PERTURBATION ANALYSIS

Michael C. Fu

University of Maryland, College Park

Perturbation analysis (PA) is a sample path technique for analyzing changes in the performance of stochastic systems due to changes in system parameters. In terms of stochastic simulation — the main setting for the application of PA — the objective is to estimate *sensitivities* of the performance measures of interest with respect to system parameters while obtaining estimates of performance itself, without the need for additional simulation runs. The primary application is gradient estimation during the simulation of discrete-event systems, e.g., queueing and inventory systems. Besides their importance in sensitivity analysis, these gradient estimators are a critical component in gradient-based simulation optimization methods.

Let $l(\theta)$ be a *performance measure* of interest with *parameter* (possibly vector) of interest θ. We are interested in those systems where $l(\theta)$ cannot be easily obtained through analytical means and therefore must be estimated from sample paths, e.g., via stochastic simulation. We denote by $L(\theta, \omega)$ the *sample performance* obtained from a sample path realization ω such that $l(\theta) = E[L(\theta, \omega)]$. The goal of PA is to efficiently estimate the effects on l of a perturbation $\theta \longrightarrow \theta + \Delta\theta$, using information from a sample path ω at θ. We distinguish between two major classes of problems that PA addresses:

- $\Delta\theta \to 0$: estimating the gradient $\dfrac{dl(\theta)}{d\theta}$, when l is differentiable in θ;
- $\Delta\theta \neq 0$: estimating changes due to a *finite* perturbation, i.e., $l(\theta + \Delta\theta)$.

In the former case, no perturbation is ever actually introduced into the system (or simulation), although the idea of a perturbation may be employed as a heuristic tool in preliminary analysis.

To sort out the abundance of acronyms in the PA field, we provide a brief definition of each corresponding approach, along with a reference. Among gradient estimation techniques, by far the most well-known is infinitesimal perturbation analysis (**IPA**), which simply uses the sample derivative $dL/d\theta$ to estimate $dl/d\theta$. It is straightforward to implement and very computationally efficient; however, as we shall discuss shortly in more detail, its applicability is not universal. The books by Ho and Cao (1991), Glasserman (1991), and Cao (1994) cover IPA in detail. A very general and well-developed extension of IPA is smoothed perturbation analysis (**SPA**), based on the ideas of conditional expectation (Gong and Ho 1987). Although its applicability is quite broad, its implementation is usually very problem dependent. The book by Fu and Hu (1997) covers this method in full generality. Other gradient estimation techniques include rare perturbation analysis (**RPA**), originally based on the thinning of point processes (Brémaud and Vázquez-Abad 1992); structural IPA (**SIPA**), dealing specifically with structural parameters (Dai and Ho 1995); discontinuous perturbation analysis (**DPA**), based on the use of generalized functions (the Dirac-delta function) to model discontinuities in the sample performance function (Shi 1996); and augmented IPA (**APA**), another extension of IPA different from SPA (Gaivoronski et al. 1992). Techniques to estimate the effect of a *finite* perturbation in the parameter include finite perturbation analysis (**FPA**) — Ho et al. (1983); extended perturbation analysis (**EPA**) — Ho and Li (1988); and the augmented chain method — Cassandras and Strickland (1989). A related technique is the standard clock (**SC**) method, based on the uniformization of Markov chains (Vakili 1991). The books by Ho and Cao (1991) and Cassandras (1993) provide further references. In the remainder of this entry, we will concentrate on the gradient techniques IPA and SPA, since they are the most widely known and extensively developed of the PA techniques.

Applicability of IPA: We discuss the applicability of IPA through the use of some simple examples, at the same time contrasting the approach with the likelihood ratio/score function (LR/SF) and weak derivative (WD) estimators. We begin by considering the expectation of a single random variable X. We write this expectation in two forms:

$$\begin{aligned} E[X] &= \int_0^\infty x f(x;\theta)dx, \\ &= \int_0^1 X(\theta;u)du, \end{aligned}$$

where f is the PDF of X. In the first interpretation, the parameter appears inside the density, whereas in the second interpretation it appears inside the random variable defined on an underlying $U(0,1)$ random number. For example, the latter could be the inverse transform $X = F^{-1}$, where F is the CDF of X.

Differentiating $E[X]$, assuming the interchange of expectation and differentiation is permissible (via the dominated convergence theorem), we get

$$\frac{dE[X]}{d\theta} = \int_0^\infty x \frac{df(x;\theta)}{d\theta}dx, \qquad (1)$$

$$= \int_0^1 \frac{dX(\theta;u)}{d\theta}du. \qquad (2)$$

Notice, however, that the conditions for the exchange will be quite different for the two interpretations. In the first interpretation, corresponding to the LR/SF and WD estimators, the conditions will be placed on the underlying density; in the case of discrete-event stochastic simulation, this means the input distributions. Since the input distributions must be known in order to perform the simulation, it is relatively easy to check the conditions. In the second interpretation, corresponding to PA estimators, the conditions will be placed on the sample performance function that is usually defined on an output stochastic process of the system. Since the output process is being simulated precisely because much is unknown about it, the conditions are often more difficult to check.

