Measure-Valued Differentiation of Positive Systems: The State-of-the-Art as of Summer 2003

Bernd Heidergott
Vrije Universiteit Amsterdam

Positive systems are stochastic systems in which the relevant variables assume nonnegative values. These systems are quite common in applications where variables represent positive quantities such as probabilities, populations, times, consumption of goods etc. Examples of positive systems are queueing systems, traffic models, Markov chains, stochastic (max,+)-linear systems etc.

The performance of a positive system may depend on a design parameter, say θ, representing, for example, the mean service time at a server in an queueing network or the travel time in a transportation system. When it comes to sensitivity analysis, control or optimisation, one seeks to compute derivatives. The problem when looking for derivatives in positive systems is the following. Consider a positive differentiable mapping *f* from *R* to *R*, i.e., $f(\theta) \geq 0, \theta \in R$. In general, the derivative of *f* with respect to θ, denoted by $f'(\theta)$, won't enjoy the property that $f'(\theta) \geq 0$ for any $\theta \in R$. In other words, the differentiation operator maps positive mappings onto general ones and by taking derivatives, one usually leaves the set of positive mappings. This property of the differentiation operator poses a severe problem when differentiation/optimisation is applied in the context of positive systems.

With the advent of measure-valued differentiation (MVD) a mathematical framework for overcoming this obstacle is available. The key contribution of MVD is that it defines for stochastic operators, such as probability measures, Markov kernels and random variables, derivatives with respect to θ (the design parameter) as differences of stochastic operators of the same kind. For example, the MVD derivative of a probability measure with respect to θ is given as the difference between two appropriate probability measures. Hence, derivatives of positive operators are (pairs of) positive operators. It is safe to say that the MVD framework constitutes a paradigm shift in the analysis of positive systems. The epistemological impact of MVD is the following:

1. Formulae containing *differences can be interpreted as derivatives* in the MVD sense. This links for example potentials and derivatives. Moreover, the deviation operator, which evaluates the difference between a Markov chain with given initial state and its stationary version, occurs in the MVD expression for the derivative of the stationary distribution. Here, a link between stability theory and gradient estimation/optimisation is established.

2. MVD formulae can be read/interpreted in two ways. First, as a mathematical expression for derivatives providing theoretical insight. Secondly, any derivative formula/expression possesses a direct interpretation as unbiased gradient estimator. Hence, MVD overcomes the dichotomy between proving the existence of a derivative and finding an unbiased estimator for estimating it. In other words, *MVD is as well a mathematical theory of derivatives of positive operators as it is a theory of simulation/estimation*.

1. Approach

In this section, the MVD approach will be explained for probability distributions and Markov kernels. Furthermore, application of MVD to stochastic (max,+)-systems will be discussed. Finally, a general set-up will be formulated and an overall research goal will be described.

1b. An Introduction to Measure-Valued Differentiation

1b.1. Measure-Valued Differentiation of Probability Distributions

The general set-up for MVD is as follows. Let $\Theta \subset R$ be an open interval and μ_θ, for $\theta \in \Theta$, a probability measure on a measurable space (S, T). Denote the set of mappings $g : S \to R$ with finite μ_θ integral for any

$\theta \in \Theta$ by $L^1(\mu_\theta : \theta \in \Theta)$. Let $D \subset L^1(\mu_\theta : \theta \in \Theta)$. Typical choices for *D* are the set of continuous bounded mappings, the set of mappings that are bounded by some particular function (which will often be a polynomial), or the set of bounded measurable mappings. The probability measure μ_θ is called *D*-differentiable if a finite signed measure μ_θ' exists such that

$$\forall g \in D : \frac{d}{d\theta} \int g(x)\mu_\theta(dx) = \int g(x)\mu_\theta'(dx) .$$

