The Annals of Probability
1997, Vol. 25, No. 3, 1536–1540

BOOK REVIEW

S. P. MEYN AND R. L. TWEEDIE, *Markov Chains and Stochastic Stability*. Springer, New York, 1993. xvi + 548 pp.

REVIEW BY VOLKER WIHSTUTZ

University of North Carolina, Charlotte

In *Markov Chains and Stochastic Stability*, which appeared in 1993 in the Springer textbook series *Communication and Control Engineering*, S. P. Meyn and R. L. Tweedie aim to develop a theoretical basis for studying discrete-time Markov processes in general state space as they occur in a wide range of applications, including Markov models in times series and in control and system theory or models with regeneration times. The authors, in their own words, seek to provide "a guide to the general space Markov chain theory and methods for practitioners" as well as "a thorough and rigorous exposition of the results." In the reviewer's opinion, this double goal has been achieved, thanks to both the careful style and the thorough organization of the rich content of the book.

Style. Having in mind practitioners as well as mathematicians interested in theory, the authors took care that readers should not be at a loss. The chapters start with a short assessment of what has been achieved in the field and what needs to be done next, together with a preview of the highlights of the chapter. Every new step is carefully motivated—often by practical needs—and the content of a result is elucidated by thoughtful comments, made to the point and sharpening the essential. (One consequence of this is the length of the book, which comprises about 550 pages in small TEX print.)

Scope. This is an analysis of Markov models in time series, in control and system theory or with regeneration time with the help of the theory represented here, the theory of Markov chains on a general state space; that is, discrete-time Markov processes on a measurable space with countably generated σ-algebra of measurable sets. This includes countable state spaces. A powerful tool (splitting) is used to transfer through regeneration techniques the known results from a countable space to a general state space, in this way unifying the theory.

Received February 1997.

Besides that, the general state space situation is compared to the situation in a topological space (space with a locally compact, separable and metrizable topology). Since Feller chains are too weak for a satisfactory theory, but strong Feller chains too strong from the point of view of applications, the appropriate in-between continuity property had to be found and was found with T-chains.

With a few exceptions, the theory of Markov chains is developed for ψ-irreducible chains; that is, under the condition of a weak, but reasonable communication structure; namely, there is a probability measure ψ such that any set with positive ψ-measure can be reached from every point at some time (with positive probability). In the case of a topological space, the counterpart to ψ-irreducibility is open set irreducibility, where open sets substitute for ψ-positive sets.

Organizing principle. In order to represent the theory in a systematic way and organize the vast amount of old and new results, the theory is developed along steps of increasing strength of stability, where stability is used in a very broad sense, with the lowest level (Part I) being the simple fact that the chain does not fall apart into two separate chains, but enjoys the weak communication structure of ψ-irreducibility.

The second level (Part II) questions how can we be sure that "the chain returns to the 'center' of the space" and "whether it might happen in a finite mean time"; there recurrence properties are discussed. On the third level (Part III) the concern is with the "way the chain 'settles down' to a stable or stationary regime." There the issue is convergence of the n-step transition probabilities to a given invariant probability measure together with the speed of convergence.

Methodological approach. The general state space Markov chain Φ with its (1-step) transition probability $P(x, A)$ and the n-step transition $P^n(x, A)$ is discussed together with and by means of associated *sampled chains* Φ_a defined by the transition probability $K_a(x, A) = \sum_n a(n) P^n(x, A)$, where $a(n)$ is a probability distribution (sampling distribution) on $\mathbb{Z}_+$. In particular, the m-skeleton $\Phi^m = K_{\delta_m}$ is considered, where $a(n) = \delta_m(n)$ is the Dirac measure on $\mathbb{Z}_+$ with footpoint m, and the resolvent chain Φ_{a_ε}, where the transition probability is given by the resolvent

$$K_{a_\varepsilon}(x, A) = (1-\varepsilon)\sum_n \varepsilon^n P^n(x, A), \qquad a(n) = (1-\varepsilon)\varepsilon^n, \qquad 0 < \varepsilon < 1.$$

To achieve a unified approach, the authors make systematic use of three tools, that is, of three kinds of sets distinguished by their uniform property of Doeblin condition type: (pseudo)-atom, small set and petite set.

A (pseudo)-*atom* is a set α out of which all transitions are identical; that is, there is a probability measure ν_α such that, for all x in α, $P(x, A) = \nu_\alpha(A)$.

A set C is called *small* (or ν_m-small), if there exists a natural number m and a measure $\nu_m \neq 0$ such that uniformly for x in C: $P^m(x, A) \geq \nu_m(A)$; and a set C is called *petite* (or ν_a-petite), if there exists a sampling distribution $a(n)$ of $\mathbb{Z}_+$ and a measure $\nu_a \neq 0$, such that uniformly for x for C: $K_a(x, A) \geq \nu_a(A)$.

If the whole state space X is small, we have a version of Doeblin's condition. Any ψ-positive set A contains a ψ-positive ν_m-small set C with $\nu_m(C) > 0$. For any ψ-irreducible chain there is an m-skeleton Φ^m with a ν_1-small set and for any $0 < \varepsilon < 1$, the associated resolvent chain has a ν_1-small set, thus can be split (see the splitting technique described below).

For a ψ-irreducible periodic chain, all petite sets are small sets. Petite sets in a general state space play the role of compact sets in a topological state space. In fact, for an open set irreducible T-chain, petite sets and compact sets are the same.

