VU Research Portal

A paradigm for derivatives of positive systems

Heidergott, B.F.

published in

Positive Systems (Lecture Notes in Control and Information Sciences)
2003

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Heidergott, B. F. (2003). A paradigm for derivatives of positive systems. In L. Benvenuti, A. de Santis, & L. Farina (Eds.), *Positive Systems (Lecture Notes in Control and Information Sciences)* (pp. 312-318). Springer-Verlag.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

A Paradigm for Derivatives of Positive Systems

Bernd Heidergott

Vrije Universiteit Amsterdam, Department of Econometrics and Operations
Research, De Boelelaan 1105, 1081 HV Amsterdam, the Netherlands,
bheidergott@feweb.vu.nl

Abstract. We develop a framework for differentiation of positive operators, such as Markov kernels, through interpreting derivatives of positive operators as differences between positive operators. This new paradigm allows to deal with differentiability issues while retaining the framework of positive systems.

1 Introduction

In this paper we show how the dichotomy between positivity and differentiation can be overcome through a concept of "weak" differentiation. The basic idea will be to write the derivative of a positive operator as re-scaled difference of two positive operators. Our main object of study will be Markov chains. The first part of the paper illustrates our concept of differentiation with finite state Markov chains: "weak differentiability" for finite Markov chains is introduced and it is shown that differentiability of the transition matrix of a finite state Markov chain implies differentiability of its stationary distribution. The proof elaborates on the product rule of differentiation (for real-valued mappings) and is different from the proofs put forward in [8] and [4], respectively, where this result has been shown by using the fact that the stationary distribution of a Markov chain is an invariant distribution of the Markov kernel. In the second part of the paper, a review of the theory of "weak differentiation" for general Markov chains will be given. Eventually, we discuss the situation for general positive operators and identify topics of further research.

2 Finite state Markov chains

Let $\Theta = (a, b) \subset \mathbb{R}$, with $a < b$. Let $X_\theta(n) \in \{1, \dots, N\}$ be a discrete-time Markov chain depending on a control parameter θ with deterministic initial value $X_\theta(0) = x_0 \in \{1, \dots, N\}$. Let P_θ denote the transition probability matrix of $X_\theta(n)$, i.e.:

$$(P_\theta)_{ij} = P(X_\theta(n+1) = j \mid X_\theta(n) = i), \quad i, j \in \{1, \dots, N\},$$

for $n \geq 0$. For example, $X_\theta(n)$ may model the queue length in a M/M/1 queue where θ represents the service rate.

2.1 Differentiability

Assume that the elements of P_θ are differentiable and denote the derivative of P_θ by P'_θ , i.e.,

$$(P'_\theta)_{ij} = \frac{d}{d\theta}(P_\theta)_{ij}, \quad 1 \leq i, j \leq N. \quad (1)$$

Notice that while P_θ acts as a mapping on the set of probability vectors in $\mathbb{R}^n$, the image space of P'_θ contains vectors with negative elements. In other words, for any probability distribution $\mu = (\mu_1, \dots, \mu_N)$ (i.e., $\sum \mu_k = 1$ and $\mu_k \geq 0$), μP_θ is again a probability distribution whereas $\mu P'_\theta$ fails to be one. However, as we will show in the following, $\mu P'_\theta$ can be written as difference between positive vectors. The key observation is that a matrix C_{P_θ} and Markov kernels P_θ^+ and P_θ^- exists such that $P'_\theta = C_{P_\theta}(P_\theta^+ - P_\theta^-)$. Typically, C_{P_θ} turns out to be a diagonal matrix with identical elements on the diagonal, which yields $\mu P'_\theta = C_{P_\theta}(\nu^+ - \nu^-)$ with $\nu^+ = \mu P_\theta^+$ and $\nu^- = \mu P_\theta^-$ probability vectors.