The key observation of [Pflg96] is that μ_θ' can be written as re-scaled difference between two probability measures, and a triple $(c_\theta, \mu_\theta^+ \mu_\theta^-)$, with μ_θ^+ and μ_θ^- probability measures and c_θ a constant, is called a *D*-derivative of μ_θ if

$$\mu_\theta' = c_\theta(\mu_\theta^+ - \mu_\theta^-) .$$

Hence, if $(c_\theta, \mu_\theta^+ \mu_\theta^-)$ is a *D*-derivative of μ_θ, then

$$\forall g \in D : \frac{d}{d\theta} \int g(x)\mu_\theta(dx) = c_\theta\left(\int g(x)\mu_\theta^+(dx) - \int g(x)\mu_\theta^-(dx)\right).$$

Higher-order derivatives are defined in the same vein. For example, let μ_θ be the exponential distribution with rate $\theta \in \Theta = (0,\infty)$ and denote by ν_θ the Gamma-$(2,\theta)$-distribution with rate $\theta \in \Theta$, i.e., ν_θ is the distribution of the sum of two independent samples from μ_θ. Let *D* be the set of real-valued mappings that are bounded by a polynomial, then a *D*-derivative of μ_θ is obtained from $(\theta^{-1}, \mu_\theta, \nu_\theta)$, see [Pflg96] and [HeidVazPflg03]. The above definition is an extension of the concept of "weak differentiation" in [Pflg96], where *D* is fixed to the set of continuous bounded mappings.

Any *D*-derivative can be immediately translated to an unbiased gradient estimator. Specifically, the generic estimator for

$$\frac{d}{d\theta} \int g(x)\mu_\theta(dx)$$

uses the difference of two experiments: one experiment driven by μ_θ^+ and the other driven by μ_θ^-. Taking the difference between the outcomes of the "+"-experiment and the "-"-experiment and re-scaling it by c_θ yields an unbiased estimator for the gradient. While this is the most general set-up, there are often other interpretations of the above formula in terms of gradient estimators available, see [HeidVaz00] for details. In particular, the well-known family of perturbation analysis gradient estimators are instances of MVD. For details on perturbation analysis, see [HoCao91], [Glas91] or [Cao94].

A MVD product rule for probability measures exits. This product rule is an analogue to the product rule of differentiation in analysis, and having computed the *D*-derivative of a probability measure μ_θ, a measure-valued derivative of the (independent) finite product of μ_θ can be obtained, see [HeidVaz00] and, for the case of continuous bounded mappings, [Pflg96].

1b.2. Measure-Valued Differentiation of Markov Kernels

MVD can be applied directly to stochastic processes, using a general state space *S* that may represent whole trajectories of dynamical processes. However, direct use of this approach would require knowledge of the underlying probability measure μ_θ, which may be impossible to evaluate. Instead, a representation is used that isolates the dependency on θ at each transition of a Markov chain, as will be explained in the following.

The definition of *D*-differentiability of probability measures readily extends to Markov kernels P_θ as follows. We say that P_θ is *D*-differentiable if a transition kernel P_θ' on *(S,T)* exists such that for all $y \in S$:

$$\forall g \in D : \frac{d}{d\theta} \int g(x)P_\theta(y;dx) = \int g(x)P_\theta'(y;dx) .$$

A triple $(c_\theta(\cdot), P_\theta^+, P_\theta^-)$, with $c_\theta(y)$ a random variable and P_θ^+ and P_θ^- Markov kernels on *(S,T)*, is called a D-derivative of P_θ if

$$P_\theta' = c_\theta\left(P_\theta^+ - P_\theta^-\right).$$

Hence, provided that P_θ is *D*-differentiable with *D*-derivative $(c_\theta(\cdot), P_\theta^+, P_\theta^-)$, it holds that

$$\forall g \in D: \frac{d}{d\theta} \int g(x) P_\theta(y;dx) = c_\theta(y)\Big(\int g(x) P_\theta^+(y;dx) - \int g(x) P_\theta^-(y;dx)\Big),$$

for any $y \in S$. For a detailed discussion on *D*-differentiability and *D*-derivatives for Markov kernels, see [HeidHorWeis03]. Sufficient condition for the existence *D*-derivatives of Markov kernels are given in [HeidHorWeis03].