It is one of the hallmarks of Meyn and Tweedie that with these tools they are able to exploit Nummelin's splitting technique in order to generate a mechanism for transferring the known theory of Markov chains on a countable state space to a general state space, with a chain (the split chain) that possesses a pseudo-atom as an intermediary step. In the easiest situation, the mechanism runs as follows: establish the relationship between Φ and a suitable sampled chain Φ_a with respect to property Π (ideally, $\Phi \in \Pi$ iff $\Phi_a \in \Pi$; if Φ_a has a ν_1-simple set, it can be split); find the relationship between Φ_a and its split version $\hat{\Phi}_a$ with respect to Π (ideally, $\Phi_a \in \Pi$ iff $\hat{\Phi}_a \in \Pi$); investigate the split chain which has—this is the salient point—a pseudo-atom or renewal point, so that regeneration techniques can be used. This way, the power of renewal theory is made available for Markov chains on a general state space, a power that is fully exploited in order to achieve the strong and new results on the ergodic behavior of Markov chains (Part III).

Another praiseworthy feature of the book is the systematic use of a tool which originally was created for continuous time systems, namely the use of the Lyapunov drift $\Delta V(x) = V(F(x)) - V(x)$ for the deterministic system $\Phi_{n+1} = F(\Phi_n)$ and its stochastic counterpart $\Delta V(x) = E_x V(\Phi_1) - V(x)$, together with a discrete-time version of Dynkin's formula. Use of the drift criterion that ΔV is negative outside a petite set (replacing compact sets) has a two-fold advantage: on the theoretical level it elucidates analogies to both deterministically degenerate and continuous-time Markov processes, and on the practical level it helps to verify the property in question.

Content of the three parts. In Part I, "Communication and Regeneration" (Chapters 1–7), after having made the reader familiar with a wide range of Markov models that are used in practice, the authors develop their methodological approach and their tools described above while analyzing the general structure of ψ-irreducible chains and applying the results, as found so far, to the practical models.

Part II, "Stability Structures" (Chapters 8–12), discusses recurrence properties. The ψ-irreducibility of a chain is strengthened to recurrence and then to Harris recurrence in a standard fashion, while transience and positivity (positive recurrence) are introduced in a nonstandard way, allowing a very systematic and transparent discussion.

Transient chains are defined by stressing the uniformity of that property. (To discover and expose uniformity structures is on the hidden agenda of the discussion in the book.) A set A is called uniformly transient, if uniformly for x in A: $E_x(\eta_A) \leq M < \infty$, η_A the occupation time in A; and a chain Φ is called transient, if the state space can be covered by countably many uniformly transient sets.

The concept of positivity is introduced, for example, not via finite mean return time, but by requiring the existence of an invariant probability measure. The starting point is a subinvariant measure, $\mu(A) \geq \mu(P(\cdot, A))$, which always exists as a strictly subinvariant measure for transient chains. So the question is to find conditions which ensure the invariance and the finiteness of the subinvariant σ-finite measure.

Part III, "Convergence" (Chapters 13–19), is dedicated to the convergence of the n-step transition probabilities, $P^n(x, A)$, mostly to a given invariant probability measure π. Here the splitting technique and renewal theory show their full potential to generalize the results from countable state spaces to general state spaces and to unify long-standing facts with those found recently (often by the authors themselves).

In its strongest and most succinct form, this convergence appears as $\|P^n - \pi\|_V \leq kr^n$, $0 < k$, $0 < r < 1$, where $\|\cdot\|_V$ is the operator norm of $P^n - \pi$ viewed as an operator with $(P^n - \pi)f(x) = P^n(x, f) - \pi(f) = \int f(y)[P^n(x, dy) - \pi(dy)]$ on the space of functions f with finite V-norm: $|f|_V = \sup_x[|f(x)|/V(k)] < \infty$, $V(x) \geq 1$. The functions space allows for unbounded functions and the operator norm takes into account the strong uniformity of the convergence of $P^n(x, A)$ in both A and x.

For $V = \text{constant} = 1$, one has what is called uniform ergodicity: $\|P^n - \pi\|_V \to 0$. One of the highlights of the book is a cycle of equivalent statements, which says that the fact of converging to zero is equivalent to converging with exponential speed, which in turn is equivalent to Doeblin's condition, as well as to conditions in terms of petite sets and the drift.

Since the Markov chain properties discussed in the book are systematically described with the help of small sets, petite sets and the drift, seemingly incomparable properties are made comparable (e.g., in terms of their behavior with respect to petite sets) and turn out to constitute a spectrum of stability properties. The reader who was not convinced of the value of this book by its title may be convinced here.

The third part is completed by an instructive discussion about the limiting behavior of the Markov paths themselves (rather than their transition probabilities) in terms of the classical triple LLN, CLT and LIL, a characterization of positive (versus null) recurrence, and generalizations of the drift condition that are easier to verify or to construct than the 1-step drift in practical situations of applying the theory.

In this way, the book ends as it starts, with a view to applications: applications that both illustrate and motivate the theory throughout this very readable, comprehensive and illuminating textbook.

DEPARTMENT OF MATHEMATICS
UNIVERSITY OF NORTH CAROLINA
CHARLOTTE, NORTH CAROLINA 28223