As an example, consider an exponential random variable X with mean θ. Then $E[X] = \theta$ and $dE[X]/d\theta = 1$. The respective PDF and one random variable representation are given by

$$\begin{aligned} f(x;\theta) &= \frac{1}{\theta}e^{-x/\theta}\mathbf{1}\{x>0\}, \\ X(\theta;u) &= -\theta \ln u, \end{aligned}$$

where $\mathbf{1}\{\cdot\}$ denotes the indicator function. Differentiating, we get

$$\begin{aligned} \frac{df(x;\theta)}{d\theta} &= \left[\frac{x}{\theta^2}\frac{1}{\theta}e^{-x/\theta} - \frac{1}{\theta^2}e^{-x/\theta}\right]\mathbf{1}\{x>0\} \\ &= f(x;\theta)\left[\frac{x}{\theta^2} - \frac{1}{\theta}\right] \\ &= \frac{1}{\theta e}\left[\frac{e}{\theta}\left(1-\frac{x}{\theta}\right)e^{-x/\theta}\mathbf{1}\{0<x\le\theta\}\right. \\ &\quad \left. -\frac{e}{\theta}\left(\frac{x}{\theta}-1\right)e^{-x/\theta}\mathbf{1}\{x>\theta\}\right], \\ \frac{dX(\theta;u)}{d\theta} &= -\ln u = \frac{X(\theta;u)}{\theta}. \end{aligned}$$

The second equation for the density derivative (which is itself *not* a density) expresses the quantity as the difference of two densities multiplied by a constant, known as a "weak derivative" representation. Substituting each of the three expressions into the corresponding equations (1) or (2), yields three unbiased derivative estimators:

$$\begin{aligned} \text{LR/SF:} &\quad \frac{X}{\theta}\left(\frac{X}{\theta}-1\right), \\ \text{WD:} &\quad \frac{1}{\theta e}\left[X^{(2)} - X^{(1)}\right], \\ \text{IPA:} &\quad \frac{X}{\theta}, \end{aligned}$$

where $X^{(1)}$ and $X^{(2)}$ are random variables with PDFs $\frac{e}{\theta}\left(\frac{x}{\theta}-1\right)e^{-x/\theta}$, $x>\theta$, and $\frac{e}{\theta}\left(1-\frac{x}{\theta}\right)e^{-x/\theta}$, $0<x\le\theta$, respectively.

Extending to a function of the underlying random variable, we have

$$\begin{aligned} \frac{dE[L(X)]}{d\theta} &= \int_0^\infty L(x)\frac{df(x;\theta)}{d\theta}dx, \\ &= \int_0^1 \frac{dL}{dX}\frac{dX(\theta;u)}{d\theta}du. \end{aligned}$$

The conditions for interchanging expectation and differentiation are unaltered when differentiating the underlying density, since that portion remains unchanged, whereas they are more involved for the sample path derivative. Basically, in order for the chain rule to be applicable, we need some sort of continuity to hold for the sample performance function with respect to the underlying random variable. This translates into requirements on the form of the performance measure and on the dynamics of the underlying stochastic system such that the interchange

$$\frac{dE[L]}{d\theta} = E\left[\frac{dL}{d\theta}\right] \tag{3}$$

holds. Roughly speaking, sample pathwise continuity of L with respect to θ will result in the interchange being valid. An important structural condition for determining the applicability of IPA for general discrete-event systems modeled as generalized semi-Markov processes is the *commuting condition* (see Glasserman 1991).

The main idea of smoothed perturbation analysis (SPA) is to use conditional expectation to "smooth" out discontinuities in L that cause IPA to fail. This is achieved by selecting a set of sample path quantities $\mathcal{Z}$, called the *characterization*, such that $E[L|\mathcal{Z}]$ – as opposed to L itself – will satisfy the interchange in (3):

$$\frac{dE\left[E[L|\mathcal{Z}]\right]}{d\theta} = E\left[\frac{dE[L|\mathcal{Z}]}{d\theta}\right]. \tag{4}$$

Applying SPA is analogous to the variance reduction technique of conditional Monte Carlo, consisting of two main steps: choosing an appropriate $\mathcal{Z}$ and calculating $dE[L|\mathcal{Z}]/d\theta$. For generalized semi-Markov processes, as well as for other stochastic systems, this is fully explored in Fu and Hu (1997).