Examining the situation in (1) more closely, one notices that $\sum_j (P_\theta)_{ij} = 1$ implies that $\sum_j (P'_\theta)_{ij} = 0$, for $1 \leq i \leq N$. In words, because P_θ has row sums equal to one (and thus independent of θ), P'_θ has row sum zero, or, equivalently:

$$\sum_j \max((P'_\theta)_{ij}, 0) = \sum_j \max(-(P'_\theta)_{ij}, 0),$$

for any row i . For $1 \leq i \leq N$, let $c_{P_\theta}(i) = \sum_j \max((P'_\theta)_{ij}, 0)$, then the matrices P_θ^+ and P_θ^- defined through

$$(P_\theta^+)_{ij} = \begin{cases} \frac{\max((P'_\theta)_{ij}, 0)}{c_{P_\theta}(i)} & \text{for } c_{P_\theta}(i) > 0 \\ (P_\theta)_{ij} & \text{for } c_{P_\theta}(i) = 0 \end{cases}$$

and

$$(P_\theta^-)_{ij} = \begin{cases} \frac{\max(-(P'_\theta)_{ij}, 0)}{c_{P_\theta}(i)} & \text{for } c_{P_\theta}(i) > 0 \\ (P_\theta)_{ij} & \text{for } c_{P_\theta}(i) = 0 \end{cases}$$

are transition matrices, i.e., their row sum equals one. Moreover, the derivative of P_θ has the following representation

$$P'_\theta = C_{P_\theta}(P_\theta^+ - P_\theta^-), \quad (2)$$

where

$$(C_\theta)_{ij} = \begin{cases} c_{P_\theta}(i) & \text{for } c_{P_\theta}(i) > 0 \text{ and } j = i \\ 1 & \text{for } c_{P_\theta}(i) = 0 \text{ and } j = i \\ 0 & \text{otherwise.} \end{cases}$$

The representation of P'_θ in (2) allows to interpret the derivative of the transition matrix as re-scaled difference of two transition matrices.

Remark 1. Let P, Q denote transition matrices on a common state space. In the theory of singularly perturbed Markov chains, the situation is studied when $P_\theta = \theta(Q - P) + P$, for $\theta \in [0, 1]$, see Chapter 4 in [2] for details on singularly perturbed Markov chains. Hence, $Q - P$ is the derivative of P_θ with respect to θ and formulae for singularly perturbed Markov chains can be interpreted as particular derivative expressions. For an interpretation of the above model in terms of infinitesimal perturbation analysis we refer to [3].

Example 1. Let $X_\theta(n)$ be the discrete-time queue length process of an M/M/1/N queue with arrival rate λ and service rate θ , with $\theta > \lambda > 0$. The transition matrix is then given in matrix form by

$$P_\theta = \begin{bmatrix} 0 & 1 & 0 & & & \\ \frac{\theta}{\lambda+\theta} & 0 & \frac{\lambda}{\lambda+\theta} & 0 & & \\ 0 & \frac{\theta}{\lambda+\theta} & 0 & \frac{\lambda}{\lambda+\theta} & 0 & \dots \\ & & & \ddots & & \\ & & & & 1 & 0 \end{bmatrix}$$

The matrix P_θ is differentiable with respect to θ with derivatives

$$\frac{d}{d\theta}(P_\theta)_{ij} = \begin{cases} \frac{-\lambda}{(\lambda+\theta)^2} & \text{for } 2 \leq i \leq N-1, j = i+1 \\ \frac{\lambda}{(\lambda+\theta)^2} & \text{for } 2 \leq i \leq N-1, j = i-1 \\ 0 & \text{otherwise.} \end{cases}$$

Let $C_{\lambda,\theta}$ be a matrix with diagonal elements $\frac{1}{(\lambda+\theta)^2}$ and zero elements elsewhere, then

$$\frac{d}{d\theta}P_\theta = C_{\lambda,\theta}(P^+ - P^-), \quad (3)$$

with

$$P^+ = \begin{bmatrix} 0 & 1 & & & \\ 1 & 0 & & & \\ & & \ddots & & \\ & & & 1 & 0 \\ & & & & 1 & 0 \end{bmatrix} \quad P^- = \begin{bmatrix} 0 & 1 & & & \\ & 0 & 1 & & \\ & & & \ddots & \\ & & & & 0 & 1 \\ & & & & 1 & 0 \end{bmatrix},$$

where the definitions of the elements of the first and the last rows of $C_{\lambda,\theta}$, P^+ and P^- have been chosen in order to obtain a simple representation. Notice that P^+ and P^- are transition matrices. The triple $(C_{\lambda,\theta}, P^+, P^-)$ may serve as matrix-valued representation of P'_θ .

