1
Introduction

We found it difficult to choose a title for this book. Clearly we are not covering the theory of Markov processes, Gaussian processes, and local times in one volume. A more descriptive title would have been "A Study of the Local Times of Strongly Symmetric Markov Processes Employing Isomorphisms That Relate Them to Certain Associated Gaussian Processes." The innovation here is that we can use the well-developed theory of Gaussian processes to obtain new results about local times.

Even with the more restricted title there is a lot of material to cover. Since we want this book to be accessible to advanced graduate students, we try to provided a self-contained development of the Markov process theory that we require. Next, since the crux of our approach is that we can use sophisticated results about the sample path properties of Gaussian processes to obtain similar sample path properties of the associated local times, we need to present this aspect of the theory of Gaussian processes. Furthermore, interesting questions about local times lead us to focus on some properties of Gaussian processes that are not usually featured in standard texts, such as processes with spectral densities or those that have infinitely divisible squares. Occasionally, as in the study of the p-variation of sample paths, we obtain new results about Gaussian processes.

Our third concern is to present the wonderful, mysterious isomorphism theorems that relate the local times of strongly symmetric Markov processes to associated mean zero Gaussian processes. Although some inkling of this idea appeared earlier in Brydges, Fröhlich and Spencer (1982) we think that credit for formulating it in an intriguing and usable format is due to E. B. Dynkin (1983), (1984). Subsequently, after our initial paper on this subject, Marcus and Rosen (1992d), in which we use Dynkin's Theorem, N. Eisenbaum (1995) found an unconditioned isomorphism that seems to be easier to use. After this Eisenbaum, Kaspi,

Cambridge University Press
978-0-521-86300-1 - Markov Processes, Gaussian Processes, and Local Times
Michael B. Marcus and Jay Rosen
Excerpt
More information

Marcus, Rosen and Shi (2000) found a third isomorphism theorem, which we refer to as the Generalized Second Ray–Knight Theorem, because it is a generalization of this important classical result.

Dynkin's and Eisenbaum's proofs contain a lot of difficult combinatorics, as does our proof of Dynkin's Theorem in Marcus and Rosen (1992d). Several years ago we found much simpler proofs of these theorems. Being able to present this material in a relatively simple way was our primary motivation for writing this book.

The classical Ray–Knight Theorems are isomorphisms that relate local times of Brownian motion and squares of independent Brownian motions. In the three isomorphism theorems we just referred to, these theorems are extended to give relationships between local times of strongly symmetric Markov processes and the squares of associated Gaussian processes. A Markov process with symmetric transition densities is strongly symmetric. Its associated Gaussian process is the mean zero Gaussian process with covariance equal to its 0-potential density. (If the Markov process, say X, does not have a 0-potential, one can consider $\widehat{X}$, the process X killed at the end of an independent exponential time with mean $1/\alpha$. The 0-potential density of $\widehat{X}$ is the α-potential density of X.)

As an example of how the isomorphism theorems are used and of the kinds of results we obtain, we mention that we show that there exists a jointly continuous version of the local times of a strongly symmetric Markov process if and only if the associated Gaussian process has a continuous version. We obtain this result as an equivalence, without obtaining conditions that imply that either process is continuous. However, conditions for the continuity of Gaussian processes are known, so we know them for the joint continuity of the local times.

M. Barlow and J. Hawkes obtained a sufficient condition for the joint continuity of the local times of Lévy processes in Barlow (1985) and Barlow and Hawkes (1985), which Barlow showed, in Barlow (1988), is also necessary. Gaussian processes do not enter into the proofs of their results. (Although they do point out that their conditions are also necessary and sufficient conditions for the continuity of related stationary Gaussian processes.) This stimulating work motivated us to look for a more direct link between Gaussian processes and local times and led us to Dynkin's isomorphism theorem.

We must point out that the work of Barlow and Hawkes just cited applies to all Lévy processes whereas the isomorphism theorem approach that we present applies only to symmetric Lévy processes. Nevertheless, our approach is not limited to Lévy processes and also opens up the possibility of using Gaussian process theory to obtain many other interesting properties of local times.

Another confession we must make is that we do not really understand the actual relationship between local times of strongly symmetric Markov processes and their associated Gaussian processes. That is, we have several functional equivalences between these disparate objects and can manipulate them to obtain many interesting results, but if one asks us, as is often the case during lectures, to give an intuitive description of how local times of Markov processes and Gaussian process are related, we must answer that we cannot. We leave this extremely interesting question to you. Nevertheless, there now exist interesting characterizations of the Gaussian processes that are associated with Markov processes. We say more about this in our discussion of the material in Chapter 13.