1b.3. Two Basic Applications of Measure-Valued Differentiation

The simplest example of a *D*-differentiable Markov kernel is the following. Let *P* and *Q* be two Markov kernels on *(S, T)* and let *D* be a set of real-valued mappings that are integrable with respect to *P* and *Q*. Consider the Markov kernel

$$P_\theta = \theta Q + (1-\theta) P,$$

for $\theta \in \Theta = [0,1]$. Mixing two Markov kernels in the above way is called a **Bernoulli scheme**. It is easily checked that P_θ is *D*-differentiable with *D*-derivative $(1, Q, P)$, for any $\theta \in \Theta = [0,1]$. Although the above model is rather simple, many problems that are of interest in applications can be modelled by it. For example, *P* may represent a queueing network with exponential inter-arrival time while *Q* models the identical queueing system with, say, uniformly distributed inter-arrival times. At $\theta = 0$, P_θ models the variant of the queueing system with Poisson arrival stream and the *D*-derivative of P_θ can be interpreted as derivative of the transition dynamic of the queueing network in direction of the uniform distribution as distribution of the inter-arrival times. In other words, the *D*-derivative of P_θ expresses the effect of replacing an exponential inter-arrival time distribution by a uniform distribution. The Bernoulli scheme thus allows for **non-parametric sensitivity analysis**.

A variant of the Bernoulli scheme is the so-called **variability expansion**. Here, *P* is taken to be the deterministic counterpart of *Q*. For example, if *Q* is the transition kernel of the waiting times in an G/G/1 queue, then *P* is the transition obtained from replacing the stochastic inter-arrival time and the stochastic service time by their means. At $\theta = 0$, P_θ models the deterministic variant of the waiting times (which is a simple deterministic mapping) and the *D*-derivative of P_θ expresses the effect of replacing mean values by the true random variables (read: how sensitive is the performance with respect to replacing random variables by their means?). In other words, variability expansion allows **combining deterministic and stochastic analysis**

1b.4. The Relation between MVD of Probability Distributions and Markov Kernels

Often the Markov chain driving the system dynamic can be influenced through *input distributions*. For example, in a queueing system the overall Markov chain depends on the service time and inter-arrival time distributions of the system. Typically such input distributions are of simple nature, such as the Bernoulli distribution for modelling stochastic routing or the exponential distribution for modelling time variables. The Markov kernel of the Markov chain describing the system process then reflects the interaction between the simple input distributions.

Let the Markov kernel depend on θ only through an input distribution, say, μ_θ. Under quite general conditions, *D*-differentiability of μ_θ implies *D*-differentiability of the Markov kernel and a *D*-derivative of the Markov kernel is directly obtained from a *D*-derivative of μ_θ, see [HeidVaz00] and [HeidVazPflg03]. This simplifies the complexity of applying measure-valued derivatives dramatically: one only has to study the measure-valued derivative of the input distribution(s), a measure-valued derivative of the associated Markov kernel is then given through a simple formula in canonical form. In other words: measure-valued differentiability on the *input level* yields measure-valued differentiability on the *operator level* (here, on the level of the Markov kernel). It becomes thus possible to integrate a library of *D*-derivatives of simple distributions into a simulation package and automatically generate an unbiased gradient estimator on the operator level.

In addition, the decomposition is helpful in establishing robust on-line sensitivity estimation for Markov processes: only the distribution of the controlled variables μ_θ are assumed known, while the distribution of the rest of the underlying variables in the kernel may be unknown. This way a process under control may be observed and those observations may be used to drive the estimation as well. In particular, for many problems a series of perturbations are added to the observations in order to model stochastic noise, making the normal distribution a particularly useful model. MVD reflects the hierarchical structure inherited by the structure of the problem and it is a promising approach for on-line control of positive systems.

