9 THE POISSON EQUATION FOR COUNTABLE MARKOV CHAINS: PROBABILISTIC METHODS AND INTERPRETATIONS

Armand M. Makowski
Adam Shwartz

Abstract: This paper considers the Poisson equation associated with time-homogeneous Markov chains on a countable state space. The discussion emphasizes probabilistic arguments and focuses on three separate issues, namely (i) the existence and uniqueness of solutions to the Poisson equation, (ii) growth estimates and bounds on these solutions and (iii) their parametric dependence. Answers to these questions are obtained under a variety of recurrence conditions.

Motivating applications can be found in the theory of Markov decision processes in both its adaptive and non-adaptive formulations, and in the theory of Stochastic Approximations. The results complement available results from Potential Theory for Markov chains, and are therefore of independent interest.

9.1 INTRODUCTION

Let $P \equiv (p_{xy})$ be the one-step transition matrix for a time-homogeneous Markov chain $\{X_t,\ t = 0, 1, \dots\}$ taking values in some countable space $\mathbb{X}$. This paper is devoted to the corresponding *Poisson* equation with forcing function $r : \mathbb{X} \to \mathbb{R}$, namely

$$h(x) + w = r(x) + \sum\nolimits_y p_{xy} h(y), \quad x \in \mathbb{X} \tag{9.1}$$

for scalar w and mapping $h : \mathbb{X} \to \mathbb{R}$. This equation arises naturally in a variety of problems associated with Markov chains as the following examples indicate.

1. As shown in Section 9.3, solving the Poisson equation provides a means to evaluate the long-run average cost w associated with the cost function r [36]: If (9.1) has a solution (h, w) and some mild growth conditions are satisfied, then Lemma 9.2 states that

$$w = \lim_t \mathbb{E}_\mu \left[\frac{1}{t+1} \sum_{s=0}^{t} r(X_s) \right] \tag{9.2}$$

where μ is the initial distribution and $\mathbb{E}_\mu$ is the corresponding expectation operator. The function h measures the sensitivity of the cost to the initial state, and represents a second-order effect captured through the "deviation matrix" [13]. This function h can also serve as a "Lyapunov function" in establishing ergodicity [26], and plays a key role in proving the convergence of the policy improvement algorithm [26]. Approximate solutions can be used for simulations—see Chapter 10.

2. During the last decade there has been widespread interest in stochastic approximation algorithms as a means to solve increasingly complex engineering problems [1, 5, 16, 17]. As a result, focus has shifted from the original Robbins-Monro algorithm to (projected) stochastic approximations driven by Markovian "noise" or "state" processes. Properties of solutions to an appropriate Poisson equation play an essential role when establishing the a.s. convergence of such adaptive algorithms [1, 18, 22, 24, 25, 39].

3. In the context of Markov decision processes (MDPs), the need for adaptive policies can arise in response to both modeling uncertainties and computational limitations [40]. Several adaptive policies have been proposed as "implementations" to a Markov stationary policy, and shown to yield the same cost performance [3, 18, 19, 23, 40]. Here too, the analysis requires precise information on the solution to the Poisson equation associated with the non-adaptive policy [40].

In many of these applications, it is natural to view the forcing function r and the transition matrix P as parameterized, say by some parameter θ (which may be loosely interpreted as a control variable). The requisite analysis then typically exploits smoothness properties (in θ) of the solution h together with various growth estimates (in x) for h. In addition, estimates on the moments of $\{h(X_t),\ t = 0, 1, 2, \dots\}$ are required, with the added difficulty that the resulting process $\{X_t,\ t = 0, 1, 2, \dots\}$ is not necessarily Markovian (say, under the given stochastic approximation scheme or adaptive policy).

Our main objective is to develop methods for addressing the concerns above in a systematic fashion. Whenever possible, we emphasize a probabilistic viewpoint as we focus mostly on the following three issues:

1. Existence and uniqueness of solutions to the Poisson equation (9.1);
2. Growth estimates and bounds on these solutions; and
3. Conditions for smoothness in the parameter of these solutions when dealing with the parametric case, as would arise when establishing the a.s.

convergence of stochastic approximations and the self-tuning property of adaptive policies.

Answers to these questions are given under a variety of recurrence conditions. As we try to keep the exposition relatively self-contained, we have included some standard material on the Poisson equation. In addition to its tutorial merit, the discussion given here provides a unified treatment to many of the issues associated with the Poisson equation, e.g. existence, uniqueness and representation of solutions. This is achieved by manipulating a single *martingale* naturally induced by the Poisson equation.

Questions of existence and uniqueness of solutions to (9.1) have obvious and natural points of contact with the Potential Theory for Markov chains [15, 29]. Unfortunately many situations of interest in applications, say in the context of MDPs, are not readily covered by classical Potential Theory. Indeed, the classical theory treats the purely transient and recurrent cases separately, with drastically different results for each situation. This approach is thus of limited use in the above-mentioned situations, where the recurrence structure of the Markov chain is typically far more complex in that it combines both transient and recurrent states. Here, in contrast with the analytical approach of classical Potential Theory, emphasis has been put on giving an explicit representation of the solution to (9.1) with a clear probabilistic interpretation.