The isomorphism theorems can be applied to very general classes of Markov processes. In this book, with the exception of Chapter 13, we consider Borel right processes. To ease the reader into this degree of generality, and to give an idea of the direction in which we are going, in Chapter 2 we begin the discussion of Markov processes by focusing on Brownian motion. For Brownian motion these isomorphisms are old stuff but because, in the case of Brownian motion, the local times of Brownian motion are related to the squares of independent Brownian motion, one does not really leave the realm of Markov processes. That is, we think that in the classical Ray–Knight Theorems one can view Brownian motion as a Markov process, which it is, rather than as a Gaussian process, which it also is.

Chapters 2–4 develop the Markov process material we need for this book. Naturally, there is an emphasis on local times. There is also an emphasis on computing the potential density of strongly symmetric Markov processes, since it is through the potential densities that we associate the local times of strongly symmetric Markov processes with Gaussian processes. Even though Chapter 2 is restricted to Brownian motion, there is a lot of fundamental material required to construct the σ-algebras of the probability space that enables us to study local times. We do this in such a way that it also holds for the much more general Markov processes studied in Chapters 3 and 4. Therefore, although many aspects of Chapter 2 are repeated in greater generality in Chapters 3 and 4, the latter two chapters are not independent of Chapter 2.

In the beginning of Chapter 3 we study general Borel right processes with locally compact state spaces but soon restrict our attention to strongly symmetric Borel right processes with continuous potential densities. This restriction is tailored to the study of local times of Markov processes via their associated mean zero Gaussian processes. Also, even though this restriction may seem to be significant from the perspective of the general theory of Markov processes, it makes it easier to introduce the beautiful theory of Markov processes. We are able to obtain many deep and interesting results, especially about local times, relatively quickly and easily. We also consider h-transforms and generalizations of Kac's Theorem, both of which play a fundamental role in proving the isomorphism theorems and in applying them to the study of local times.

Chapter 4 deals with the construction of Markov processes. We first construct Feller processes and then use them to show the existence of Lévy processes. We also consider several of the finer properties of Borel right processes. Lastly, we construct a generalization of Borel right processes that we call local Borel right processes. These are needed in Chapter 13 to characterize associated Gaussian processes. This requires the introduction of Ray semigroups and Ray processes.

Chapters 5–7 are an exposition of sample path properties of Gaussian processes. Chapter 5 deals with structural properties of Gaussian processes and lays out the basic tools of Gaussian process theory. One of the most fundamental tools in this theory is the Borell, Sudakov–Tsirelson isoperimetric inequality. As far as we know this is stated without a complete proof in earlier books on Gaussian processes because the known proofs relied on the Brun–Minkowski inequality, which was deemed to be too far afield to include its proof. We give a new, analytical proof of the Borell, Sudakov–Tsirelson isoperimetric inequality due to M. Ledoux in Section 5.4.

Chapter 6 presents the work of R. M. Dudley, X. Fernique and M. Talagrand on necessary and sufficient conditions for continuity and boundedness of sample paths of Gaussian processes. This important work has been polished throughout the years in several texts, Ledoux and Talagrand (1991), Fernique (1997), and Dudley (1999), so we can give efficient proofs. Notably, we give a simpler proof of Talagrand's necessary condition for continuity involving majorizing measures, also due to Talagrand, than the one in Ledoux and Talagrand (1991). Our presentation in this chapter relies heavily on Fernique's excellent monograph, Fernique (1997).

Chapter 7 considers uniform and local moduli of continuity of Gaussian processes. We treat this question in general in Section 7.1. In most of the remaining sections in this chapter, we focus our attention on real-valued Gaussian processes with stationary increments, $\{G(t), t \in R^1\}$, for which the increments variance, $\sigma^2(t-s) := E(G(t)-G(s))^2$, is relatively smooth. This may appear old fashioned to the Gaussian purist but it is exactly these processes that are associated with real-valued Lévy processes. (And Lévy processes with values in R^n have local times only when $n = 1$.) Some results developed in this section and its applications in Section 9.5 have not been published elsewhere.

Chapters 2–7 develop the prerequisites for the book. Except for Section 3.7, the material at the end of Chapter 4 relating to local Borel right processes, and a few other items that are referenced in later chapters, they can be skipped by readers with a good background in the theory of Gaussian and Markov processes.