2. The Product Rule of MVD

On the operator-level, the key feature of MVD is the *product-rule of MVD*. In the following we address the product rule for Markov kernels. Let P_θ^n denote the *n*-fold product of Markov kernel P_θ and assume that P_θ is *D*-differentiable. Under appropriate conditions, *D*-differentiability of P_θ implies:

$$\left(P_\theta^n\right)' = \sum_{j=0}^{n-1} P_\theta^j \left(P_\theta\right)' P_\theta^{n-j-1},$$

see [HeidVaz00]. Notice that the above formula is a mathematical convenient way of expressing the fact that for any $y \in S$:

$$\forall g \in D : \frac{d}{d\theta} \int g(x) P_\theta^n(y;dx) = \sum_{j=0}^{n-1} \int g(x) \left(P_\theta^j \left(P_\theta\right)' P_\theta^{n-j-1} \right)(y;dx) .$$

Inserting a *D*-derivative of P_θ into the above formula immediately yields the following instance of a *D*-derivative of P_θ^n :

$$\left(P_\theta^n\right)' = \left(c_\theta, \sum_{j=0}^{n-1} P_\theta^j P_\theta^+ P_\theta^{n-j-1}, \sum_{j=0}^{n-1} P_\theta^j P_\theta^- P_\theta^{n-j-1} \right).$$

As explained in Section 1b.4, a *D*-derivative of P_θ can usually be deduced from that of an input distribution. By virtue of the product rule, a generic unbiased gradient estimator for derivatives of the finite horizon performance of a positive system with respect to θ can be incorporated into a simulation package. In particular, the input distributions and their *D*-derivatives can be given *a priori*, whereas the actual Markov kernel is problem dependent.
To illustrate the importance of the product rule, revisit the Bernoulli scheme in Section 1b.3. At $\theta = 0$, $P_\theta = P$ models the variant of the queueing system with Poisson arrival stream. Let $X(k)$ denote the queue length process governed by transition kernel P and let $X(k;j)$ be defined as follows: the transition of $X(k;j)$ to $X(k+1;j)$ is for $k \neq j$ driven by *P* and for $j = k$ the transition is obtained by replacing the inter-arrival time distribution by a uniform distribution, i.e., the state transition is driven by *Q*. Then, the above equation reads in explicit form

$$\forall g \in D : \frac{dP}{dQ} E\left[g(X(n))\right] = \sum_{j=0}^{n-1} E[g(X(n;j))] - nE[g(X(n))],$$

where we write *dP/dQ* to express the fact that the expression on the left-hand side of the above formula can be interpreted as the derivative of $E\left[g(X(n))\right]$ in direction of *Q*. The effect that one of the variants coincides with the nominal one occurs rather frequently in applications and increases the applicability of MVD to simulation. For examples on this, see [Heid01b], [HeidVaz00] and [HeidVaz01].

3. MVD for Stationary Distributions

Let π_θ denote the unique invariant distribution of P_θ (existence is assumed). Provided that P_θ is *D*-differentiable, it holds under appropriate stability conditions that

$$\left(\pi_\theta\right)' = \pi_\theta \left(P_\theta\right)' \sum_{n \geq 0} \left(P_\theta^n - \Pi_\theta\right),$$

where Π_θ denotes the transition kernel that maps any probability distribution unto π_θ . A first proof of this relation can be found in [Pflg96], where *D* is taken to be the set of continued bounded mappings, and a proof for general *D* is given in [HeidHorWeis02]. The operator

$$D_\theta = \sum_{n \geq 0} \left(P_\theta^n - \Pi_\theta\right)$$

is well-known in the theory of Markov chains where it is called *deviation operator*, or, in case the state-space of the Markov chain is (at most) denumerable, *deviation matrix*. See for example [CoolDoor02] for a detailed discussion of the relation between the deviation matrix, the fundamental matrix and group inverse, respectively, the Drazin inverse. See also [Cao98b] and [HeidCao02] for a discussion of the role of the deviation matrix in perturbation analysis. Elaborating on the deviation operator, MVD yields

$$(\pi_\theta)' = \pi_\theta (P_\theta)' D_\theta .$$

For the Bernoulli scheme with finite state-space, sufficient conditions for analyticity of π_θ with respect to θ have been established in [Cao98a]. Specifically, it holds for sufficiently small Δ that

$$\pi_{\theta+\Delta} = \pi_\theta \sum_{n \geq 0} \Delta^n \Big((P_\theta)' D_\theta \Big)^n .$$

Based on MVD, the above result could be extended to general-state space Markov chains and general *D*-derivatives (leading to a more complex formula), see [HeidHor03]. In particular, lower bounds for the radius of convergence of the Taylor series have been establish in the aforementioned paper. Ongoing research is on applications to the so-called power series algorithm (PSA), an analytical approach for computing stationary characteristics of Markovian queueing systems, see [Blnc91], [Blnc92a], [Blnc92b] and [vdHt96]. Specifically, the key problem in applying PSA is to establish analyticity of the stationary characteristics, and finding a lower bound for the radius of convergence for the Taylor series expansions may help overcoming this obstacle of PSA. Only in special cases can the deviation matrix be obtained in a closed form, see [Kl98] and [KlSpk01], and part of the proposed research will be to develop efficient estimation algorithms.