Queueing Example: We illustrate IPA and SPA estimators for a single-server, first come, first-served (FCFS) queue. Let A_n be the interarrival time between the $(n-1)$th and nth customer (i.i.d. with PDF f_1 and CDF F_1), X_n the service time of the nth customer (i.i.d. with PDF f_2 and CDF F_2), and T_n the system time (in queue plus in service) of the nth customer. We consider the case where θ is a parameter in the service time distribution, and the sample performance of interest is the average system time over the first N customers $\overline{T}_N = \frac{1}{N}\sum_{n=1}^N T_n$. The system time of a customer for a FCFS single-server queue satisfies the well-known recursive Lindley equation:

$$T_{n+1} = X_{n+1} + (T_n - A_{n+1})^+. \tag{5}$$

The IPA estimator is obtained by differentiating (5):

$$\frac{dT_{n+1}}{d\theta} = \frac{dX_{n+1}}{d\theta} + \frac{dT_n}{d\theta}\mathbf{1}\{T_n \ge A_{n+1}\}, \tag{6}$$

where

$$\frac{dX}{d\theta} = -\frac{dF_2(X;\theta)/d\theta}{dF_2(X;\theta)/dX}.$$

For example, for scale parameters, such as if θ is the mean of an exponential distribution, we have $dX/d\theta = X/\theta$. Using the above recursion, the IPA estimator for the derivative of average system time is given by

$$\begin{aligned}\frac{d\overline{T}_N}{d\theta} &= \frac{1}{N}\sum_{n=1}^{N}\frac{dT_n}{d\theta} \qquad (7)\\ &= \frac{1}{N}\sum_{m=1}^{M}\sum_{i=n_{m-1}+1}^{n_m}\sum_{j=n_{m-1}+1}^{i}\frac{dX_j}{d\theta},\end{aligned}$$

where M is the number of busy periods observed and n_m is the index of the last customer served in the mth busy period ($n_0 = 0$). Implementation of the estimator involves keeping track of two running quantities, one for (6) and another for the summation in (7); thus, the additional computational overhead is minimal, and *no alteration of the underlying simulation is required.* IPA is also applicable to multi-server queues and Jackson-like queueing networks (Jackson networks without the exponential distribution assumptions).

The implicit assumption used in deriving an IPA estimator is that small changes in the parameter will result in small changes in the sample performance. Thus, in the above, this means that the boundary condition in (6) is unchanged by differentiation. In general, the interchange (3) will hold if the sample performance is continuous with respect to the parameter. For the Lindley equation, although T_{n+1} in (5) has a "kink" at $T_n = A_{n+1}$, it is still continuous at that point. This intuitively explains why IPA works. Unfortunately, the "kink" means that the derivative given by (6) has a discontinuity at $T_n = A_{n+1}$, so that IPA will fail for the second derivative.

For the FCFS single-server queue, SPA works nicely for the second derivative of mean system time, resulting in the following estimator:

$$\left(\frac{d^2\overline{T}_N}{d\theta^2}\right)_{SPA} = \frac{1}{N}\sum_{m=1}^{M}\sum_{i=n_{m-1}+1}^{n_m}\sum_{j=n_{m-1}+1}^{i}\frac{d^2X_j}{d\theta^2} + \frac{1}{M}\sum_{m=1}^{M}\frac{f_1(T_{n_m})}{1-F_1(T_{n_m})}\left(\sum_{i=n_{m-1}+1}^{n_m}\frac{dX_i}{d\theta}\right)^2,$$

where $d^2X/d\theta^2$ is well-defined when $F_2(X;\theta)$ is twice differentiable.

Inventory Example: We now illustrate IPA and SPA for a single-item periodic review (s,S) inventory system, in which once every period the inventory level is reviewed and, if necessary, orders are placed to replenish depleted inventory. An (s,S) ordering policy specifies that an order be placed when the level of inventory on hand plus that on order (called *inventory position*) falls below the level s, and that the amount of the order be the difference between S and the present inventory position, i.e., order amounts are placed "up to S." With average inventory as our performance measure of interest, we present gradient estimates with respect to the policy parameters s and $q = S - s$. Note that the parameters in this example are *structural*, as opposed to *distributional* in the previous queueing example.

We will consider the model where all excess demand is backlogged and eventually filled, and orders are immediately received (zero lead time), so that inventory level and inventory position coincide. We assume that at the end of a period, demand is satisfied *before* the order placement decision is made. Let D_n be the demand in period n (i.i.d. with PDF f and CDF F), and V_n be the inventory level in period n after demand satisfaction. This quantity satisfies a recursive equation somewhat analogous to the Lindley equation:

$$V_{n+1} = \begin{cases} V_n - D_{n+1} & \text{if } V_n \geq s, \\ S - D_{n+1} & \text{if } V_n < s. \end{cases} \qquad (8)$$

Our sample performance is the average inventory level over N periods given by $\overline{V}_N = \frac{1}{N}\sum_{n=1}^{N} V_n$.

