VU Research Portal

A Note on Gradient Estimation for Maintenance Systems

Heidergott, B.F.; Yuan, T.

published in

IEEE Transactions on Automatic Control
2007

DOI (link to publisher)

[10.1109/TAC.2007.908333](https://doi.org/10.1109/TAC.2007.908333)

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., & Yuan, T. (2007). A Note on Gradient Estimation for Maintenance Systems. *IEEE Transactions on Automatic Control*, 52, 2141-2145. <https://doi.org/10.1109/TAC.2007.908333>

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

IV. CONCLUSION

In this note, a new method is developed to test stability of piecewise discrete-time linear systems based on a piecewise Lyapunov function. It is shown that the stability can be determined by solving a set of LMIs. The approach can be extended to performance analysis of such systems as in [2] and [3] for their continuous counterparts.

ACKNOWLEDGMENT

The author would like to thank the Associate Editor and the reviewers for a number of constructive comments based on which the note has been greatly improved.

REFERENCES

- [1] J. I. Imura and A. van der Schaft, "Characterization of well-posedness of piecewise-linear systems," *IEEE Trans. Automat. Contr.*, vol. 45, pp. 1600–1619, 2000.
- [2] M. Johansson and A. Rantzer, "Computation of piecewise quadratic Lyapunov functions for hybrid systems," *IEEE Trans. Automat. Contr.*, vol. 43, pp. 555–559, 1998.
- [3] A. Rantzer and M. Johansson, "Piecewise linear quadratic optimal control," *IEEE Trans. Automat. Contr.*, vol. 45, pp. 629–637, 2000.
- [4] A. Hassibi and S. Boyd, "Quadratic stabilization and control of piecewise-linear systems," in *Proc. Amer. Control Conf.*, Philadelphia, PA, 1998, pp. 3659–3664.
- [5] S. Pettersson and B. Lennartson, "LMI approach for stability and robustness of hybrid systems," in *Proc. Amer. Control Conf.*, Albuquerque, NM, 1997, pp. 1714–1718.
- [6] —, "Exponential stability analysis of nonlinear systems using LMIs," in *Proc. 36th IEEE Conf. Decision Control*, San Diego, CA, 1997, pp. 199–204.
- [7] O. Slupphaug and B. A. Foss, "Constrained quadratic stabilization of discrete-time uncertain nonlinear multi-model systems using piecewise affine state feedback," *Int. J. Control*, vol. 72, no. 7, pp. 686–701, 1999.
- [8] F. A. Cuzzola and M. Morari, "A generalized approach for analysis and control of discrete-time piecewise affine and hybrid systems," in *Hybrid Systems: Computation and Control—Lecture Notes in Computer Sciences 2034*. New York: Springer-Verlag, 2001, pp. 189–203.
- [9] D. Mignone, G. Ferrari-Trecate, and M. Morari, "Stability and stabilization of piecewise affine and hybrid systems: An LMI approach," in *Proc. 39th IEEE Conf. Decision Control*, Sydney, Australia, 2000, pp. 504–509.
- [10] A. Bemporad, G. Ferrari-Trecate, and M. Morari, "Observability and controllability of piecewise affine and hybrid systems," *IEEE Trans. Automat. Contr.*, vol. 45, pp. 1864–1876, Oct. 2000.
- [11] N. B. O. L. Pettit and P. E. Wellstead, "Analyzing piecewise linear dynamical systems," *IEEE Control Syst. Mag.*, vol. 15, pp. 43–50, 1995.
- [12] M. Kantner, "Robust stability of piecewise linear discrete-time systems," in *Proc. Amer. Control Conf.*, Albuquerque, NM, 1997, pp. 1241–1245.
- [13] E. D. Sontag, "Nonlinear regulation: The piecewise linear approach," *IEEE Trans. Automat. Contr.*, vol. AC-26, pp. 346–357, 1981.
- [14] J. M. Goncalves, "Constructive Global Analysis of Hybrid Systems," Ph.D. dissertation, Massachusetts Inst. Technol., Cambridge, MA, 2000.
- [15] S. Boyd, L. Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in Systems and Control Theory*. Philadelphia, PA: SIAM, 1994.

