Queueing Syst (2011) 68:305–311
DOI 10.1007/s11134-011-9238-x

Wide-sense regeneration for Harris recurrent Markov processes: an open problem

Peter W. Glynn

Received: 10 May 2011 / Revised: 10 May 2011 / Published online: 6 July 2011
© Springer Science+Business Media, LLC 2011

Abstract Harris recurrence is a widely used tool in the analysis of queueing systems. For discrete-time Harris chains, such systems automatically exhibit wide-sense regenerative structure, so that renewal theory can be applied to questions related to convergence of the transition probabilities to the equilibrium distribution. By contrast, in continuous time, the question of whether all Harris recurrent Markov processes are automatically wide-sense regenerative is an open problem. This paper reviews the key structural results related to regeneration for discrete-time chains and continuous time Markov processes, and describes the key remaining open problem in this subject area.

Keywords Harris recurrence · Markov chains · Markov processes · Regeneration · Renewal theory

Mathematics Subject Classification (2000) 60J05 · 60J25 · 60K05

1 Introduction

In constructing a queueing model, the first order of business is typically to establish the conditions under which the system is stable. The stability (or instability) of the model has both important performance engineering implications and fundamental mathematical consequences. When the model inputs are stationary sequences, stability is often established using tools based on ergodic theory and monotonicity; see, for example, [11]. A much richer class of tools exists when the model inputs form independent and identically distributed (iid) sequences, in which case the system can

P.W. Glynn (✉)
Department of Management Science and Engineering, Stanford University, Stanford, CA 94305, USA
e-mail: glynn@stanford.edu

typically be studied as a Markov chain or Markov process. Particularly nice mathematical properties ensue when the chain or process can be established to be Harris recurrent. For example, Harris recurrence is known to imply existence of invariant measures, as well as laws of large numbers and central limit theorems.

Harris recurrence underlies much of the extensive stability theory for queues that have been developed over the last twenty years. For example, the very useful connections that have been developed between fluid model stability and queueing network stability rest on a framework in which (deterministic) fluid stability is shown to imply (positive) Harris recurrence of the associated (stochastic) queueing model; see [4]. Harris recurrence has also been exploited in the study of stability of Brownian queueing models, as investigated (for example) in [9]. Finally, it should be noted that a middle ground between the tractability of assuming iid model inputs and the generality of stationary process inputs arises when the inputs are assumed themselves to be Harris recurrent Markov processes. Providing conditions under which the queue is then Harris recurrent underlies the approach developed in [15].

Harris recurrence also can be fruitfully exploited in the development of numerical algorithms for studying queues, specifically Monte Carlo simulation-based methods. In particular, Harris recurrence implies that the process exhibits one-dependent regenerative structure, which can then be utilized to (significantly) simplify the construction of confidence intervals for equilibrium expectations; see [6] for details.

In view of the important connections between queueing theory and Harris recurrence, we discuss in this paper a key mathematical question that remains open in our understanding of Harris recurrence in continuous time: When is the theory of renewal equations applicable to the study of such processes? This essentially comes down to determining whether all such processes are wide-sense regenerative. In the remainder of this paper, we review in Sect. 2 the key results in the theory of Harris recurrence for discrete-time Markov chains. Section 3 then describes the corresponding theory in continuous time, and presents our open problem.

2 Harris recurrent Markov chains

Let $X = (X_n : n \geq 0)$ be a (time-homogeneous) Markov chain taking values in a separable metric space S, endowed with Borel σ-algebra $\mathscr{S}$. We denote the associated probability and expectation operator under which $X_0 = x \in S$ by $\mathrm{P}_x(\cdot)$ and $\mathrm{E}_x(\cdot)$, respectively. Such a Markov chain X is said to be *Harris recurrent* if there exists a non-trivial σ-finite measure η on $(S, \mathscr{S})$ for which $\eta(A) > 0$ implies that

$$\sum_{i=1}^{\infty} I(X_n \in A) = \infty \quad \mathrm{P}_x \text{ a.s.} \tag{2.1}$$

for each $x \in S$. A clearly equivalent characterization of Harris recurrence is to require that whenever $\eta(A) > 0$,

$$\tau_A < \infty \quad \mathrm{P}_x \text{ a.s.} \tag{2.2}$$

for each $x \in S$, where $\tau_A = \inf\{n \geq 0 : X_n \in A\}$ is the first hitting time of A.