A particular interesting case for applications will be variability expansion, see Section 1b.4. Developing π_θ, for $\theta = 1$ into a Taylor series at $\theta = 0$, will provide a scheme for approximating stationary performance characteristics through mixed expressions of stochastic and deterministic entities. This offers a potential for decreasing the complexity of computing stationary characteristics. This approach has already provided good results in the finite horizon case, see [Heid02] for an application to the computation of expected values of transient waiting times in a G/G/1 queue and [HeidBmSch02] for an application to model predictive control.

4. MVD for (Max,+)-Linear Systems

The above presentation focussed on the application of MVD to Markov kernels. In the following we briefly review the application of MDV to (max,+) algebra. The basic steps are similar to those for Markov chains and this similarity is best seen when operator language is used. For details on (max,+) algebra, see [Cngh79] or [BCOQ92]. A state-of-the-art review of applications of (max,+) algebra to transportation problems is given in [HeiddVrs01].

Elaborating on MVD, one can establish sufficient conditions such that *D*-differentiability of the elements of a matrix over the (max,+) algebra implies *D*-differentiability of the matrix itself, for appropriately defined sets *D*. See [Heid01a]. Like in the Markov chain case, MVD provides a hierarchical approach to perturbation analysis for (max,+)-linear systems. One of the highlights of this approach is that if a particular service time in a (max,+)-linear queueing network is *D*-differentiable [input level], then the matrix modelling the network dynamic is *D*-differentiable [operator level] and, by virtue of the product rule of MVD, the state-vector of the system is *D*-differentiable. This fact can then be translated into expressions for the derivative of the expected value of the finite time-horizon performance of the system measured by performance functions out of *D*, see [Heid01a].

A key performance indicator for (max,+)-linear systems is the so-called *Lyapunov exponent* . For example, in a queueing application, the Lyapunov exponent is given by the inverse throughput of the queueing network. Only in special cases can the Lyapunov exponent be obtained in a closed form and the approach for computing the Lyapunov exponent pre-dominant in the literature is Taylor series expansion. Key reference are [BcHng00a], [BcHng00b] and [GbHng00]. Here, MVD provides the means of unifying the results for Taylor series expansions of (max,+) systems and Markov chains

Literature references

Papers listed as "submitted" can be accessed via http://staff.feweb.vu.nl/bheidergott .

[BCOQ92] F. Baccelli, G. Cohen, G. Olsder, and J.-P. Quadrat. *Synchronization and Linearity*. John Wiley and Sons, New York, 1992.

[BcHng00a] F. Baccelli and D. Hong. Analytic expansions of (max,+) Lyapunov exponents. *Ann. Appl. Probab*., 10:779-827, 2000.

[BcHng00b] F. Baccelli and D. Hong. Analyticity of Iterates of Random Non-Expansive Maps**.** *Adv. Appl. Prob.* 32:193-220, 2000.

[Blnc91] J.P.C. Blanc. The power-series algorithm applied to cyclic polling systems, *Stochastic Models* **7**:527-545, 1991.

[Blnc92a] J.P.C. Blanc. The power-series algorithm applied to the shortest-queue model, *Operations Research* 40:157-167, 1992.

[Blnc92b] J.P.C. Blanc. Performance evaluation of polling systems by means of the power-series algorithm, *Annals of Operations Research* 35:155-186, 1992.

[Cao98a] X.-R. Cao. The Maclaurin Series for Performance Functions of Markov Chains. *Advances in Applied Probability*, 30:676-692, 1998.