A Note on the Relation Between Weak Derivatives and Perturbation Realization

Bernd Heidergott and Xi-Ren Cao

Abstract—This note studies the relationship between two important approaches in perturbation analysis (PA)—perturbation realization (PR) and weak derivatives (WDs). Specifically, we study the relation between PR and WDs for estimating the gradient of stationary performance measures of a finite state-space Markov chain. Will show that the WDs expression for the gradient of a stationary performance measure can be interpreted as the expected PR factor where the expectation is carried out with respect to a distribution that is given through the weak derivative of the transition kernel of the Markov chain. Moreover, we present unbiased gradient estimators.

Index Terms—Markov chains, perturbation analysis (PA), weak derivatives (WDs).

I. INTRODUCTION

Today, *perturbation analysis* (PA) is the most widely accepted gradient estimation technique; see [5]–[7] for details. In this note, we work in particular with the interpretation of PA via *perturbation realization* (PR) factors, see [1]. The aim of our analysis is to establish a connection between PR and the concept of *weak derivatives* (WDs), see [8]. Whereas PA is a sample-path based approach, WDs are a measure theoretic approach to gradient estimation.

WDs translate the analysis of the gradient into a particular splitting of the sample path into two subpaths and observing these subpaths until they couple, that is, until the perturbation dies out. The basic principle for PA with PR is as follows. A small change in parameters induces a sequence of changes (either small perturbations in timing, or big jumps in states) in a sample path; the effect of such a change on a performance in a long term can be measured by the PR factors, which can be estimated on a single sample path. Thus, the performance gradient can be obtained by the expectation (in some sense depending on the problem) of the realization factor.

In this note, we study the gradient of stationary performance measures of (discrete time) finite state-space Markov chains via WDs and PR. Our analysis will show that the WDs expression for the gradient of a stationary performance measure of finite state Markov chain can be interpreted as the expected PR factor where the expectation is carried out with respect to a distribution that is given through the weak derivative of the transition probability matrix of the Markov chain.

The note is organized as follows. Section II provides a short introduction to PR and WDs. In Section III, we illustrate the relation between the PA via PR and the weak derivative estimator for the stationary performance of a finite state-space Markov chain. In Section IV, we show the application of realization factors to the weak derivative of the transition matrix. In Section V, we deduce unbiased estimators from the

Manuscript received December 12, 2000; revised September 14, 2001 and January 10, 2002. Recommended by Associate Editor X. Zhou. The work of B. Heidergott was supported by Deutsche Forschungsgemeinschaft under Grant He3139/1-1. The work of X.-R. Cao was supported by a Grant from Hong Kong RGC.

B. Heidergott was with The Hong Kong University of Science and Technology, Kowloon, Hong Kong. He is now with the Department of Mathematics and Computing Science, Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands (e-mail: b.heidergott@tue.nl).

X.-R. Cao is with the Department of Electrical and Electronic Engineering, The Hong Kong University of Science and Technology, Kowloon, Hong Kong (e-mail: eecao@ust.hk).

Publisher Item Identifier 10.1109/TAC.2002.800648.

results obtained in Section III. We conclude the note with a discussion on the relation between realization factors and WDs.

II. BACKGROUND ON PERTURBATION REALIZATION AND WEAK DERIVATIVES

A. PA via PR

The basic principle for PA via PR is to decompose the performance sensitivity into the effect of a set of perturbations (big or small) in a sample path, which can be measured precisely by a quantity called PR factor. The idea was first applied to infinitesimal perturbations in queueing networks [1], and has been further developed to the case of discrete time Markov chains in [3] and [4].

Let $\mathbf{X} = \{X_k; k \geq 0\}$ be an ergodic Markov chain with finite state-space $S = \{1, \dots, M\}$ and transition probability matrix P . Let $f: S \rightarrow \mathbb{R}$ be a performance function and write f in vectorial notation by $f = (f_1, \dots, f_M)^T$, with $f_i = f(i)$ for $1 \leq i \leq M$, where “ T ” represents the transpose. We denote the unique stationary distribution of $\mathbf{X}$ by $\pi = (\pi_1, \dots, \pi_M)$, and the stationary performance of $\mathbf{X}$ is thus given by

$$\eta = E_\pi(f) = \sum_{i=1}^M \pi_i f_i = \pi f.$$

Let Q be a nonzero square matrix with

$$Qe = 0 \text{ for } e = (1, 1, \dots, 1)^T \quad (1)$$

and assume that a neighborhood of $\delta = 0$, denoted by Θ , exists, so that for any $\delta \in \Theta$ the matrix $P(\delta) = P + \delta Q$ is a transition probability matrix on S . Denote the performance measure associated with $P(\delta)$ by $\eta_Q(\delta)$ (which implies $\eta = \eta_Q(0)$). The derivative of η in the direction of Q is defined as

$$\frac{d\eta_Q}{d\delta} := \left. \frac{d\eta_Q(\delta)}{d\delta} \right|_{\delta=0} = \lim_{\delta \rightarrow 0} \frac{\eta_Q(\delta) - \eta}{\delta}.$$