In Chapter 8 we prove the three main isomorphism theorems that we use. Even though we are pleased to be able to give simple proofs that avoid the difficult combinatorics of the original proofs of these theorems, in Section 8.3 we give the combinatoric proofs, both because they are interesting and because they may be useful later on.

Chapter 9 puts everything together to give sample path properties of local times. Some of the proofs are short, simply a reiteration of results that have been established in earlier chapters. At this point in the book we have given all the results in our first two joint papers on local times and isomorphism theorems (Marcus and Rosen, 1992a, 1992d). We think that we have filled in all the details and that many of the proofs are much simpler. We have also laid the foundation to obtain other interesting sample path properties of local times, which we present in Chapters 10–13.

In Chapter 10 we consider the p-variation of the local times of symmetric stable processes $1 < p \leq 2$ (this includes Brownian motion). To use our isomorphism theorem approach we first obtain results on the p-variation of fractional Brownian motion that generalize results of Dudley (1973) and Taylor (1972) that were obtained for Brownian motion. These are extended to the squares of fractional Brownian motion and then carried over to give results about the local times of symmetric stable processes.

Chapter 11 presents results of Bass, Eisenbaum and Shi (2000) on the range of the local times of symmetric stable processes as time goes to infinity and shows that the most visited site of such processes is transient. Our approach is different from theirs. We use an interesting bound for the behavior of stable processes in a neighborhood of the origin due to Molchan (1999), which itself is based on properties of the reproducing kernel Hilbert spaces of fractional Brownian motions.

In Chapter 12 we reexamine Ray's early isomorphism theorem for the h-transform of a transient regular symmetric diffusion, Ray (1963) and give our own, simpler version. We also consider the Markov properties of the local times of diffusions.

In Chapter 13, which is based on recent work of N. Eisenbaum and H. Kaspi that appears in Eisenbaum (2003), Eisenbaum (2005), and Eisenbaum and Kaspi (2006), we take up the problem of characterizing associated Gaussian processes. To obtain several equivalencies we must generalize Borel right processes to what we call local Borel right processes. In Theorem 13.3.1 we see that associated Gaussian processes are just a little less general than the class of Gaussian processes that have infinitely divisible squares. Gaussian processes with infinitely divisible squares are characterized in Griffiths (1984) and Bapat (1989). We present their results in Section 13.2.

We began our joint research that led to this book over 19 years ago. In the course of this time we received valuable help from R. Adler, M. Barlow, H. Kaspi, E. B. Dynkin, P. Fitzsimmons, R. Getoor, E. Giné, M. Talagrand, and J. Zinn. We express our thanks and gratitude to them. We also acknowledge the help of P.-A. Meyer.

In the preparation of this book we received valuable assistance and advice from O. Daviaud, S. Dhamoon, V. Dobric, N. Eisenbaum, S. Evans, P. Fitzsimmons, C. Houdré, H. Kaspi, W. Li, and J. Rosinski. We thank them also.

We are also grateful for the continued support of the National Science Foundation and PSC–CUNY throughout the writing of this book.

1.1 Preliminaries

In this book $\mathcal{Z}$ denotes the integers both positive and negative and $\mathbb{N}$ or sometimes $\mathcal{N}$ denotes the the positive integers including 0. R^1 denotes the real line and R_+ the positive half line (including zero). $\overline{R}$ denotes the extended real line $[-\infty, \infty]$. R^n denotes n-dimensional space and $|\cdot|$ denotes Euclidean distance in R^n. We say that a real number a is positive if $a \geq 0$. To specify that $a > 0$, we might say that it is strictly positive. A similar convention is used for negative and strictly negative.

Measurable spaces: A measurable space is a pair $(\Omega, \mathcal{F})$, where Ω is a set and $\mathcal{F}$ is a sigma-algebra of subsets of Ω. If Ω is a topological space, we use $\mathcal{B}(\Omega)$ to denote the Borel σ-algebra of Ω. Bounded $\mathcal{B}(\Omega)$ measurable functions on Ω are denoted by $\mathcal{B}_b(\Omega)$.

Let $t \in R_+$. A filtration of $\mathcal{F}$ is an increasing family of sub σ-algebras $\mathcal{F}_t$ of $\mathcal{F}$, that is, for $0 \leq s < t < \infty$, $\mathcal{F}_s \subset \mathcal{F}_t \subset \mathcal{F}$ with $\mathcal{F} = \cup_{0 \leq t < \infty} \mathcal{F}_t$.

(Sometimes we describe this by saying that $\mathcal{F}$ is filtered.) To emphasize a specific filtration $\mathcal{F}_t$ of $\mathcal{F}$, we sometimes write $(\Omega, \mathcal{F}, \mathcal{F}_t)$.

