by asking for upper/lower bounds for the heat kernel $p(t,x,y)$ in terms of some explicit functions $(t,x,y) \mapsto g_{\mathbf{c}}(t,x,y)$ which are expressed in terms of t and basic geometric quantities, including the volume functions $V(x,r)$, $V(y,r)$, $r>0$, the distance function $d(x,y)$, and perhaps other similar quantities. Here, $\mathbf{c}$ represents a positive constant (more generally, finite set of positive constants) that enters the definition of the function g and may be different in upper and lower bounds. In case when one knows or expects that $\int_M p(t,x,y)dy = 1$ (i.e., heat diffusion on M is conservative), it is highly desirable that the functions $g_{\mathbf{c}}$ used to estimate p satisfy $\varepsilon_{\mathbf{c}} \leq \int_M g_{\mathbf{c}}(t,x,y)dy \leq \varepsilon_{\mathbf{c}}^{-1}$, for some $\varepsilon_{\mathbf{c}} > 0$.

This challenge has many facets. Here, we will focus only on one of them. Namely, we are going to focus on manifolds which lack basic homogeneity but can be decomposed

into simpler pieces that are Harnack. The simplest basic example of such is the catenoid which, roughly speaking, is the connected sum of two planes. The catenoid is doubling but does not satisfy the Poincaré inequality at all scales and locations; see the next section.

First, we briefly describe explicitly a different type of challenging example. The Lie group Sol is the model for one of the eight "geometries" that are the building blocks of manifolds in dimension 3 (Perelman's theorem, formerly Thurston conjecture; the Heisenberg group mentioned earlier is another one of these eight). We can describe Sol as the matrix group

$$\begin{pmatrix} e^x & 0 & y \\ 0 & e^{-x} & z \\ 0 & 0 & 1 \end{pmatrix}, \quad x, y, z \in \mathbb{R}.$$

We equip Sol with the left-invariant metric associated with the orthonormal basis of $\mathbb{R}^3$ viewed as the tangent space of Sol at the identity, id. The rough on-diagonal behavior of the heat kernel $p(t, \mathrm{id}, \mathrm{id})$ is described by the function $\exp(-t^{1/3})$ for large t, but no off-diagonal estimate of the type described above is known. The simplest way to see that Theorem 4.4 does not apply (positively) to this example is to note that the volume of large balls grows exponentially fast with the radius. Hence the doubling condition fails. In a similar spirit, sharp off-diagonal estimates of the heat kernel of the universal cover of a compact manifold whose fundamental group is a finitely generated solvable group of exponential volume growth is a challenge that goes well beyond existing techniques.

6. MANIFOLDS MADE OF NICE PIECES AND RECONSTRUCTION: BASIC EXAMPLES

Consider the connected sum of two Euclidean spaces, $M = \mathbb{R}^n \# \mathbb{R}^n$, $n \geq 2$, equipped with a Riemannian metric, that is, the Euclidean metric away from the central collar gluing the two copies of $\mathbb{R}^n$ together. This manifold is made of two very nice Harnack pieces, the two Euclidean spaces. It is doubling at all scales and locations, but the Poincaré inequality fails to hold for balls of large radius centered at the collar. In fact, the second-lowest Neumann eigenvalue for such a large central ball is of order $1/(r^2 \log r)$ when $n = 2$ and $1/r^n$ when $n > 2$ (for the Poincaré inequality to hold, we need $1/r^2$). Write M as the disjoint union of a compact part K (the collar), and the two disconnected ends E_1, E_2, both equal to $\mathbb{R}^n \setminus \mathbf{B}$, where $\mathbf{B}$ is a ball centered at the origin in $\mathbb{R}^n$, of radius large enough so that the metric of M on each E_i is Euclidean. For $x, y \in M$, consider the following geometric quantities:

- $|x| = \sup_{z \in K}\{d(x, z)\}$;
- $d_+(x, y)$, the infimum of the lengths of smooth curves joining x to y in M having a *nonempty intersection* with K;
- $d_\emptyset(x, y)$, the infimum of the lengths of smooth curves joining x to y in M having a *empty intersection* with K.

For $n > 2$, set

$$g_{M,c}(t,x,y) = \frac{1}{ct^{n/2}}\left(\frac{1}{|x|^{n-2}} + \frac{1}{|y|^{n-2}}\right)\exp\left(-c\frac{d_+(x,y)^2}{t}\right) + \frac{1}{ct^{n/2}}\exp\left(-c\frac{d_\emptyset(x,y)^2}{t}\right).$$

It is proved in [18] that there are constants c_1, c_2 such that, for all $(t,x,y) \in (0,+\infty)\times M\times M$, the heat kernel of M, $p_M(t,x,y)$, satisfies

$$g_{M,c_1}(t,x,y) \le p_M(t,x,y) \le g_{M,c_2}(t,x,y). \tag{6.1}$$

More complicated formulae of the same type apply when $M = M_1\#\cdots\#M_k$ with each $M_i = \mathbb{R}^{n_i}\times\mathbb{S}^{N-n_i}$ for $n_i \ge 3$. In particular, in this case, for any fixed point $o \in M$, there are constants c_1, C_1 such that, for all $t > 1$, $c_1 t^{-\min\{n_i\}/2} \le p_M(t,o,o) \le C_1 t^{-\min\{n_i\}/2}$.