In this setup, a perturbation means that the Markov chain is perturbed from one state i to another state j . For example, consider the case where $q_{ki} = -\delta$, $q_{kj} = \delta$, and $q_{kl} = 0$ for all $l \neq i, j$. Suppose that in the original sample path the system is in state k and jumps to state i , then in the perturbed path it may jump to state j instead. Thus, we study two independent Markov chains $\mathbf{X} = \{X_n; n \geq 0\}$ and $\mathbf{X}' = \{X'_n; n \geq 0\}$ with $X_0 = i$ and $X'_0 = j$; both of them have the same transition matrix P . The *realization factor* is defined as [4]:

$$d(i, j) = E \left[\sum_{n=0}^{\infty} (f(X'_n) - f(X_n)) \middle| X_0 = i, X'_0 = j \right] \quad (2)$$

for $i, j = 1, \dots, M$. Thus, $d(i, j)$ represents the long term effect of a change from i to j on the system performance.

If P is irreducible, then with probability one the two sample paths of $\mathbf{X}$ and $\mathbf{X}'$ will merge together. That is, there is a random number $L(i, j)$ such that

$$X'_{L(i,j)} = X_{L(i,j)}$$

provided that $X_0 = i$, $X'_0 = j$. Therefore, from the Markov property

$$E \left[\sum_{n=L(i,j)}^{\infty} (f(X'_n) - f(X_n)) \middle| X_0 = i, X'_0 = j \right] = 0$$

and (2) becomes

$$d(i, j) = E \left[\sum_{n=0}^{L(i,j)} (f(X'_n) - f(X_n)) \middle| X_0 = i, X'_0 = j \right] \quad (3)$$

for $i, j = 1, \dots, M$. The matrix $D \in \mathbb{R}^{M \times M}$, with $D_{ij} = d(i, j)$, is called a *realization matrix*. It is shown in [4] and [3] that D satisfies

the Lyapunov equation $D - PDP^T = F$, where $F = ef^T - fe^T$, and $e = (1, 1, \dots, 1)^T$, and the performance derivative is

$$\frac{d\eta_Q}{d\delta} = \pi Q D^T \pi^T. \quad (4)$$

Furthermore, $D_{ij} = g_j - g_i$, where $g = (g_1, \dots, g_M)$ is the potential vector satisfying the Poisson equation $(I - P + e\pi)g = f$, and (4) can be written as

$$\frac{d\eta_Q}{d\delta} = \pi Q \sum_{l=0}^{\infty} (P^l - e\pi) f. \quad (5)$$

In Markov chain literature, the matrix $\sum_{l=0}^{\infty} (P^l - e\pi) = (I - P + e\pi)^{-1} - e\pi$ is sometimes called the *deviation* or the *fundamental matrix*.

B. WDs

WDs provide an approach to write gradients as differences between expectation with respect to appropriately chosen probability measures. More formally, let $(\mathbf{E}, \mathcal{E})$ denote a Polish measurable space and let $\{\mu_\theta; \theta \in \Theta\}$, with $\Theta := (a, b) \subset \mathbb{R}$, be a family of probability measures on $(\mathbf{E}, \mathcal{E})$. We call μ_θ *weakly differentiable at θ* if a signed finite measure μ'_θ exists, such that for any continuous bounded real-valued functions f on $(\mathbf{E}, \mathcal{E})$ it holds that

$$\lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \left(\int_{\mathbf{E}} f(s) \mu_{\theta+\Delta}(ds) - \int_{\mathbf{E}} f(s) \mu_\theta(ds) \right) = \int_{\mathbf{E}} f(s) \mu'_\theta(ds).$$

Note that μ'_θ is not a probability measure. To see this, take $f = 1$, which implies $\int_{\mathbf{E}} \mu'_\theta(ds) = 0$. Hence, μ'_θ has positive and negative parts. However, any finite signed measure can be written as difference between two probability measures (apply, for example, the Hahn–Jordan decomposition).