For $i, j \in \{1, \dots, N\}$, the probabilities $P(X_\theta(n+1) = j | X_\theta(1) = i)$ are given through the elements of n th power of P_θ , denoted by $P_\theta^n := (P_\theta)^n$, where P_θ^0 is the identity matrix, i.e., $P(X_\theta(n+1) = j | X_\theta(1) = i) = (P_\theta^n)_{ij}$. We have assumed that P_θ is differentiable. Therefore, P_θ^n is differentiable as well. Specifically,

$$\frac{d}{d\theta} P_\theta^n = \sum_{j=0}^{n-1} P_\theta^j P_\theta' P_\theta^{n-j-1}. \quad (4)$$

Example 2. We revisit the M/M/1/N queue as introduced in Example 1. Inserting (3) in (4) and noticing that $C_{\lambda, \theta}$ is a matrix that only has elements on its diagonal and these elements are identical, yields

$$\frac{d}{d\theta} P_\theta^n = C_{\lambda, \theta} \left(\sum_{j=0}^{n-1} P_\theta^j P^+ P_\theta^{n-j-1} - \sum_{j=0}^{n-1} P_\theta^j P^- P_\theta^{n-j-1} \right).$$

In words, the derivative of the n th power of a differentiable transition matrix admits a representation like (2) as well.

2.2 Differentiating a stationary distribution

In this section we show that, under some mild additional conditions, differentiability of P_θ implies differentiability of the unique invariant distribution of P_θ (existence is assumed here), denoted by π_θ , and that the derivative of π_θ can be obtained as difference between appropriate Markov chains. We denote by

$$\Pi_\theta = \lim_{N \rightarrow \infty} \frac{1}{N+1} \sum_{n=0}^N P_\theta^n$$

the ergodic projector associated to P_θ . Specifically, Π_θ is a matrix with rows equal to π_θ and it holds that $\pi_\theta = \mu \Pi_\theta$, for any initial distribution μ . Assume that μ is independent of θ . Hence,

$$\frac{d}{d\theta} \pi_\theta = \mu \frac{d}{d\theta} \Pi_\theta$$

In the following we calculate

$$\frac{d}{d\theta} \left(\lim_{N \rightarrow \infty} \frac{1}{N+1} \sum_{n=0}^N P_\theta^n \right).$$

The key conditions for our analysis is the following.

The space of transition probabilities on $\{1, \dots, N\}$ can be equipped with a norm, denoted by $\|\cdot\|$, such that for an open neighborhood $\Theta_0 \subset \Theta$ of θ it holds that:

- (C1) $\|\pi_\theta\|$ is finite on Θ_0 (local stability),
 (C2) finite constants $c_{\theta'}$ and $\rho_{\theta'}$, with $\sup_{\theta' \in \Theta_0} c_{\theta'} < \infty$ and $\sup_{\theta' \in \Theta_0} \rho_{\theta'} \triangleq \rho < 1$, exist such that

$$\forall \theta' \in \Theta_0 : \quad \|P_{\theta'}^m - \Pi_{\theta'}\| \leq c_{\theta'} \rho_{\theta'}^m,$$

(local geometric ergodicity at uniform rate)

- (C3) P'_θ is Lipschitz continuous at θ , i.e.,

$$\forall \theta' \in \Theta_0 : \quad \|P'_\theta - (P_{\theta'})'\| < |\theta - \theta'| K',$$

for some finite number K , and $\|P'_\theta\|$ is finite on Θ_0 .