The case of the connected sum of two planes, $\mathbb{R}^2\#\mathbb{R}^2$, is different because Brownian motion on $\mathbb{R}^2$ is recurrent (an open ball is visited with probability 1 from any starting point; equivalently, there is no positive Green's function). It is proved in [13, 18] (some technical elementary manipulations are required to turn the results of [13, 18] into the statement given here) that the heat kernel for $M = \mathbb{R}^2\#\mathbb{R}^2$ satisfies (6.1) with

$$g_{\mathbb{R}^2\#\mathbb{R}^2,c}(t,x,y) = \frac{1}{ct}\left(\frac{\log(e(1+t/|x|^2))}{\log(1+t+|x|^2)} + \frac{\log(e(1+t/|y|^2))}{\log(1+t+|y|^2)}\right)\exp\left(-c\frac{d_+(x,y)^2}{t}\right) + \frac{1}{ct}\exp\left(-c\frac{d_\emptyset(x,y)^2}{t}\right).$$

In this formula, the second term, $\frac{1}{ct}\exp(-cd_\emptyset(x,y)^2/t)$, is 0 if x, y are in different planes, and it always dominates if x, y are in the same plane.

7. MANIFOLDS WITH FINITELY MANY NONPARABOLIC HARNACK ENDS

Our goal now is to generalize as much as possible the results described above in model cases. Consider the connected sum $M = M_1\#\cdots\#M_k$ of k complete noncompact weighted Riemannian manifolds with boundary. So M may have a nonempty "boundary" $\delta M \subset M$ along which it is modeled locally by the half-space $\mathbb{R}^n_+$, and M equipped with its Riemannian distance is a complete metric space. The weight σ is a positive smooth function (in fact, continuity is more than enough). The heat equation on this weighted manifold, and the heat kernel p_M, are associated with the Dirichlet space

$$\left(M, \mu, \int_M |\nabla f|^2 d\mu, W_0^1(M, d\mu)\right)$$

where dx is the Riemannian measure, $\mu(dx) = \sigma(x)dx$, and $W_0^1(M)$ is the closure of smooth compactly supported functions on M under the norm $(\int_M (|f|^2 + |\nabla f|^2)d\mu)^{1/2}$. By definition, we can write M as the disjoint union $M = K \cup E_1 \cup \cdots \cup E_k$ where K is compact and $\overline{E_i}$ are smooth manifolds with boundary isometric to $M_i \setminus K_i$ for some compact K_i in M_i. Each E_i inherits a weight $\sigma_i = \sigma|_{E_i}$. In more classical terms, the Laplacian of a

smooth function f on (M,μ) is $\frac{1}{\sigma}\mathrm{div}(\sigma\nabla f)(x)$ at point $x \in M \setminus \delta M$, and the heat equation is taken with Neumann boundary condition along δM. Although we informally refer to the M_i or the E_i, $1 \le i \le k$, as the "ends" of M, it is not necessarily the case in the setting described above that they represent the full list of the topological ends of M as any one of them could possibly split if a very large ball is removed.

In this section, we make two fundamental assumptions:

(HE) Each weighted manifold M_i is Harnack at all scales and locations.

(NPE) Each weighted manifold M_i, $1 \le i \le k$, is nonparabolic.

Regarding (NPE), note that the dichotomy parabolic/nonparabolic (nonparabolic means the "existence of a positive Green function") is identical to the dichotomy recurrent/transient (recurrence means an open ball is visited with probability 1). Moreover, a manifold satisfying (HE) is nonparabolic if and only if $\int_1^\infty \frac{ds}{V(x,\sqrt{s})} < +\infty$. See [11] for a comprehensive review. Under the two assumptions (HE)–(NPE), each E_i is indeed a representative of an end of M in the classical sense because nonparabolic Harnack manifolds can only have one end [15].

For any $x, y \in M$, define $|x|, d_+(x,y), d_\emptyset(x,y)$ in terms of the compact set K as before. Also, set

$$i_x = \begin{cases} i & \text{if } x \in E_i, 1 \le i \le k, \\ 0 & \text{if } x \in K. \end{cases}$$

For $x \in E_i$, set $V_i(x,r) = \mu(B(x,r) \cap E_i)$ and $V_i(r) = \mu(B(o,r) \cap (K \cup E_i))$ where o is a fixed central point in K. Set

$$V_0(r) = \min\{V_i(r) : 1 \le i \le k\}$$

and

$$H_{(M,\mu)}(t,x) = H(t,x) = \min\left\{1, \frac{|x|^2}{V_{i_x}(|x|)} + \left(\int_{|x|^2}^t \frac{ds}{V_{i_x}(\sqrt{s})}\right)_+\right\}. \tag{7.1}$$

To understand the behavior of $H(x,t)$, note that (HE) and (NPE) imply $\int^{+\infty} \frac{ds}{V_{i_x}(\sqrt{s})} < +\infty$. Whenever $V_{i_x}(r)/V_{i_x}(s) \ge c(r/s)^\alpha$ with $\alpha > 2$, the integral $\int_{|x|^2}^t \frac{ds}{V_{i_x}(\sqrt{s})}$ is dominated by the term $\frac{|x|^2}{V_{i_x}(|x|)}$. This integral becomes relevant when the end containing x, E_{i_x}, is only barely nonparabolic, for instance, if $V_{i_x}(r)$ grows like $r^2(\log r)^2$.