However, in the presence of the separability of S (and the consequent countable generation of $\mathscr{S}$), an additional interesting and non-trivial re-formulation of Harris recurrence is possible. In particular, Harris recurrence is then equivalent to existence of an $\mathscr{S}$-measurable subset $K \subseteq S$ (called a *small set* in [12]) for which there exists $\lambda > 0$, a probability φ on S, and $m \geq 1$ such that:

(i) $\tau_K < \infty$ P_x a.s. for each $x \in S$.
(ii) $\mathrm{P}_x(X_m \in B) \geq \lambda\varphi(B)$ for $x \in K$ and $B \in \mathscr{S}$.

In the presence of the above re-formulation, [13] and [2] recognized that condition (ii) permits one to write the m-step transition probabilities in the form

$$\mathrm{P}_x(X_m \in \cdot) = \lambda\varphi(\cdot) + (1-\lambda)Q(x,\cdot) \tag{2.3}$$

for $x \in K$, where λ necessarily must lie in $(0, 1]$ and $(Q(x, B) : B \in \mathscr{S})$ is a probability on S for each $x \in K$. In view of (2.3), one can identify regenerative structure for X as follows:

Suppose that X visits K at time τ. Flip a λ-coin. If the λ-coin comes up "heads", distribute $X_{\tau+m}$ according to φ; if the λ-coin comes up "tails", distribute $X_{\tau+m}$ according to $Q(X_\tau, \cdot)$. Then, generate $(X_{\tau+1}, \ldots, X_{\tau+m-1})$, conditional on $(X_\tau, X_{\tau+m})$.

Conditional on the coin coming up "heads", $X_{\tau+m}$ has a distribution φ that is independent of the position X_τ m time units earlier. Thus, the time $T = \tau + m$ is a randomized stopping time that exhibits regenerative structure. Of course, condition (i) ensures that infinitely many such regeneration times $T_0, T_1, T_2, \ldots$ occur at which X_{T_i} has distribution φ. The regeneration times $(T_n : n \geq 0)$ constructed in this way have the following properties:

Property 1 Set $T_{-1} = 0$ and let $W_j = ((X_k, T_j - T_{j-1}) : T_{j-1} \leq k < T_j)$ be the jth cycle induced by $(T_n : n \geq 0)$. Then, the sequence $(W_j : j \geq 0)$ is a one-dependent sequence of random elements (in the sense that $(W_k : k \leq j)$ is independent of $(W_k : k \geq j + 2)$). Furthermore, $(W_k : k \geq 1)$ is a sequence of identically distributed random elements. In view of this cycle structure, X is said to possess *one-dependent regenerative cycles*.

Property 2 For each $n \geq 0$, $(X_{T_n+k} : k \geq 0)$ is independent of T_n (and, of course, $(X_{T_n+k} : k \geq 0)$ is identically distributed in n as a consequence of Property 1). This means, by definition, that X is a *wide-sense regenerative process*.

When $m = 1$, the cycles $(W_j : j \geq 0)$ are actually independent, so that X is then a (classically) *regenerative process*. (In fact, [14] proves that cycle independence implies that condition (ii) then must hold with $m = 1$.) The existence of one-dependent cycle structure then permits one to easily establish laws of large numbers, central limit theorems, laws of the iterated logarithm and other limit results that are valuable both theoretically and computationally (in simulation of such processes) under the weakest possible conditions; see [7] and [8] for a discussion of necessary and sufficient conditions for such limit theorems when X is classically regenerative.

On the other hand, in the presence of wide-sense regeneration, it follows that if $\mathrm{P}(X_0 \in \cdot) = \varphi(\cdot)$ (so that $T_0 = 0$ and X is a non-delayed regenerative process), then

$$\mathrm{E}f(X_n) = \mathrm{E}f(X_n)I(T_1 > n) + \sum_{j=1}^{n} \mathrm{E}f(X_{n-j})\mathrm{P}(T_1 = j)$$

for (measurable) $f : S \to \mathbb{R}_+$. This permits one to apply the full arsenal of renewal theory to the study of the Markov chain X. In particular, renewal-theoretic methods offer a powerful device for obtaining conditions under which X_n converges to a limit X_∞ in total variation distance, as well as in computing associated rates of convergence.

3 Harris recurrent Markov processes

Let $X = (X(t) : t \geq 0)$ be a (time-homogeneous) strong Markov process taking values in a separable metric space S, endowed with Borel σ-algebra $\mathscr{S}$, and possessing sample paths that are right continuous with left limits. Analogously to discrete time, X is said to be *Harris recurrent* (in continuous time) if there exists a non-trivial σ-finite measure η on $(S, \mathscr{S})$ for which $\eta(A) > 0$ implies that

$$\int_0^\infty I(X(t) \in A)\,\mathrm{d}t = \infty \quad \mathrm{P}_x \text{ a.s.} \tag{3.1}$$

for each $x \in S$; see [3] for this definition. As shown in [10], (3.1) is equivalent to existence of a non-trivial σ-finite measure ν on $(S, \mathscr{S})$ for which $\nu(A) > 0$ implies that

$$\tau_A < \infty \quad \mathrm{P}_x \text{ a.s.} \tag{3.2}$$

for each $x \in S$, where $\tau_A = \inf\{t \geq 0 : X(t) \in A\}$. Note that in continuous time, the measure ν need not coincide with η as must be the case in discrete time; see (2.1) and (2.2). (For example, a one-dimensional recurrent diffusion visits points infinitely often, and yet spends zero time there.)