A typical choice for $\|\cdot\|$ is the supremum norm on $\mathbb{R}^n$, which implies that $\|P_\theta^n\| \leq 1$, for any n , and $\|\Pi_\theta\| \leq 1$ provided that π_θ exists. By (4),

$$\lim_{N \rightarrow \infty} \frac{d}{d\theta} \left(\frac{1}{N+1} \sum_{n=0}^N P_\theta^n \right) = \lim_{N \rightarrow \infty} \frac{1}{N+1} \sum_{n=1}^N \sum_{j=0}^{n-1} P_\theta^j P'_\theta P_\theta^{n-j-1}$$

and the fact that $P_{\theta'}'$ has row sum zero implies

$$\frac{1}{N+1} \sum_{n=1}^N \sum_{j=0}^{n-1} P_\theta^j P'_\theta P_\theta^{n-j-1} = \frac{1}{N+1} \sum_{n=1}^N \sum_{j=0}^{n-1} P_\theta^j P'_\theta (P_\theta^{n-j-1} - \Pi_\theta). \quad (5)$$

By conditions (C1) – (C3), for any N , the supremum norm of the expression on the right-hand side of the above equation is bounded by $c \|P'_\theta\| \frac{1}{1-\rho}$, which is finite. Hence, the limit exists and we compute

$$\begin{aligned} & \lim_{N \rightarrow \infty} \frac{1}{N+1} \sum_{n=1}^N \sum_{j=0}^{n-1} P_\theta^j P'_\theta (P_\theta^{n-j-1} - \Pi_\theta) \\ &= \lim_{N \rightarrow \infty} \frac{1}{N+1} \sum_{n=0}^N P_\theta^n \sum_{j=0}^{\infty} P'_\theta (P_\theta^j - \Pi_\theta) \\ &= \Pi_\theta P'_\theta \sum_{j=0}^{\infty} (P_\theta^j - \Pi_\theta) = \Pi_\theta P'_\theta D_\theta, \end{aligned}$$

with

$$D_\theta = \sum_{j=0}^{\infty} (P_\theta^j - \Pi_\theta),$$

where D_θ is known as *deviation matrix* in the literature, see for example [7]. We have thus shown that

$$\lim_{N \rightarrow \infty} \frac{d}{d\theta} \left(\frac{1}{N+1} \sum_{n=0}^N P_\theta^n \right) = \Pi_\theta \sum_{j=0}^{\infty} P'_\theta D_\theta$$

and elaborating on (C1) - (C3) it follows that

$$\Pi'_\theta = \lim_{N \rightarrow \infty} \frac{d}{d\theta} \left(\frac{1}{N+1} \sum_{n=0}^N P_\theta^n \right). \quad (6)$$

For a proof use the fact that we may choose Θ_0 small enough such that $\sup_{\theta' \in \Theta} \|P_{\theta'}\| =: K < \infty$. Hence, the expression in equation (5) is uniformly bounded on Θ_0 in N . By Lipschitz continuity of P'_θ and P_θ (which follows from $K < \infty$), we obtain that the expression in equation (5) is uniformly continuous as well. The theorem of Arzela-Ascoli applies and the right-hand side of (6) converges uniformly, which implies that interchanging the order of differentiation and limit is justified. Hence,

$$\frac{d}{d\theta} \pi_\theta = \mu \Pi_\theta P_\theta' D_\theta = \pi_\theta \Pi_\theta' D_\theta, \quad (7)$$

or, equivalently,

$$\frac{d}{d\theta} \pi_\theta = \pi_\theta C_{P_\theta} P_\theta^+ D_\theta - \pi_\theta C_{P_\theta} P_\theta^- D_\theta.$$

In words, the derivative of the stationary distribution (the fix-point of the positive operator P_θ) can be represented as the difference of two well-defined positive systems. Specifically, the above result recovers the result in [9] for the case of finite state space. The above formula can be translated in various ways into unbiased gradient estimators for the stationary performance, see [6, 4] for details.

Example 3. We revisit the M/M/1 example. If the system is stable on Θ_0 with Θ_0 an open neighborhood of θ , then $\pi'_\theta = \pi_\theta C_{\lambda, \theta} P^+ D_\theta - \pi_\theta C_{\lambda, \theta} P^- D_\theta$.