Theorem 7.1 ([18]). *Assuming that* (HE) *and* (NPE) *are satisfied, the heat kernel* $p_{(M,\mu)}(t,x,y)$ *satisfies the two-sided estimate* (6.1) *with*

$$g_{M,c}(t,x,y) = \left(\frac{H(x,t)H(y,t)}{cV_0(\sqrt{t})} + \frac{H(y,t)}{cV_{i_x}(\sqrt{t})} + \frac{H(x,t)}{cV_{i_y}(\sqrt{t})}\right)\exp\left(-c\frac{d_+(x,y)^2}{t}\right) + \frac{1}{cV_{i_x}(x,\sqrt{t})}\exp\left(-c\frac{d_\emptyset(x,y)^2}{t}\right). \tag{7.2}$$

Note that the last term comes into play only when x and y are in the same end. For $t \in (0,1)$, it is possible to show that

$$\frac{1}{c_1 V(x,\sqrt{t})} \exp\left(-c_1 \frac{d(x,y)^2}{t}\right) \le g_{M,c}(t,x,y) \le \frac{1}{c_2 V(x,\sqrt{t})} \exp\left(-c_2 \frac{d(x,y)^2}{t}\right),$$

as expected, though this takes a bit of technical work.

For fixed x_0, y_0 and large t, $p_M(t,x_0,y_0)$ behaves as $\frac{1}{V_0(\sqrt{t})}$, that is, it is controlled by the volume growth of the smallest end at scale $\sqrt{t}$ (it is possible that the "the smallest end" changes depending on the scale at which the question is asked).

8. NONPARABOLIC MANIFOLDS WITH FINITELY MANY HARNACK ENDS

In this section, the manifold $M = M_1 \# \cdots \# M_k$ is a complete noncompact weighed Riemannian manifold with boundary as before, and we continue to make assumption (HE) that each end is a Harnack manifold. The manifold M is nonparabolic if and only if at least one of the ends M_i, $1 \le i \le k$, is nonparabolic. So, we may weaken assumption (NPE) to

(NP) At least one of the weighted manifolds M_i, $1 \le i \le k$, is nonparabolic.

However, under these circumstances, we need to make a further assumption in order to obtain sharp heat kernel estimates. Namely, we assume the following:

(RCA*) All the ends M_i, $1 \le i \le k$, that are parabolic must satisfy the relatively connected annulus condition (RCA): There exists a constant $A > 1$ such that, for any $R > A^2$ and any two points $x, y \in E$ with $|x| = |y| = R$, there is a continuous curve connecting x to y in $\{z : R/A \le |z| \le AR\}$.

In words, we assume that, in any parabolic end E_{i_0} of M, two points at a distance about R from the central part K can be connected without going too far toward infinity (no further than AR) and without coming back too close to the central part K (no closer than R/A). This condition is key to obtaining the results below. Note that again, under these assumptions, $E_1, \ldots, E_k$ are indeed representative of the ends of M in the classical sense as any one of them is a manifold with only one end in the classical sense: the nonparabolic ones because they are Harnack, and the parabolic ones because of condition (RCA).

Each E_i is an incomplete manifold with boundary such that $\delta E_i = \delta M \cap E_i$ and $\partial E_i = \overline{E_i} \setminus E_i \subset K$. A *harmonic profile* for E_i is a function u_i which is positive in E_i, vanishes along ∂E_i, and is harmonic in E_i (this includes the condition that u_i has vanishing normal derivative along δE_i). It is known that such a function exists, is continuous on $\overline{E_i}$, and is unique up to multiplication by a positive real (recall hypothesis (HE)). Moreover, there is a constant $c > 0$ such that, for all $x \in E_i$ with $|x|$ large enough (e.g., $|x| \ge 2(1 + \operatorname{diam}(K))$),

$$c \int_1^{|x|^2} \frac{ds}{V_i(\sqrt{s})} \le u_i(x) \le c^{-1} \int_1^{|x|^2} \frac{ds}{V_i(\sqrt{s})}.$$

This fact (see [15,18,20]) depends crucially upon the hypotheses (HE) and (RCA*). It implies that the harmonic profile of a nonparabolic end E_i is bounded and bounded away from 0 in E_i away from K while the harmonic profile of a parabolic end tends to infinity at infinity.

A *harmonic profile* for $M = K \cup E_1 \cup \cdots \cup E_k$ is a positive harmonic function h on M (this implies it has vanishing normal derivative along δM) which, in each E_i, behaves as u_i at infinity. Again, it is known that such a function exists under assumption (NP), see [18,30].

We use this positive harmonic function h on M to consider the new weighted manifold, (M, μ_{h^2}), where

$$\mu_{h^2}(dx) = h^2(x)\mu(dx) = h^2(x)\sigma(x)dx,$$

whose Dirichlet form is

$$\int_M |\nabla f|^2 d\mu_{h^2} = \int_M |\nabla f|^2 h^2 d\mu.$$

Because h is harmonic,

$$\int_M |\nabla f|^2 h^2 d\mu = \int_M |\nabla(hf)|^2 d\mu, \quad f \in \mathcal{C}_c^\infty(M).$$

This means that the heat equation associated with

$$\left(M, \mu_{h^2}, \int_M |\nabla f|^2 d\mu_{h^2}, W_0^1(M, \mu_{h^2})\right)$$

is $\partial_t u - \frac{1}{h}\Delta_\sigma(hu) = 0$, and the associated heat kernel $p_{(M,\mu_{h^2})}$ is given by

$$p_{(M,\mu_{h^2})} = \frac{1}{h(x)h(y)} p_{(M,\mu)}(t, x, y). \tag{8.1}$$

In probability theory, the use of this relation is often referred to as the "Doob transform" technique after Joseph Doob. For us, its significance is that, assuming we know the profile h, it is possible to turn estimates of $p_{(M,\mu_{h^2})}$ into estimates of $p_{(M,\mu)}$. This is useful because of the following theorem.