Note that if $(\Gamma_n : n \geq 0)$ is the sequence of jump times of a unit rate Poisson process independent of X, then $(X(\Gamma_n) : n \geq 0)$ is a time-homogeneous Markov chain taking values in S. Furthermore, it is easily verified that (3.1) implies that

$$\sum_{n=0}^{\infty} I\big(X(\Gamma_n) \in A\big) = \infty \quad \mathrm{P}_x \text{ a.s.}$$

for each $x \in S$, so that $(X(\Gamma_n) : n \geq 0)$ is a Harris recurrent Markov chain (in discrete time). We can then exploit the one-dependent regenerative structure of $(X(\Gamma_n) : n \geq 0)$ to conclude that $(X(t) : t \geq 0)$ contains one-dependent regenerative cycles. In other words, every Harris recurrent Markov process satisfies

Property 1′ Set $T_{-1} = 0$. There exists a sequence of randomized stopping times $(T(n) : n \geq 0)$ for which $(W_j : j \geq 0)$ is a one-dependent sequence of identically distributed random elements, where $W_j = ((X(s), T(j) - T(j-1)) : T(j-1) \leq s < T(j))$.

Turning to the analog of Property 2, suppose that there exists a subset $K \in \mathscr{S}$, $\lambda > 0$, a probability φ on S, and $t > 0$ such that:

(i′) $\tau_K < \infty$ P_x a.s. for each $x \in S$.
(ii′) $\mathrm{P}_x(X(t) \in B) \geq \lambda\varphi(B)$ for $x \in K$ and $B \in \mathscr{S}$.

In the presence of (i′) and (ii′), we can directly exploit the regenerative construction of Sect. 2 to establish that $X(\tau_K + t)$ has distribution φ with probability λ, independent of the position $X(\tau_K)$. Because of condition (i′), infinitely many regeneration times $(T^*(n) : n \geq 0)$ at which $X(T^*(n))$ has distribution φ (independent of $T^*(n)$ and $X(T^*(n) - t)$) can be constructed. (The independence of $T^*(n)$ and $X(T^*(n))$ follows from the fact that $T^*(n)$ is determined by $(X(s) : 0 \leq s \leq T^*(n) - t)$ and the sequence of coin flips used up to time $T^*(n) - t$.) Thus, conditions (i′) and (ii′) guarantee that Property 2′ holds:

Property 2′ For each $n \geq 0$, $(X(T^*(n) + s) : s \geq 0)$ is independent of $T^*(n)$ (and, of course, $(X(T^*(n) + s) : s \geq 0)$ is identically distributed in n), so that X is a wide-sense regenerative process.

Given the presence of wide-sense regeneration, it follows that if $T^*(0) = 0$ (so that X is a non-delayed regenerative process), then

$$\mathrm{E}f\big(X(r)\big) = \mathrm{E}f\big(X(r)\big)I\big(T^*(1) > r\big) + \int_{(0,t]} \mathrm{E}f\big(X(r-u)\big)\mathrm{P}\big(T^*(1) \in \mathrm{d}u\big)$$

for (measurable) $f : S \to \mathbb{R}_+$, so that (once again) the full body of renewal theory can be applied to the analysis of X. In view of this, (i′) and (ii′) are sometimes taken as the natural starting point for a theory of recurrence in continuous time (rather than (3.1)); see, for example, [1], p. 198. This raises the following question.

Open Problem: Does every Harris recurrent Markov process (satisfying (3.1)) necessarily satisfy (i′) and (ii′) (or, more generally, exhibit wide-sense regeneration)?

In addition to its natural interest as a key question in the recurrence theory for Markov processes, it should be noted that in some applications, verification of (3.1) or (3.2) is significantly easier than direct verification of (i′) and (ii′). This occurs because the natural mechanism for specifying a Markov process in continuous time is through its infinitesimal generator (see, for example, [5]), in which case the transition probabilities $(\mathrm{P}_x(X(t) \in B) : x \in K, B \in \mathscr{S})$ will typically not be known explicitly.