We call a triple $(c_\theta, \mu_\theta^+, \mu_\theta^-)$, where $\mu_\theta^\pm$ are probability measures on $(\mathbf{E}, \mathcal{E})$ and c_θ is a finite number, a weak derivative of μ_θ if for any continuous bounded function f on $(\mathbf{E}, \mathcal{E})$ it holds that

$$\begin{aligned} \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \left(\int_{\mathbf{E}} f(s) \mu_{\theta+\Delta}(ds) - \int_{\mathbf{E}} f(s) \mu_\theta(ds) \right) \\ = \int_{\mathbf{E}} f(s) \mu'_\theta(ds) \\ = c_\theta \left(\int_{\mathbf{E}} f(s) \mu_\theta^+(ds) - \int_{\mathbf{E}} f(s) \mu_\theta^-(ds) \right). \end{aligned} \quad (6)$$

The probability measure μ_θ^+ is called the (normalized) positive part of μ'_θ and μ_θ^- is called the (normalized) negative part of μ'_θ , respectively. Note that the previous presentation of μ'_θ is not unique.

Example 1: Let μ_θ be the uniform distribution on the interval $[0, \theta]$ for $0 < \theta \leq a$, with $a < \infty$. For any continuous $f: [0, a] \rightarrow \mathbb{R}$, it holds that

$$\begin{aligned} \frac{d}{d\theta} \int f(x) \mu_\theta(dx) \\ = \frac{d}{d\theta} \left(\frac{1}{\theta} \int_0^\theta f(x) dx \right) \\ = \frac{1}{\theta} f(\theta) - \frac{1}{\theta^2} \int_0^\theta f(x) dx \\ = \frac{1}{\theta} \left(\int f(x) \delta_\theta(dx) - \int f(x) \mu_\theta(dx) \right) \end{aligned}$$

where δ_x denotes the Dirac measure in x . Hence, $(1/\theta, \delta_\theta, \mu_\theta)$ is a weak derivative of μ_θ . Observe that no measure ν on $[0, a]$ exists, such that μ_θ^+ and μ_θ^- are absolutely continuous with respect to ν .

We now turn to Markov chains on the finite state-space $S = \{1, \dots, M\}$ and consider the family $P(\delta) = P + \delta Q$ of transition probability matrices introduced in Section II-A. Each row of P can be viewed as a probability measure on the state space and hence

there exist WDs with respect to δ . For any i, j , let $Q_{ij}^+ = \max(Q_{ij}, 0)$ and $Q_{ij}^- = \max(-Q_{ij}, 0)$. Note that $Qe = 0$ implies that

$$c(i) := \sum_{j=1}^M Q_{ij}^+ = \sum_{j=1}^M Q_{ij}^-$$

for all i . From a measure theoretic point of view, $(Q_{ij}^\pm: 1 \leq j \leq M)$ are finite measures on S with total mass $c(i)$. Re-scaling $(Q_{ij}^\pm: 1 \leq j \leq M)$ by $c(i)$ yields probability measures $(P_{ij}^\pm: 1 \leq j \leq M)$ on S . Note that we have to avoid dividing by zero, since $c(i)$ might be zero. Therefore, we set

$$P_{ij}^\pm = \begin{cases} \frac{Q_{ij}^\pm}{c_i} & \text{for } c_i > 0 \\ P_{ij} & \text{for } c_i = 0. \end{cases}$$

It is easily checked that for all i, j it holds that

$$Q_{ij} = c_i (P_{ij}^+ - P_{ij}^-). \quad (7)$$

Recall that $P(\delta)$ is affine linear in δ and that Q is just the derivative of $P(\delta)$ with respect to δ , which implies that

$$\frac{d}{d\delta} P(\delta)_{ij} = Q_{ij} \stackrel{(7)}{=} c_i (P_{ij}^+ - P_{ij}^-). \quad (8)$$