Theorem 8.1. *Assume that* (M, μ) *satisfies* (HE), (NP), *and* (RCA*). *Then* (M, μ_{h^2}) *satisfies* (HE) *and* (NPE).

What this theorem says is that the *weighted* Riemannian manifold $(M, h^2\mu)$,

$$M = M_1 \# \cdots \# M_k = K \cup E_1 \cup \cdots \cup E_k,$$

is a connected sum of Harnack weighted manifolds, and, moreover, each of them is nonparabolic. The proof that each $(\overline{E_i}, \mu_{h^2}|_{E_i})$ is a Harnack manifold proceeds by showing that doubling and Poincaré inequalities hold at all scales and locations; see [17,20].

Theorems 7.1–8.1 and (8.1) lead to a sharp two-sided estimate for the heat kernel of the nonparabolic weighted manifold (M, μ), $\mu(dx) = \sigma(x)dx$ in terms of the functions

$$g_{(M,\mu),c}(t, x, y) = h(x)h(y)g_{(M,\mu_{h^2}),c}(t, x, y), \tag{8.2}$$

where $g_{(M,\mu_{h^2}),c}$ is given by (7.2).

Theorem 8.2. *Assume that* (M,μ) *satisfies* (HE), (NP), *and* (RCA*). *Then, there exist* c_1, c_2 *such that, for all* $t>0$, $x,y\in M$, *we have*

$$g_{(M,\mu),c_1}(t,x,y)\le p_{(M,\mu)}(t,x,y)\le g_{(M,\mu),c_2}(t,x,y) \tag{8.3}$$

where $g_{(M,\mu),c}$ *is given by* (8.2).

It takes quite a bit of work to unpack this statement. How explicit the obtained estimates are depends very much on our ability to understand the function h, the harmonic profile of M. One key point to notice is that everything depends on the volume growth function of each end and our ability to compute, for $r>1$, quantities such as

$$r\mapsto 1+\int_1^{r^2}\frac{ds}{V_i(\sqrt{s})},$$

which controls u_i (hence h in E_i), and

$$(r,t)\mapsto \int_{r^2}^{t}\frac{ds}{(1+\int_1^s\frac{d\tau}{V_i(\sqrt{\tau})})^2V_i(\sqrt{s})},\quad r^2<t,$$

which are needed to control the function $H_{(M,\mu_{h^2})}$ in E_i. These computations are the trickiest for ends that are near the threshold separating parabolic from nonparabolic ends, e.g., when $V_i(r)$ grows as r^2 up to a slowly-varying factor (think of $r^2[\log(1+r)]^\alpha$ with $\alpha\in\mathbb{R}$). It is worth noting that the result holds without restriction on the behavior of the volume growth through the parabolic/nonparabolic threshold, as long as (M,μ) itself is nonparabolic. The simplest general result concerns the long-time behavior of $p_{(M,\mu)}(t,x_0,y_0)$ for fixed $x_0,y_0\in M$ which is that

$$\min_{1\le i\le k}\left\{\frac{c}{(1+\int_1^t\frac{ds}{V_i(\sqrt{s})})^2V_i(\sqrt{t})}\right\}\le p_{(M,\mu)}(t,x_0,y_0)\le \min_{1\le i\le k}\left\{\frac{C}{(1+\int_1^t\frac{ds}{V_i(\sqrt{s})})^2V_i(\sqrt{t})}\right\}.$$

Example 8.3. To illustrate what this says, consider the case when $M=K\cup E_1\cup E_2\cup E_3$ is a solid 3-dimensional body with 3 ends that can be described as follows:

- E_1 is a half-cylinder of radius 1 around the the bottom part of the z-axis,
$$E_1=\{(x,y,z):x^2+y^2\le 1, z<-1\};$$
- E_2 is essentially a solid planar slab around the xy-plane,
$$E_2=\{(x,y,z):x^2+y^2>2, -1\le 2z\le 1\};$$
- E_3 is essentially a solid half-cone of revolution of positive aperture around the positive z-axis, say
$$E_3=\{(x,y,z):x^2+y^2\le z, z>1\};$$
- K is a compact set joining these ends smoothly together, and the measure on M is Lebesgue measure (i.e., $\sigma\equiv 1$).

In this description, $\overline{E_i}$ are smooth manifolds with corners but this can easily be fixed. The attentive reader will note that the enumeration of the ends corresponds precisely to their

volume growth, with $V_1(r)$ growing linearly, $V_2(r)$ growing quadratically, and $V_3(r)$ growing as r^3. These ends are all Harnack, satisfy (RCA), and E_3 is nonparabolic so that M is nonparabolic. The other two ends are parabolic. The harmonic profile h of M satisfies $c|x| \le h(x) \le c^{-1}|x|$ in E_1, $c\log(1+|x|) \le h(x) \le c^{-1}\log(1+|x|)$ in E_2, and $c \le h(x) \le c^{-1}$ in E_3. It follows that $V_{i,h^2}(r)$ grows as r^3 in both E_1 and E_3, while $V_2(r)$ grows as $r^2\log^2 r$. For $o = (0,0,0) \in M$ and $t > 1$, these computations give

$$\frac{c}{t\log^2(1+t)} \le p_{(M,\mu)}(t,o,o) \le \frac{C}{t\log^2(1+t)}.$$

Now, consider the following two questions (see Figure 2):

(a) At time $t > 1$, where is $p_M(t,o,x)$ approximately the largest?