3 General state-space Markov chains

In this section we review the theory of differentiation for Markov chains on a general state-space S . Let $(S, \mathcal{T})$ denote a measurable space, i.e., $\mathcal{T}$ is a σ -field over S , and consider a family of Markov kernels $(P_\theta : \theta \in \Theta)$ on $(S, \mathcal{T})$, with $\Theta = (a, b) \subset \mathbb{R}$, for $a < b$. Let $L^1(P_\theta; \Theta) \subset \mathbb{R}^S$ denote the set of measurable mappings $g : S \rightarrow \mathbb{R}$ such that $\int_S P_\theta(s; du) |g(u)|$ is finite for all $\theta \in \Theta$ and $s \in S$.

A first complication arises when one tries to define what "differentiability" of P_θ should mean. The following definition has been fruitful in applications. Let $\mathcal{D} \subset L^1(P_\theta; \Theta)$. We call P_θ $\mathcal{D}$ -differentiable if a transition kernel P'_θ exists such that for any $s \in S$ and any $g \in \mathcal{D}$

$$\frac{d}{d\theta} \int_S P_\theta(s; du) g(u) = \int_S P'_\theta(s; du) g(u). \quad (8)$$

For example, the Markov kernel P_θ in Example 1 is $\mathbb{R}^N$ -differentiable. Let $C_b(S)$ denote the set of continuous bounded mappings from S to $\mathbb{R}$. Then, $C_b(S) \subset \mathcal{D}$ implies that P'_θ in (8) is uniquely defined. Notice that uniqueness of P'_θ comes for free in the finite state-space case.

Any triple $(c_{P_\theta}(\cdot), P_\theta^+, P_\theta^-)$, with $P_\theta^\pm$ Markov kernels and c_{P_θ} a measurable mapping from S to $\mathbb{R}$, that satisfies

$$\int_S P'_\theta(s; du) g(u) = c_{P_\theta}(s) \left(\int_S P_\theta^+(s; du) g(u) - \int_S P_\theta^-(s; du) g(u) \right),$$

for any $g \in \mathcal{D}$, is called a $\mathcal{D}$ -derivative of P_θ . Notice that $\mathcal{D}$ -derivatives are not unique. For example, the Markov kernel in Example 1 has $\mathbb{R}^N$ -derivative (c_{P_θ}, P^+, P^-) with P^+ and P^- as defined in Example 1 and $c_{P_\theta}(s) = \lambda/(\lambda + \theta)^2$.

Does $\mathcal{D}$ -differentiability of P_θ already imply the existence of a $\mathcal{D}$ -derivative of P_θ ? For the *finite state space* case, the answer is "yes" as we have shown in Section 2.1. For a general state-space, however, the situation is more complicated. Provided that $(S, \mathcal{T})$ is such that $\mathcal{T}$ is *countable*, it holds that if $\mathcal{D}$ contains for any $A \in \mathcal{T}$ its indicator function, then $\mathcal{D}$ -differentiability of P_θ implies the existence of a $\mathcal{D}$ -derivative. For *general* $(S, \mathcal{T})$, we have the following result. Denote the *total variation norm* of a transition kernel Q on $(S, \mathcal{T})$ by

$$\|Q\|_{tv} \triangleq \sup_{s \in S} \sup_{\substack{f \in C_b(S) \\ |f| \leq 1}} \int f(z) Q(s; dz). \quad (9)$$

If $C_b(S) \subset \mathcal{D}$ and if $\|P'_\theta\|_{tv} < \infty$, then $\mathcal{D}$ -differentiability of P_θ implies the existence of a $\mathcal{D}$ -derivative, see [5].

The key ingredients for our proof of differentiability of the stationary distribution in Theorem 1 was that (i) a product rule of differentiation holds, and that (ii) there exists a norm, say $\|\cdot\|$, such that $\|P'_\theta\|$ is finite, P'_θ is Lipschitz and P_θ is geometrically ergodic with coefficient ρ_θ , such that $\sup_{\theta' \in \Theta_0} \rho_{\theta'} \triangleq \rho < 1$ for some neighborhood Θ_0 of θ , i.e., conditions (C1) - (C3) hold.