This probabilistic approach allows for a relatively elementary treatment of questions of existence and uniqueness, under a rather general recurrence structure. We accomplish this by focusing on the discrete space case, and by keeping the assumptions as transparent as possible. The intuition developed here applies to the general state-space case, under mild conditions on the existence of petite sets—see Chapter 10 and [10, 26, 27]. Results are obtained in various degrees of completeness for both finite and countably infinite state spaces; recurrence structures include multiple positive recurrent classes, and transient classes. A representation for h is derived in detail in the case of a single positive recurrent class under integrability conditions involving the forcing function r. The derivation uses elementary methods, and provides intuition into more general situations. This representation is shown to also hold in countable case with multiple classes, and readily lends itself to establishing natural bounds on the growth rate of h (as a function of the state), and to investigating smoothness properties in the parameterized problem.

Similar results are given in [10] for the ergodic case on general state spaces. In addition, when the forcing function r is positive and "increasing" (i.e. when its sub-level sets are compact), there is an elegant theory that relates geometric ergodicity to the Poisson equation; details and references can be found in Chapter 10. As evidenced by the references section, there is a very large literature on the Poisson equation; of particular note is the monograph by Nummelin [28]. In the context of MDPs, bounded solutions are discussed by Ross [34, 35], Gubenko and Shtatland [11] and Yushkevich [42]. One of the first treatments of unbounded solutions is available in Robinson [31, 32] (with details in Chapter 5).

The paper is organized as follows: The set-up is given in Section 9.2 together with the basic martingale associated with (9.1). Various uniqueness results on the the solution (w, h) are discussed in Section 9.3. We give two decomposition results in Section 9.4; the first is based on the decomposition of the state space $\mathbb{X}$ into its recurrent and transient classes, while the second is an analog of the standard Green decomposition and relies on an expansion of the forcing function in terms of more "elementary" functions. To set the stage for the countably infinite case, we briefly recall an algebraic treatment of the finite-state case in Section 9.5. In Section 9.6, under a single positive recurrent class assumption, an explicit representation for the solution is developed in terms of the recurrence time to some distinguished state. An example is developed in Section 9.7 to illustrate the material of previous sections. Bounds and extensions to unbounded forcing functions and multiple recurrent classes are given in Section 9.8. Equipped with this probabilistic representation of solutions, we can now investigate the smoothness properties of solutions to the parameterized problem; methods for proving continuity and Lipschitz continuity are developed in Sections 9.9 and 9.10, respectively.

To close, we note that most of the ideas which are discussed here in the context of countable Markov chains have extensions to fairly general state spaces. This is achieved by means of the so-called *splitting technique* [10, 26, 27, 28] which in essence guarantees the existence of an atom on an enlarged state space; details can be found in Chapter 10.

9.2 THE POISSON EQUATION AND ITS ASSOCIATED MARTINGALE

First, a few words on the notation used throughout the paper: The set of all real numbers is denoted by $\mathbb{R}$ and $\mathbf{1}[A]$ stands for the indicator function of a set A. Unless otherwise stated, $\lim_t$, $\underline{\lim}_t$ and $\overline{\lim}_t$ are taken with t going to infinity. Moreover, the infimum over an empty set is taken to be ∞ by convention. The Kronecker mapping $\delta : \mathbb{X} \times \mathbb{X} \to \mathbb{R}$ is defined by $\delta(x, y) = 1$ if $x = y$, and $\delta(x, y) = 0$ otherwise. Finally, the notation $\sum_{x \in \mathbb{X}}$ is often abbreviated as $\sum_x$.

9.2.1 *The set-up*

The notion of a Markov chain we adopt in this paper is more general than the elementary one used in most applications. We do so with the view of broadening the applicability of the material developed here, especially to problems of adaptive control for Markov chains [18, 19, 22, 23, 39, 40].

The state space is a countable, and we assume the existence of a measurable space $(\Omega, \mathcal{F})$ large enough to carry all the probabilistic elements considered in this paper. In particular, let $\{\mathcal{F}_t,\ t = 0, 1, \ldots\}$ denote a filtration of $\mathcal{F}$, i.e. a monotone increasing sequence of σ-fields contained in $\mathcal{F}$ such that $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$ for all $t = 0, 1, \ldots$, and let $\{X_t,\ t = 0, 1, \ldots\}$ be a sequence of $\mathbb{X}$-valued rvs which are $\mathcal{F}_t$-adapted, i.e. the rv X_t is $\mathcal{F}_t$-measurable for all $t = 0, 1, \ldots$.

The Markovian structure of interest is defined by postulating the existence of a family $\{\mathbb{P}_x,\ x \in \mathbb{X}\}$ of probability measures on $\mathcal{F}$ such that for all x and