Let $\mathcal{M}$ and $\mathcal{N}$ denote two σ-algebras of subsets of Ω. We use $\mathcal{M} \vee \mathcal{N}$ to denote the σ-algebra generated by $\mathcal{M} \cup \mathcal{N}$.

Probability spaces: A probability space is a triple $(\Omega, \mathcal{F}, P)$, where $(\Omega, \mathcal{F})$ is measurable space and P is a probability measure on Ω. A random variable, say X, is a measurable function on $(\Omega, \mathcal{F}, P)$. In general we let E denote the expectation operator on the probability space. When there are many random variables defined on $(\Omega, \mathcal{F}, P)$, say $Y, Z, \ldots$, we use E_Y to denote expectation with respect to Y. When dealing with a probability space, when it seems clear what we mean, we feel free to use E or even expressions like E_Y without defining them. As usual, we let ω denote the elements of Ω. As with E, we often use ω in this context without defining it.

When X is a random variable we call a number a a median of X if

$$P(X \le a) \ge \tfrac{1}{2} \qquad \text{and} \qquad P(X \ge a) \ge \tfrac{1}{2}. \tag{1.1}$$

Note that a is not necessarily unique.

A stochastic process X on $(\Omega, \mathcal{F}, P)$ is a family of measurable functions $\{X_t, t \in I\}$, where I is some index set. In this book, t usually represents "time" and we generally consider $\{X_t, t \in R_+\}$. $\sigma(X_r; r \le t)$ denotes the smallest σ-algebra for which $\{X_r; r \le t\}$ is measurable. Sometimes it is convenient to describe a stochastic process as a random variable on a function space, endowed with a suitable σ-algebra and probability measure.

In general, in this book, we reserve $(\Omega, \mathcal{F}, P)$ for a probability space. We generally use $(S, \mathcal{S}, \mu)$ to indicate more general measure spaces. Here μ is a positive (i.e., nonnegative) σ-finite measure.

Function spaces: Let f be a measurable function on $(S, \mathcal{S}, \mu)$. The $L^p(\mu)$ (or simply L^p), $1 \le p < \infty$, spaces are the families of functions f for which $\int_S |f(s)|^p \, d\mu(s) < \infty$ with

$$\|f\|_p := \left(\int_S |f(s)|^p \, d\mu(s) \right)^{1/p}. \tag{1.2}$$

Sometimes, when we need to be precise, we may write $\|f\|_{L^p(S)}$ instead of $\|f\|_p$. As usual we set

$$\|f\|_\infty = \sup_{s \in S} |f(s)|. \tag{1.3}$$

These definitions have analogs for sequence spaces. For $1 \le p < \infty$, ℓ_p

CAMBRIDGE

Cambridge University Press
978-0-521-86300-1 - Markov Processes, Gaussian Processes, and Local Times
Michael B. Marcus and Jay Rosen
Excerpt
More information

is the family of sequences $\{a_k\}_{k=0}^{\infty}$ of real or complex numbers such that $\sum_{k=0}^{\infty}|a_k|^p<\infty$. In this case, $\|a_k\|_p := (\sum_{k=0}^{\infty}|a_k|^p)^{1/p}$ and $\|a_k\|_\infty := \sup_{0\le k<\infty}|a_k|$. We use ℓ_p^n to denote sequences in ℓ_p with n elements.

Let m be a measure on a topological space $(S,\mathcal{S})$. By an approximate identity or δ-function at y, with respect to m, we mean a family $\{f_{\epsilon,y}\,;\,\epsilon>0\}$ of positive continuous functions on S such that $\int f_{\epsilon,y}(x)\,dm(x)=1$ and each $f_{\epsilon,y}$ is supported on a compact neighborhood K_ϵ of y with $K_\epsilon\downarrow\{y\}$ as $\epsilon\to 0$.

Let f and g be two real-valued functions on R^1. We say that f is asymptotic to g at zero and write $f\sim g$ if $\lim_{x\to 0}f(x)/g(x)=1$. We say that f is comparable to g at zero and write $f\approx g$ if there exist constants $0<C_1\le C_2<\infty$ such that $C_1\le\liminf_{x\to 0}f(x)/g(x)$ and $\limsup_{x\to 0}f(x)/g(x)\le C_2$. We use essentially the same definitions at infinity.

Let f be a function on R^1. We use the notation $\lim_{y\uparrow\uparrow x}f(y)$ to be the limit of $f(y)$ as y increases to x, for all $y<x$, that is, the left-hand (or simply left) limit of f at x.