(b) Can we find balls $B_t = B(x_t, \frac{1}{4}\sqrt{t})$ with $|x_t| \le 4\sqrt{t}$ such that

$$\lim_{t\to+\infty} p_M(t,o,B_t) = \lim_{t\to+\infty} \int_{B(x_t,\frac{1}{4}\sqrt{t})} p_M(t,o,y)dy = 0?$$

Such balls contain an unusually small amount of heat given their sizes and locations.

The answer to the first question is that the heat kernel $p_M(t,o,x)$ is the largest when x is relatively close to $(0,0,-\sqrt{t})$, down in the cylinder E_1 where its approximate value is $1/t$. For comparison, note that $p_M(t,o,(0,0,+\sqrt{t}))$ is of the order of $1/t^{3/2}$. However, the ball $B((0,0,-\sqrt{t}),\frac{1}{4}\sqrt{t})$ has small volume, of order $\sqrt{t}$, so that $p_M(t,o,B((0,0,-\sqrt{t}),\frac{1}{4}\sqrt{t}))$ is approximately equal to $1/\sqrt{t}$. In the slab around the xy plane, $p_M(t,o,B((\sqrt{t},\sqrt{t},0),\frac{1}{4}\sqrt{t}))$ is approximately $1/\log t$. However, in the largest end, E_3, where the heat kernel is the smallest, $p_M(t,o,B((0,0,\sqrt{t}),\frac{1}{4}\sqrt{t}))$ is approximately 1; see Figure 2. In terms of heat diffusion, the heat kernel describes punctual temperature, and the integral over a ball is the caloric content. The caloric content of balls of a given radius

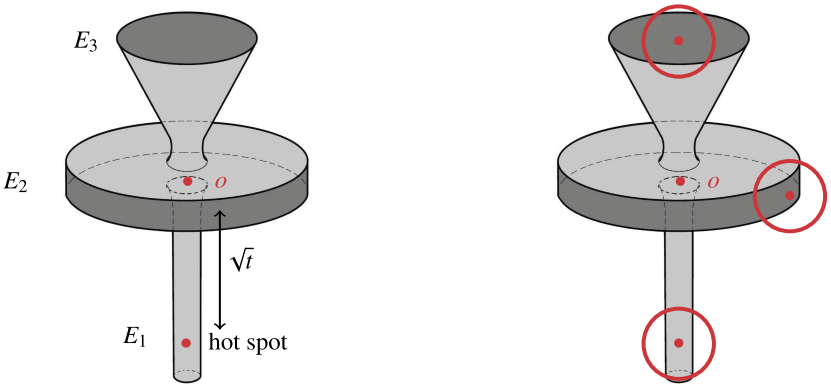

FIGURE 2
The solid body of Example 8.3: the "hot spot" and balls of interest (scale is $\sqrt{t}$).

depends on both the local temperature and the local volume growth. These computations illustrate the detailed information provided by Theorem 8.2.

9. PARABOLIC MANIFOLDS WITH FINITELY MANY HARNACK ENDS

It turns out that the case when the weighted manifold $M = M_1\#\cdots\#M_k$ is parabolic (i.e., Brownian motion is recurrent) is harder, and the treatment remains incomplete despite good results presented in [13] and a forthcoming companion paper. To give an idea of what is expected, let us consider the very simple case when each M_i is a surface of revolution in $\mathbb{R}^3$ associated with the rotation of the graph of

$$\phi : [0, +\infty) \to \big[0, +\infty), x \mapsto z = \phi_i(x) \text{ with } \phi_i(x) = x^{\alpha_i} \text{ for } x > 2 \text{ and } \alpha_i \in (0, 1\big].$$

The behavior of ϕ near 0 is $\sqrt{x}$ so that the surface M_i is smooth. These smooth surfaces in $\mathbb{R}^3$ are equipped with their natural Riemannian metric and measure. Each such surface is Harnack and (RCA) and its heat kernel satisfies

$$\frac{c}{V_i((s,\phi_i(s)),\sqrt{t})} \le p_{M_i}\big(t,(s,\phi_i(s)),(s,\phi_i(s))\big) \le \frac{C}{V_i((s,\phi_i(s)),\sqrt{t})}$$

with $V_i((s,\phi_i(s)),r)$ approximately equal to

$$\begin{cases} r^2 & \text{if } 0 \le r < \max\{1, s^{\alpha_i}\}, \\ s^{\alpha_i} r & \text{if } \max\{1, s^{\alpha_i}\} \le r \le s, \\ r^{1+\alpha_i} & \text{if } \max\{1, s\} \le r. \end{cases}$$

Reference [13] gives sharp global two-sided estimates for $M = M_1\#\cdots\#M_k$ as above. Here are some highlights:

- If for all $i \in \{1,\dots,k\}$, $\alpha_i = \alpha \in (0,1)$, then $M = M_1\#\cdots\#M_k$ is Harnack.
- If $k \ge 2$ and for all $i \in \{1,\dots,k\}$, $\alpha_i = 1$, then M is doubling but does not satisfy the Poincaré inequality in large balls centered at a fixed point o in M. For large t and a fixed point $o \in M$, $p_M(t,o,o)$ is approximately equal to $1/t$, whereas if x, y are in different ends, at distance $\sqrt{t}$ from o, then $p_M(t,x,y)$ is approximately equal to $1/(t\log^2(1+t))$.
- In all cases, $p_M(t,o,o)$ is approximately equal to $1/\max_{1\le i\le k}\{V_i(\sqrt{t})\}$ where $V_i(r) = \begin{cases} r^2 & \text{if } r \in (0,1), \\ r^{1+\alpha_i} & \text{if } r \ge 1. \end{cases}$

The simplest and most important thing to note is that $p_M(t,o,o)$ is now controlled by the volume of the largest end whereas, in the case when each end is nonparabolic (i.e., Section 7), $p_M(t,o,o)$ is controlled by the volume of the smallest end. The first observation of this phenomena in a simplified model case which appeared in [7]. It is also worth stressing that the following problem remains open (see [14] for additional details on what is known).

Problem 9.1. Prove a sharp two-sided heat kernel estimate for $M = M_1\#\cdots\#M_k$ under the assumption that each M_i is Harnack, parabolic, and satisfies (RCA).

10. MIXED BOUNDARY CONDITIONS ON HARNACK MANIFOLDS

Although we did not insist much on this aspect, the results discussed in Sections 7, 8 and 9 depend in a significant way on our ability to derive sharp heat kernel estimates with Dirichlet boundary condition on domains obtained from a Harnack manifold by deleting a compact set with nonempty interior and smooth boundary; see [15, 16, 18, 19]. One is then naturally led to consider the problem of heat kernel estimates for manifolds with boundary and mixed boundary conditions (Neumann and Dirichlet). This requires new types of hypotheses. Assume that M is a complete, weighted Riemannian manifold with boundary δM, and that Ω is an open subset of M such that $M \setminus \Omega \subseteq \delta M$. To simplify the presentation, assume that $\partial\Omega = M \setminus \Omega$ has finitely many connected components which are manifolds with boundary. Our new object of interest here is the minimal heat kernel of Ω, $p_\Omega(t, x, y)$, which is the heat kernel of the Dirichlet form $(L^2(\Omega, \mu), \int_\Omega |\nabla f|^2 d\mu, W_0^1(\Omega))$ where the domain $W_0^1(\Omega)$ is the closure of smooth, compactly-supported functions in Ω for the norm $(\int_\Omega (|f|^2 + |\nabla f|^2) d\mu)^{1/2}$. The corresponding heat equation has Dirichlet boundary condition along $\partial\Omega$ and Neumann boundary condition along the rest of the original boundary of M, δM. In [20], sharp heat kernel estimates are derived under the condition that (1) (M, μ) is Harnack, and (2) Ω is a uniform domain in M. Before we describe what *uniform* means, observe that the distance between two points x, y in Ω, $d_\Omega(x, y)$, is the same as the distance between x and y in M, $d_M(x, y)$.

To say that Ω is uniform in M with constant C is to say that, for any pair of points $x, y \in \Omega$, there is a rectifiable curve parametrized by arc length, $\gamma_{xy} : [0, T_{xy}] \to \Omega$, joining x to y, of length $T_{x,y} \leq C d_\Omega(x, y)$, and satisfying $d(\gamma_{xy}(s), M \setminus \Omega) \geq C^{-1} \min\{s, T_{xy} - s\}$, for all $s \in [0, T_{xy}]$. In words, the curve γ_{xy} is roughly of optimal length and, when moving away from x (or y) along γ_{xy}, one also moves away from the boundary in a roughly linear fashion. For instance, the open upper-half plane in the closed upper-half plane is uniform, but the open infinite strip $\Omega = \{(x, y) : -1 < y < 1, x \in \mathbb{R}\}$ is NOT uniform in its closure because one cannot escape from being close to the boundary; see Figure 3.

When Ω is uniform in M, it admits a harmonic profile h_Ω, which is positive harmonic in Ω, vanishing continuously along $\partial\Omega$ (this function has vanishing normal derivative along $\delta M \setminus \partial\Omega$).

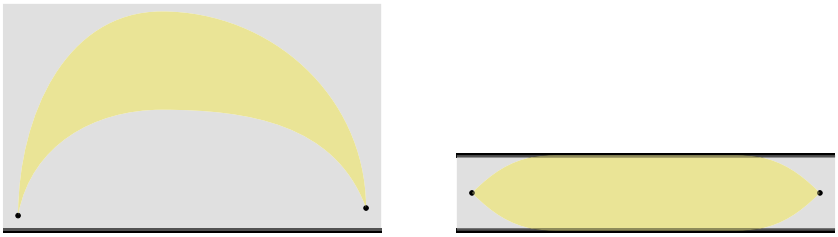

FIGURE 3
The upper-half space is uniform; the band is not.