Note that the right-hand side of the aforementioned expression is independent of δ and we set

$$P' := \frac{d}{d\delta} P(\delta).$$

Let C be a square matrix with $C_{ii} = c_i, 1 \leq i \leq M$, and otherwise zero, then (8) reads

$$P' = C (P^+ - P^-)$$

and for any $f \in \mathbb{R}^S$ it holds that

$$\frac{d}{d\delta} P(\delta)f = P'f = C (P^+f - P^-f). \quad (9)$$

If such a representation of P' exists, then $P(\delta)$ is called weakly differentiable and (C, P^+, P^-) is called a weak derivative of $P(\delta)$. It has been shown in [8] that if $P(\delta)$ is weakly differentiable and ergodic then

$$\frac{d\eta_Q(\delta)}{d\delta} = \frac{d}{d\delta} \bigg|_{\delta=0} \pi(\delta)f = \pi \sum_{l=0}^{\infty} P' P^l f \quad (10)$$

where $\pi(\delta)$ denotes the stationary distribution associated with $P(\delta)$ (which implies $\pi(0) = \pi$). In particular, weak differentiability of $P(\delta)$ implies finiteness of the right-hand side of the above expression, see [8]. In Section III, we will contrast (10) to (5) in order to establish the relation between realization probabilities and WDs.

III. DIFFERENTIATING THE STATIONARY DISTRIBUTION OF A MARKOV CHAIN

We study the performance derivative of the Markov chain $\mathbf{X}$, as defined in Section II-A. The derivative of η with respect to δ can be obtained in a closed analytical form, see (5). However, the matrix Q in (5) is not a stochastic matrix, that is, we cannot interpret Q as a transition matrix of the Markov chain $\mathbf{X}$. Set

$$A := \sum_{l=0}^{\infty} (P^l - e\pi)$$

then

$$\frac{d\eta_Q(\delta)}{d\delta} = \pi Q A f.$$

As shown in [2] and [3], the entries of A can be estimated on a single sample path, which gives rise to the following estimation procedure for $d\eta_Q/d\delta$. First, estimate A on a sample path, and then evaluate $Q A f$ by simple matrix-vector multiplication. This then yields an estimator for $d\eta_Q/d\delta$. The question of whether or not this estimator is unbiased

depends on the estimator for A . Various estimators for A both biased and unbiased are discussed in [3].

According to Section II, an alternative way of facilitating (5) for simulation is to write Q as the difference of two transition matrices and to translate (5) into the difference between two experiments. In what follows, we explain this approach in more detail.

By definition

$$P(X_{k+1} = j | X_k = i) = P_{ij}$$

and (8) implies

$$\begin{aligned} \frac{d}{d\delta} E[f(X_{k+1}) | X_k = i] \\ &= \frac{d}{d\delta} \sum_{j=1}^M P_{ij}(\delta) f_j \\ &= \sum_{j=1}^M c(i) (P_{ij}^- f_j - P_{ij}^+ f_j) \\ &= c(i) (E[f(X_{k+1}) | X_k^+ = i] - E[f(X_{k+1}) | X_k^- = i]) \end{aligned}$$

where

$$P(X_{k+1}^\pm = j | X_k^\pm = i) = P_{ij}^\pm.$$

Using (8), we now rewrite (5) as the difference between two stochastic experiments. Denote the i th row of $P^\pm$ by $p_i^\pm$, that is, $p_i^\pm = (P_{ij}^\pm: 1 \leq j \leq M)$ is a probability distribution on S for all i . By calculation

$$\begin{aligned} \frac{d\eta_Q}{d\delta} &= \pi Q \sum_{l=0}^{\infty} (P^l - e\pi) f \\ &= \lim_{n \rightarrow \infty} \pi \left(Q \sum_{l=0}^n P^l f - Q \sum_{l=0}^n e\pi f \right) \\ &\stackrel{(1)}{=} \lim_{n \rightarrow \infty} \pi Q \sum_{l=0}^n P^l f \\ &\stackrel{(7)}{=} \lim_{n \rightarrow \infty} \sum_i \pi_i c(i) (p_i^+ - p_i^-) \sum_{l=0}^n P^l f \\ &= \lim_{n \rightarrow \infty} \sum_i \pi_i c(i) \left(p_i^+ \sum_{l=0}^n P^l f - p_i^- \sum_{l=0}^n P^l f \right) \\ &= \sum_i \pi_i c(i) \sum_{l=0}^{\infty} (p_i^+ P^l f - p_i^- P^l f) \\ &= \pi \sum_{l=0}^{\infty} P' P^l f \end{aligned} \quad (11)$$

which is the expression for $d\eta_Q/d\delta$ derived using WDs, cf. (10). In particular, finiteness of the last two sums in the above row of equations follows from finiteness in (10).