Metric spaces: Let (S,τ) be a locally compact metric or pseudo-metric space. A pseudo-metric has the same properties as a metric except that $\tau(s,t)=0$ does not imply that $s=t$. Abstractly, one can turn a pseudo-metric into a metric by making the zeros of the pseudo-metric into an equivalence class, but in the study of stochastic processes pseudo-metrics are unavoidable. For example, suppose that $X=\{X(t),t\in[0,1]\}$ is a real-valued stochastic process. In studying sample path properties of X it is natural to consider $(R^1,|\cdot|)$, a metric space. However, X may be completely determined by an L^2 metric, such as

$$d(s,t):=d_X(s,t):=(E(X(s)-X(t))^2)^{1/2} \tag{1.4}$$

(and an additional condition such as $EX^2(t)=1$). Therefore, it is natural to also consider the space (R^1,d). This may be a pseudo-metric space since d need not be a metric on R^1.

If $A\subset S$, we set

$$\tau(s,A):=\inf_{u\in A}\tau(s,u). \tag{1.5}$$

We use $C(S)$ to denote the continuous functions on S, $C_b(S)$ to denote the bounded continuous functions on S, and $C_b^+(S)$ to denote the positive bounded continuous functions on S. We use $C_\kappa(S)$ to denote the continuous functions on S with compact support; $C_0(S)$ denotes the functions on S that go to 0 at infinity. Nevertheless, $C_0^\infty(S)$ denotes infinitely differentiable functions on S with compact support (whenever S

is a space for which this is defined). In all these cases we mean continuity with respect to the metric or pseudo-metric τ.

We say that a function is locally uniformly continuous on a measurable set in (S,τ) if it is uniformly continuous on all compact subsets of (S,τ). We say that a sequence of functions converges locally uniformly on (S,τ) if it converges uniformly on all compact subsets of (S,τ).

Separability: Let T be a separable metric space, and let $X=\{X(t),t\in T\}$ be a stochastic process on $(\Omega,\mathcal{F},P)$ with values in $\overline{R}^n$. X is said to be separable if there is a countable set $D\subset T$ and a P-null set $\Lambda\subset\mathcal{F}$ such that, for any open set $U\subset T$ and closed set $A\subset\overline{R}^n$,

$$\{X(t)\in A,t\in D\cap U\}/\{X(t)\in A,t\in U\}\subset\Lambda. \tag{1.6}$$

If X is separable and $U\subset T$ is an open set and Λ is as above, then $\omega\notin\Lambda$ implies

$$\begin{aligned}\sup_{t\in D\cap U}|X(t,\omega)| &= \sup_{t\in U}|X(t,\omega)| \\ \inf_{t\in D\cap U}|X(t,\omega)| &= \inf_{t\in U}|X(t,\omega)|.\end{aligned} \tag{1.7}$$

If T is a separable metric space, every stochastic process $X=\{X(t),t\in T\}$ with values in $\overline{R}^n$ has a separable version $\widetilde{X}=\{\widetilde{X}(t),t\in T\}$, that is, $P\left(\widetilde{X}(t)=X(t)\right)=1$, for all $t\in T$, and $\widetilde{X}$ is separable for some D and Λ.

If X is stochastically continuous, that is, $\lim_{t\to t_0}P(|X(t)-X(t_0)|>\epsilon)=0$, for every $\epsilon>0$ and $t_0\in T$, then any countable dense set $V\subset T$ serves as the set D in the separability condition (sometimes called the separability set). The P-null set Λ generally depends on the choice of V.

Fourier transform: We often give results with precise constants, so we need to describe what version of the Fourier transform we are using. Let $f\in L^2(R^1)$. Consistent with the standard definition of the characteristic function, the Fourier transform $\widehat{f}$ of f is defined by

$$\widehat{f}(\lambda)=\int_{-\infty}^{\infty}e^{i\lambda x}f(x)\,dx, \tag{1.8}$$

where the integral exists in the L^2 sense. The inverse Fourier transform is given by

$$f(x)=\frac{1}{2\pi}\int_{-\infty}^{\infty}e^{-i\lambda x}\widehat{f}(\lambda)\,d\lambda. \tag{1.9}$$

With this normalization, Parseval's Theorem is

$$\int_{-\infty}^{\infty} f(x)\overline{g(x)}\,dx = \frac{1}{2\pi}\int_{-\infty}^{\infty} \widehat{f}(\lambda)\overline{\widehat{g}(\lambda)}\,d\lambda. \tag{1.10}$$