It can be shown that, provided the Markov kernel satisfies a weak Lipschitz condition, the product of $\mathcal{D}$ -differentiable Markov kernels is again $\mathcal{D}$ -differentiable, see [6]. To find good candidates for the norm, we have to resort to stability theory for Markov kernels. A first choice is the total variation norm, see equation (9) for a definition. This is the choice in [8] where it is shown that, under suitable conditions, the stationary distribution is $C_b(S)$ -differentiable. Of course, $C_b(S)$ -differentiability of π is not satisfactory in applications where one is also interested in unbounded performance indicators. Fortunately, the concept of normed ergodicity allows to overcome this restriction. The key idea is to find a Lyapunov function g for P_θ and to consider $v = e^{\lambda g}$, for some positive λ . The norm is then the *weighted supremum norm* with respect to v , in symbols:

$$\|Q\|_v \triangleq \sup_{s \in S} \sup_{\substack{g \\ |g| \leq v}} \frac{\left| \int g(z) Q(s; dz) \right|}{v(s)},$$

with Q a transition kernel on $(S, \mathcal{T})$, see [4] for details. In [4] sufficient conditions are established such that P_θ is ergodic with coefficient ρ_θ , such that $\sup_{\theta' \in \Theta_0} \rho_{\theta'} \triangleq \rho < 1$, and $\|\Pi_{\theta'}\| < \infty$, for $\theta' \in \Theta_0$, where Θ_0 is neighborhood of θ (i.e., condition (C1) and (C2) hold for $\|\cdot\|_v$). Under these conditions, it holds true that if P_θ is $\mathcal{D}_v$ -differentiable with P'_θ Lipschitz continuous at θ and $\|P'_\theta\|_v$ finite, then the stationary distribution is $\mathcal{D}_v$ -differentiable as well and its derivative is given by equation (7), we refer to [4] for details. Facilitating this formula for gradient estimation is discussed in [4].

4 General positive operators

Let λ_θ denote an eigenvalue and x_θ an eigenvector (associated to λ_θ) of the positive operator T_θ , i.e., $\lambda_\theta x_\theta = T_\theta x_\theta$, $\theta \in \Theta$, for some suitable set Θ . For example, T_θ may represent the transition operator in a (max,+)– or (min,+)–linear system and λ_θ the unique eigenvalue (existence is assumed here) and x_θ an eigenvector, see [1] for details. For the analysis in the previous sections, we relied on the fact that, for Markov chains, the maximal positive eigenvalue of T_θ is independent of θ (in fact, $\lambda_\theta = 1$ for $\theta \in \Theta$). For general operators eigenvector(s) as well as eigenvalue(s) will depend on θ . The development of an approach for general positive operators is topic of future research. An application of these results might, for example, lead to a sensitivity analysis of the spectral gap of a Markov chain.

References

1. F. Baccelli, G. Cohen, G. Olsder, and J.-P. Quadrat. *Synchronization and Linearity*. John Wiley and Sons, New-York, 1992.
2. E. Feinberg and A. Schwartz (eds.) *Handbook of Markov decision processes*. Kluwer, Boston, 2002.
3. B. Heidergott and X. Cao. A note on the relation between weak derivatives and perturbation realization. *IEEE Trans. Aut. Control*, 47:1112-1115, 2002.
4. B. Heidergott, A. Hordijk and H. Weisshaupt. Measure-Valued Differentiation for Stationary Markov Chains. *EURANDOM report 2002-027*, 2002.
5. B. Heidergott, A. Hordijk and H. Weisshaupt. A Jordan type decomposition for weak derivatives of Markov kernels. *EURANDOM report 2003-001*, 2003.
6. B. Heidergott and F. Vázquez-Abad. Measure-valued differentiation for stochastic processes: the finite horizon case. *EURANDOM report 2000-033*, 2000.
7. S. Meyn and R. Tweedie. *Markov Chains and Stochastic Stability*. Springer, London, 1993.
8. G. Pflug. *Optimization of Stochastic Models*. Kluwer Academic, Boston, 1996.
9. F. Vázquez-Abad and H. Kushner. Estimation of the derivative of a stationary measure with respect to a control parameter. *J. Appl. Prob.*, 29:343–352, 1992.