[Cao98b] X.-R. Cao. The relations among potentials, perturbation analysis, and Markov decision processes. *DEDS*, 8:71-87, 1998.

[Cao94] X.-R. Cao. *Realization probabilities: The dynamics of queueing systems*. Springer, New York, 1994.

[CoolDoor02] P. Coolen-Schrijner and E.A. van Doorn. The deviation matrix of a continuous-time Markov chain. *Probab. Engrg. Inform. Sci.* 16:351-366, 2002.

[Cngh79] R. Cuninghame-Green. *Minimax algebra*. Vol. 166 of Lecture Notes in Economics and Mathematical Systems. Springer, Berlin, 1979.

[GbHng00] S. Gaubert and D. Hong. Series Expansions of Lyapunov Exponents and Forgetful Monoids. *Technical Report*, RR-3971, INRIA Rocquencourt. 2000.

[Glas91] P. Glasserman. *Gradient Estimation via Perturbation Analysis*. Kluwer Academic, Boston, 1991.

[Heid01a] B. Heidergott. A differential calculus for random matrices with applications to (max,+)-linear stochastic systems. *MOR*, vol. 26, pages 679-699, 2001.

[Heid01b] B. Heidergott. Option pricing via Monte Carlo Simulation: A weak derivative approach. *Probability in Engineering and Informational Sciences*, vol. 15, pages 335-349, 2001.

[Heid02] B. Heidergott. Variability expansion for performance characteristics of (max,plus)-linear systems. In *International Workshop on DES (WODES)*, Zaragoza, Spain, pages 245-250, 2002.

[HeidCao02] B. Heidergott and X.-R. Cao. A note on the relation between realization probabilities and weak derivatives. *IEEE Trans. on Automatic Control*, vol. 47, pages 1112-1115, 2002.

[HeidHor03] B. Heidergott and A. Hordijk. Taylor series expansions for stationary Markov chains. To appear in *Adv. Appl. Prob*., 2003.

[HeiddVrs01] B. Heidergott and R. de Vries. Towards a control theory for transportation networks. *DEDS*, vol. 11, pages 371-398, 2001.

[HeidVaz00] B. Heidergott and F. Vázquez-Abad. Measure-valued differentiation for stochastic processes: the finite horizon case*. EURANDOM Report 2000-033*, 2000. (submitted)

[HeidVaz01] B. Heidergott and F. Vázquez-Abad. Measure-valued differentiation for stochastic processes: the random horizon case*. GERAD Report G-2001-18*, 2001. (submitted)

[HeidBmSch02] B. Heidergott, T. van den Boom and B. de Schutter. Complexity reduction in MPC for stochastic max-plus-linear systems by variability expansion. In *Proceedings of the CDC*, Las Vegas, Nevada, USA, pages 3567-3572, 2002.

[HeidHorWeis02] B. Heidergott, A. Hordijk and H. Weißhaupt. Measure-valued differentiation for stationary Markov chains. *EURANDOM Report 2002-27*. 2002 (submitted)

[HeidHorWeis03] B. Heidergott, A. Hordijk and H. Weißhaupt. Derivatives of Markov kernels and their Jordan decomposition. 2003 (submitted)

[HeidVazPflg03] B. Heidergott, F. Vázquez-Abad and Pflug. Measure-valued differentiation for Markov chains. 2003 (submitted)

[vdHt96] W. van den Hout. *The Power-Series-Algorithm*. PhD thesis, Centre for Economic Research, Tilburg, 1996.

[HoCao91] Y.-C. Ho and X.-R. Cao. *Perturbation Analysis of Discrete Event Systems*. Kluwer Academic, Boston, 1991.

[Kl98] G. Koole. The deviation matrix of the $M/M/1/\infty$ and M/M/1/N queue, with applications to controlled queueing models. *Proceedings of the 37th IEEE CDC*, Tampa, FL, pp.56-59, 1998.

[KlSpk01] G. Koole and F. Spieksma. On deviation matrices for birth-death process. *Probability in the Engineering and Informational Sciences.* 15:239-258, 2001.

[Pflg96] G. Pflug. *Optimization of Stochastic Models*. Kluwer Academic, Boston, 1996.