Theorem 10.1 (See [20, THEOREM 5.11]). *Referring to the setting outlined above, assume that* (M, μ) *is Harnack and that* Ω *is uniform in* M *with harmonic profile* h_Ω. *Then there are constants* c_1, c_2 *such that*

$$g_{\Omega,c_1}(t,x,y) \le p_\Omega(t,x,y) \le g_{\Omega,c_2}(t,x,y), \tag{10.1}$$

where

$$g_{\Omega,c}(t,x,y) = \frac{h_\Omega(x)h_\Omega(y)}{c\,V_{h_\Omega^2}(x,\sqrt{t})}\exp\left(-c\frac{d_M(x,y)^2}{t}\right).$$

When reading this theorem, recall that

$$d_\Omega(x,y) = d_M(x,y) \quad \text{and} \quad V_{h_\Omega^2}(x,r) = \int_{B_M(x,r)} h_\Omega^2(y)\mu(dy).$$

The lack of symmetry between x and y is intentional; because p_Ω is symmetric, a symmetric estimate can easily be derived from that stated here. This is a satisfactory and useful result from a theoretical viewpoint, but detailed applications require estimating the profile h_Ω, which is a difficult problem.

11. MIXED BOUNDARY CONDITIONS ON MANIFOLDS WITH ENDS

The techniques used in the previous section to study uniform domains in complete Riemannian manifolds with boundary can be implemented together with the techniques of Section 7 to study the minimal heat kernel p_Ω of a domain Ω in a complete Riemannian manifold $M = K \cup E_1 \cup \cdots \cup E_k$ with boundary when $\partial\Omega \subset \delta M$. A simplistic, yet interesting example is depicted in Figure 4. Can you guess the behavior of $p_\Omega(t,o,o)$ in Figure 4? The answer is $1/(t\log^2 t)$ because, far from o, one of the three cones is free of Dirichlet boundary condition. If each cone had at least one of its sides contained in $\partial\Omega$, the behavior would depend in an explicit way on the apertures of the cones and whether each cone has one or two sides contained in $\partial\Omega$. See [6] for a detailed discussion and general results in this direction, including complete two-sided heat kernel estimates for such mixed boundary problems.

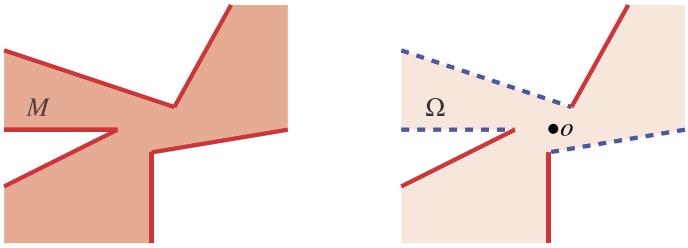

FIGURE 4
Sketch (corners should be rounded) of M (left, thick boundary lines are part of M) and Ω (right) with "Dirichlet boundary" $\partial\Omega \subseteq \delta M$ (dashed) not part of Ω.

12. ATTACHMENTS ALONG NONCOMPACT SUBMANIFOLDS

The basic ideas implemented in the study of connected sums above can be described informally in greater generality as follows. Given a (metrically complete) manifold M, identify large chunks, hopefully, finitely many, $M_1, \dots, M_k$, which, taken by themselves, are Harnack manifolds. Each chunk has attachment boundaries along which they are attached to each other to form the manifold M. Call Att_i the attachment boundary of M_i and set $\Omega_i = M_i \setminus \mathrm{Att}_i$. If each Ω_i is uniform in M_i, $1 \le i \le k$, not only can we have good estimates for the heat kernel p_{M_i} of M_i (because M_i is Harnack), but we can also estimate the minimal heat kernel p_{Ω_i} of Ω_i (this heat kernel satisfies the Dirichlet boundary condition along Att_i). For the next step, it may be necessary to make further assumptions about the manifolds M_i and their open subsets Ω_i (see, for instance, conditions (NP) and (RCA) above). Now, find a way to use the known information regarding the different large chunks M_i to reconstruct and estimate the heat kernel p_M of M. Sections 7–9 above describe how these ideas apply successfully to connected sums (i.e., compact attachments). The article [12] is, so far, the lone published attempt to carry out this approach when two large chunks are glued along a noncompact attachment boundary. Emily Dautenhahn and the author are working on applying these ideas beyond the cases treated in [12].

The results of Sections 7–12 should ultimately be developed in the more abstract context of Dirichlet spaces so as to include ends that satisfy the Harnack inequalities that appear in the context of fractals (see [3] and the references therein). This would allow for the treatment of a larger class of Riemannian manifolds, as the geometry of a Riemannian manifold at infinity can mimic that of a fractal object.

ACKNOWLEDGMENTS

A. Grigor'yan and the author have shared their interests in heat kernel estimates and their applications to potential theory, probability theory, and geometry ever since they met, a few years after having proved, independently, some version of Theorem 4.4. This article is partly the result of this long lasting collaboration and friendship. Thanks to T. Coulhon, N. Varopoulos, and to all my colleagues, collaborators, and students who helped me make progress in the study of heat kernels over the years.

FUNDING

The author's research is partially supported by a grant from the National Science Foundation (USA), DMS-2054593.

REFERENCES

[1] D. G. Aronson, Bounds for the fundamental solution of a parabolic equation. *Bull. Amer. Math. Soc.* **73** (1967), 890–896.

[2] D. G. Aronson and J. Serrin, Local behavior of solutions of quasilinear parabolic equations. *Arch. Ration. Mech. Anal.* **25** (1967), 81–122.

[3] M. Barlow and M. Murugan, Stability of the elliptic Harnack inequality. *Ann. of Math. (2)* **187** (2018), no. 3, 777–823.