Using (11), the expression

$$\pi \sum_{l=0}^{\infty} P' P^l f$$

can be estimated as follows. Let $\mathbf{X}^\pm = \{X_k^\pm: k \geq 0\}$ denote the Markov chain with (a) initial state $X_0^\pm$, (b) for $\mathbf{X}^+$ perform the first transition from X_0^+ to X_1^+ according to $p_{X_0^+}^+$ and generate all other transitions according to P , and (c) for $\mathbf{X}^-$ perform the first transition from X_0^- to X_1^- according to $p_{X_0^-}^-$ and generate all other transitions according to P . With this notation, we obtain

$$\begin{aligned} \frac{d\eta_Q}{d\delta} &= \sum_i \pi_i c(i) \sum_{l=0}^{\infty} (p_i^+ P^l f - p_i^- P^l f) \\ &= \sum_i \pi_i c(i) \sum_{l=0}^{\infty} E[f(X_l^+) - f(X_l^-) | X_0^\pm = i]. \end{aligned}$$

(12)

The aforementioned expression leads to estimation schemes for $d\eta_Q/d\delta$, as we will explain in Section V.

IV. WDS WITH PR FACTORS

In this section, we write the gradient expression via WDs as the expected values of PR factors $d(i, j)$ introduced in Section II-A. The construction of the processes $\mathbf{X}^\pm$ differs from that of $\mathbf{X}$ only through the first transition. More precisely, after the first transition $\mathbf{X}^\pm$ and $\mathbf{X}$ behave stochastically identical, in formula, for all i, j it holds that

$$P(X_{l+1}^\pm = j | X_l^\pm = i) = P(X_{l+1} = j | X_l = i) \quad (13)$$

for $l \geq 1$. Hence, we obtain

$$E \left[\sum_{l=0}^n f(X_l) \middle| X_0 = i \right] = E \left[\sum_{l=1}^{n+1} f(X_l^\pm) \middle| X_1^\pm = i \right].$$

By calculation

$$\begin{aligned} & \sum_{l=0}^{\infty} E \left[(f(X_l^+) - f(X_l^-)) \middle| X_0^\pm = i \right] \\ &= \sum_{j_1 j_2} \sum_{l=0}^{\infty} E \left[1_{X_1^+ = j_1, X_1^- = j_2} f(X_l^+) - f(X_l^-) \middle| X_0^\pm = i \right] \\ &= \sum_{j_1 j_2} E \left[1_{X_1^+ = j_1, X_1^- = j_2} d(j_1, j_2) \middle| X_0^\pm = i \right] \\ &= \sum_{j_1 j_2} d(j_1, j_2) P(X_1^+ = j_1, X_1^- = j_2 | X_0^\pm = i) \\ &= \sum_{j_1 j_2} d(j_1, j_2) P_{ij_1}^+ P_{ij_2}^- \end{aligned}$$

where $1_{X_1^+ = j_1, X_1^- = j_2}$ denotes the indicator function. The previous formula can be phrased as follows. $P_{ij_1}^+ P_{ij_2}^-$ is the joint probability with which the weak derivative of P splits the nominal process at state i to state j_1 for the “+” part and j_2 for the “−” part, respectively. Hence

$$\sum_{j_1 j_2} d(j_1, j_2) P_{ij_1}^+ P_{ij_2}^-$$

is the expected PR factor with respect to the “splitting probability” defined by the weak derivative of P . In particular, we obtain the following overall formula:

$$\frac{d\eta_Q}{d\delta} = \sum_i \pi_i c(i) \sum_{j_1 j_2} d(j_1, j_2) P_{ij_1}^+ P_{ij_2}^-.$$

Elaborating on the interpretation of Q as a scaled difference between two transition probability matrices we have written (4), respectively, (5), in way that allows to use simulation for evaluating $d\eta_Q/d\delta$. Particular estimation schemes will be addressed in Section IV.