To understand further the subtleties involved in proving that (i′) and (ii′) hold for a process satisfying (3.1), it should be observed that the key to verifying (i) and (ii) is the use of a measure-differentiation argument that can be found, for example, on pp. 103–105 of [12]. The argument proves that when (2.1) holds, (ii) can be verified for a probability φ that is absolutely continuous with respect to η. However, this can fail in continuous time.

Example 3.1 Let $X = (X(t) : t \geq 0)$ be the residual lifetime process associated with a renewal process in which the inter-renewal time distribution F has countable support on $(0, \infty)$ and is non-arithmetic. In this case, the stationary distribution π for X satisfies

$$\pi(\mathrm{d}x) = \frac{(1 - F(x))\,\mathrm{d}x}{\int_{[0,\infty)} t F(\mathrm{d}t)}$$

for $x \geq 0$, and (3.1) holds for $\eta(\mathrm{d}x) = I(x \geq 0)\,\mathrm{d}x$. Because F is supported on countably many points, so is its n-fold convolution for each $n \geq 1$. It follows that $\mathrm{P}_x(X(t) \in \cdot)$ is singular with respect to Lebesgue measure for any $x \in \mathbb{R}_+$ and $t > 0$. Consequently, (ii′) cannot hold for a probability φ that is absolutely continuous with respect to η. Note, however, that this process is both classically regenerative and wide-sense regenerative.

This example makes clear that any use of a measure-differentiation argument to verify (ii′) must have a different flavor from that used in discrete time. Example 3.1 is, in some sense, a canonical example of the difficulties that can arise. In particular, if $(X(t) : t \geq 0)$ is a Harris recurrent Markov process for which $X(t)$ converges in total variation to π as $t \to \infty$ (or, equivalently, that there exists $h > 0$ such that $X(nh)$ converges to π in total variation as $n \to \infty$), then $\mathrm{P}_x(X(t) \in \cdot)$ has a non-trivial component for t sufficiently large that is absolutely continuous with respect to π, so that one can construct a recurrent subset $K \in \mathscr{S}$ for which (ii′) holds with φ absolutely continuous with respect to π. For the residual lifetime process, spread-outness of F is known to be a necessary and sufficient condition for total variation convergence; see [16].

References

1. Asmussen, S.: Applied Probability and Queues, 2nd edn. Springer, New York (2003)
2. Athreya, K.B., Ney, P.: A new approach to the limit theory of recurrent Markov chains. Trans. Am. Math. Soc. **245**, 493–501 (1978)
3. Azéma, J., Kaplan-Duflo, M., Revuz, D.: Mesure invariante sur les classes récurrentes des processus de Markov. Z. Wahrscheinlichkeitstheor. Verw. Geb. **8**, 157–181 (1967)
4. Dai, J.G.: On positive Harris recurrence of multiclass queueing networks: A unified approach via fluid limits. Ann. Appl. Probab. **5**, 49–77 (1995)
5. Ethier, S.N., Kurtz, T.G.: Markov Processes: Characterization and Convergence. Wiley, New York (1986)
6. Glynn, P.W.: Some topics in regenerative steady-state simulation. Acta Appl. Math. **34**, 225–236 (1994)
7. Glynn, P.W., Whitt, W.: Limit theorems for cumulative processes. Stoch. Process. Appl. **47**, 299–314 (1993)
8. Glynn, P.W., Whitt, W.: Necessary conditions in limit theorems for cumulative processes. Stoch. Process. Appl. **98**, 199–209 (2002)
9. Harrison, J.M., Williams, R.J.: Brownian models of open queueing networks with homogeneous customer populations. Stochastics **22**, 77–115 (1987)
10. Kaspi, H., Mandelbaum, A.: On Harris recurrence in continuous time. Math. Oper. Res. **19**, 211–222 (1994)
11. Loynes, R.M.: The stability of a queue with non-independent inter-arrival and service times. Proc. Camb. Philos. Soc. **58**, 497–520 (1962)

12. Meyn, S.P., Tweedie, R.L.: Markov Chains and Stochastic Stability, 2nd edn. Cambridge University Press, Cambridge (2009)
13. Nummelin, E.: A splitting technique for Harris recurrent chains. Z. Wahrscheinlichkeitstheor. Verw. Geb. **43**, 309–318 (1978)
14. Nummelin, E.: General Irreducible Markov Chains and Nonnegative Operators. Cambridge University Press, Cambridge (1984)
15. Sigman, K.: One-dependent regenerative processes and queues in continuous time. Math. Oper. Res. **15**, 175–189 (1990)
16. Thorisson, H.: Coupling, Stationarity, and Regeneration. Springer, New York (2000)