[4] G. Carron and D. Tewodrose, A rigidity result for metric measure spaces with Euclidean heat kernel. *J. Éc. Polytech. Math.* **9** (2022), 101–154.

[5] T. Coulhon, Ultracontractivity and Nash type inequalities. *J. Funct. Anal.* **141** (1996), no. 2, 510–539.

[6] E. Dautenhahn and L. Saloff-Coste, Heat kernel estimates on manifolds with ends with mixed boundary condition. 2021, arXiv:2108.05790.

[7] E. B. Davies, Non-Gaussian aspects of heat kernel behaviour. *J. Lond. Math. Soc.* **55** (1997), 105–125.

[8] N. Eldredge, M. Gordina, and L. Saloff-Coste, Left-invariant geometries on SU(2) are uniformly doubling. *Geom. Funct. Anal.* **28** (2018), no. 5, 132–1367.

[9] M. Fukushima, Y. Oshima, and M. Takeda, *Dirichlet forms and symmetric Markov processes*. Second revised and extended edn., de Gruyter Stud. Math. 19, Walter de Gruyter & Co., Berlin, 2011.

[10] A. Grigor'yan, The heat equation on non-compact Riemannian manifolds. *Mat. Sb.* **182** (1991), 55–87. Engl. Transl. *Math. USSR Sb.* **72** (1992), 47–77.

[11] A. Grigor'yan, Analytic and geometric background of recurrence and non-explosion of the Brownian motion on Riemannian manifolds. *Bull. Amer. Math. Soc. (N.S.)* **36** (1999), no. 2, 135–249.

[12] A. Grigor'yan and S. Ishiwata, Heat kernel estimates on a connected sum of two copies of $\mathbb{R}^n$ along a surface of revolution. *Glob. Stoch. Anal.* **2** (2012), no. 1, 29–65.

[13] A. Grigor'yan, S. Ishiwata, and L. Saloff-Coste, Heat kernel estimates on connected sums of parabolic manifolds. *J. Math. Pures et Appl.* **113** (2018), 155–194.

[14] A. Grigor'yan, S. Ishiwata, and L. Saloff-Coste, Geometric analysis on manifolds with ends. In *Analysis and partial differential equations on manifolds, fractals and graphs*, edited by A. Grigor'yan and Y. Sun, pp. 325–344, De Gruyter, Berlin, Boston. 2021. DOI 10.1515/9783110700763-011.

[15] A. Grigor'yan and L. Saloff-Coste, Dirichlet heat kernel in the exterior of a compact set. *Comm. Pure Appl. Math.* **55** (2002), 93–133.

[16] A. Grigor'yan and L. Saloff-Coste, Hitting probabilities for Brownian motion on Riemannian manifolds. *J. Math. Pures et Appl.* **81** (2002), 115–142.

[17] A. Grigor'yan and L. Saloff-Coste, Stability results for Harnack inequalities. *Ann. Inst. Fourier* **55** (2005), 825–890.

[18] A. Grigor'yan and L. Saloff-Coste, Heat kernel on manifolds with ends. *Ann. Inst. Fourier* **59** (2009), 1917–1997.

[19] A. Grigor'yan and L. Saloff-Coste, Surgery of the Faber–Krahn inequality and applications to heat kernel bounds. *Nonlinear Anal.* **131** (2016), 243–272.

[20] P. Gyrya and L. Saloff-Coste, *Neumann and Dirichlet heat kernels in inner uniform domains*. Astérisque 336, Société Mathématique de France, 2011.

[21] P. Li and S.-T. Yau, On the Schrödinger equation and the eigenvalue problem. *Comm. Math. Phys.* **88** (1983), 309–318.

[22] U. Mosco, Composite media and asymptotic Dirichlet forms. *J. Funct. Anal.* **123** (1994), no. 2, 368–421.

[23] J. Moser, A Harnack inequality for parabolic differential equations. *Comm. Pure Appl. Math.* **16** (1964), 101–134. Correction in **20** (1967), 231–236.

[24] J. Nash, Continuity of solutions of parabolic and elliptic equations. *Amer. J. Math.* **80** (1958), 931–954.

[25] F. Porper and S. Eidel'man, Two-sided estimates of fundamental solutions of second-order parabolic equations and some applications. *Russian Math. Surveys* **39** (1984), 119–178.

[26] L. Saloff-Coste, A note on Poincaré, Sobolev and Harnack inequalities. *Duke Math. J.* **65**, *Int. Math. Res. Not. IMRN* **2** (1992), 27–38.

[27] L. Saloff-Coste, *Aspects of Sobolev-type inequalities*. London Math. Soc. Lecture Note Ser. 289, Cambridge University Press, Cambridge, 2002.

[28] K. T. Sturm, Analysis on local Dirichlet spaces. II. Upper Gaussian estimates for the fundamental solutions of parabolic equations. *Osaka J. Math.* **32** (1995), no. 2, 275–312.

[29] K. T. Sturm, Analysis on local Dirichlet spaces. III. The parabolic Harnack inequality. *J. Math. Pures Appl. (9)* **75** (1996), no. 3, 273–297.

[30] C.-J. Sung, L.-F. Tam, and J. Wang, Spaces of harmonic functions. *J. Lond. Math. Soc. (2)* **3** (2000), 789–806.

LAURENT SALOFF-COSTE

567 Malott Hall, Cornell University, Ithaca NY, 14850, USA, lps2@cornell.edu