V. ESTIMATION SCHEMES

The expression in (12) can be simplified when stopping times are used. To see this, define the coupling time of $\mathbf{X}^+$ and $\mathbf{X}^-$ by $\tau^* = \inf\{l: X_l^+ = X_l^-\}$. Then

$$\begin{aligned} \frac{d\eta_Q}{d\delta} &= \sum_{i=1}^M \pi_i c(i) \sum_{l=1}^{\infty} E [f(X_l^+) - f(X_l^-) | X_0^\pm = i] \\ &= \sum_{i=1}^M \pi_i c(i) E \left[\sum_{l=1}^{\tau^*} f(X_l^+) - \sum_{l=1}^{\tau^*} f(X_l^-) \middle| X_0^\pm = i \right]. \end{aligned}$$

There is close relation between the stopping times τ^* and $L(i, j)$, defined in Section II-A: τ^* counts the number of transitions from the last state *before splitting* until the sample paths merge, whereas $L(i, j)$ counts the number of transition until the sample paths merge provided that the sample path *has split* up to state i and j , respectively, or, more formally, $1_{X_1^+ = j_1, X_1^- = j_2} \tau^*$ is identical with $L(j_1, j_2) + 1$.

A stationary version of $\mathbf{X}$ can be constructed as follows. Fix a state j^* , start the chain $\mathbf{X}$ in j^* , denote the recurrence time to j^* by τ and let σ be uniformly distributed over $\{1, \dots, \tau\}$ independent of everything else. Let the random variable X^* have distribution

$$P(X^* = i) = \frac{E[\tau P(X_\sigma = i | \tau)]}{E[\tau]}$$

then X^* is a stationary version of the process $\mathbf{X}$, see [9]. Hence, we may replace π in the previous estimator by sampling from π , which yields

$$\frac{d\eta_Q}{d\delta} = \frac{1}{E[\tau]} E \left[\tau c(X_\sigma) E \left[\sum_{l=1}^{\tau^*} f(X_l^+) - \sum_{l=1}^{\tau^*} f(X_l^-) \middle| X_0^\pm = X_\sigma \right] \middle| X_0 = j^* \right]$$

where σ is uniformly distributed over $\{1, \dots, \tau\}$ and independent of everything else, or, equivalently

$$\frac{d\eta_Q}{d\delta} = \frac{1}{E[\tau]} E \left[\sum_{l=1}^{\tau} c(X_l) E \left[\sum_{k=1}^{\tau^*} f(X_k^+) - \sum_{k=1}^{\tau^*} f(X_k^-) \middle| X_0^\pm = X_l \right] \middle| X_0 = j^* \right].$$

Elaborating on the fact that the state-space of $\mathbf{X}$ is finite, the above expression can be estimated from a single sample path of the nominal systems using a cut-and-past type of approach; see [2] and [3] for details.

VI. DISCUSSION

We have shown the connections between PR and WD. WD naturally transfers the performance derivative into the performance differences on different sample paths and offers an explanation of the performance derivative as the expected PR factor with respect to the “splitting probability” defined by the WD of the transition kernel P . PR factors provide a mechanism for obtaining a quantitative result for the weak derivative approach. We believe that PR factors can be used for quantitative analysis of many other problems which are involved with comparison of performance difference due to parameter changes and hope that the present note offers such an example.

We conclude with the remark that the PA approach via realization factors is used in [2] to develop η into a Taylor series. A WDs-based approach to developing stationary performance measures into a Taylor series has still to be found. This is topic of further research.

REFERENCES

- [1] X.-R. Cao, *Realization Probabilities: The Dynamics of Queueing Networks*. New York: Springer-Verlag, 1994.
- [2] —, “The MacLaurin series for performance functions of Markov chains,” *Adv. Appl. Probab.*, vol. 30, pp. 676–692, 1998.
- [3] —, “The relations among potentials, perturbation analysis, and Markov decision processes,” *J. Discrete Event Dyna. Syst.*, vol. 8, pp. 71–87, 1998.
- [4] X.-R. Cao and H.-F. Chen, “Potentials, perturbation realization, and sensitivity analysis of Markov processes,” *IEEE Trans. Automat. Contr.*, vol. 42, pp. 1382–1393, 1997.
- [5] P. Glasserman, *Gradient Estimation via Perturbation Analysis*. Norwell, MA: Kluwer, 1991.
- [6] B. Heidergott, “Infinitesimal perturbation analysis for queueing networks with general service time distributions,” *QUESTA*, pp. 43–58, 1999.
- [7] Y.-C. Ho and X.-R. Cao, *Perturbation Analysis of Discrete Event Systems*. Norwell, MA: Kluwer, 1991.
- [8] G. Pflug, *Optimization of Stochastic Models*. Norwell, MA: Kluwer, 1996.
- [9] H. Thorisson, *Coupling, Stationarity, and Regeneration*. New York: Springer-Verlag, 2000.