From a sample path point of view, the key discrete event in the system is the ordering decision each period. A change in s, with q held fixed, has no effect on these decisions, so infinitesimal perturbations in s result in infinitesimal changes in the inventory level, and hence in the sample performance function $\overline{V}_N$. In particular, for a perturbation of size Δs (of any size, not necessarily infinitesimal), $V_n(s+\Delta s) = V_n(s) + \Delta s$, and hence $\partial\overline{V}_N/\partial s = 1$ is an unbiased estimator for $\partial E[\overline{V}_N]/\partial s$. Intuitively, the shape of sample paths are unaltered by changes in s if q is held constant; the entire sample path is merely shifted by the size of the change. The IPA estimator can also be obtained by simply differentiating the recursive relationship (8), noting that D_n does not depend on s or q:

$$\frac{dV_{n+1}}{d\theta} = \begin{cases} \dfrac{dV_n}{d\theta} & \text{if } V_n \geq s, \\ 1 & \text{if } V_n < s. \end{cases} \qquad (9)$$

for either $\theta = s$ or $\theta = q$. Under the assumption that $Y_1 = S = s + q$, the expression reduces to 1 for all n, which is intuitively obvious from the above discussion.

On the other hand, a change in q with s held fixed may cause a change in the set of ordering decisions, resulting in radical changes in the sample path and hence in the sample performance function $\overline{V}_N$. Thus, SPA is required to derive an unbiased gradient estimator for $\theta = q$. An SPA estimator for $\partial E[\overline{V}_N]/\partial q$

that can be easily and efficiently estimated from the original sample path is given by

$$1+\frac{1}{N}\sum_{n\leq N:V_n<s}\frac{f(V_n+D_n-s)}{1-F(V_n+D_n-s)}\left[s-E[D]-\overline{V}_N\right].$$

Historical Notes: PA was "invented" by Ho, Eyler, and Chien (1979), when the first author was consulting on a real-world buffer design problem for a Fiat Motor Company serial production line. The single-server queue example was first considered in Suri and Zazanis (1988), and the inventory example in Fu (1994). Aside from queueing and inventory, PA has also been applied to PERT networks, dams, insurance, maintenance, finance, and statistical process control (see Ho and Cao 1991, and Fu and Hu 1997 for examples and references).

See **Simulation optimization; Score function; Simulation of discrete-event stochastic systems; Monte Carlo sampling and variance reduction; Sensitivity analysis.**

References

[1] Brémaud, P. and F.J. Vázquez-Abad (1992). "On the Pathwise Computation of Derivatives with Respect to the Rate of a Point Process: The Phantom RPA Method," *Queueing Systems: Theory and Applications*, **10**, 249–270.

[2] Cao, X.R. (1994). *Realization Probabilities: The Dynamics of Queueing Systems*, Springer Lecture Notes in Control and Optimization, **194**, Springer-Verlag, New York.

[3] Cassandras, C.G. (1993). *Discrete Event Systems: Modeling and Performance Analysis*, Irwin.

[4] Cassandras, C.G., and S.G. Strickland, S.G. (1989). "On-Line Sensitivity Analysis of Markov Chains," *IEEE Trans. Automatic Control*, **AC-34**, 76-86.

[5] Dai, L. and Y.C. Ho, Y.C. (1995). "Structural Infinitesimal Perturbation Analysis for Derivative Estimation in Discrete Event Dynamic Systems," *IEEE Trans. Automatic Control*, **40**, 1154-1166.

[6] Fu, M.C. (1994). "Sample Path Derivatives for (s, S) Inventory Systems," *Operations Research*, **42**, 351-364.

[7] Fu, M.C. and J.Q. Hu (1997). *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*, Kluwer Academic.

[8] Gaivoronski, A., L.Y. Shi and R.S. Sreenivas (1992). "Augmented Infinitesimal Perturbation Analysis: An Alternate Explanation," *Discrete Event Dynamic Systems: Theory and Applications*, **2**, 121-138.

[9] Glasserman, P. (1991). *Gradient Estimation Via Perturbation Analysis*, Kluwer Academic.

[10] Gong, W.B. and Y.C. Ho (1987). "Smoothed Perturbation Analysis of Discrete-Event Dynamic Systems," *IEEE Trans. Automatic Control*, **AC-32**, 858-867.

[11] Ho, Y.C. and X.R. Cao (1991). *Perturbation Analysis of Discrete Event Dynamic Systems*, Kluwer Academic.

[12] Ho, Y.C., X.R. Cao and C.G. Cassandras (1983). "Infinitesimal and Finite Perturbation Analysis for Queueing Networks," *Automatica*, **19**, 439-445.

[13] Ho, Y.C., M.A. Eyler and T.T. Chien (1979). "A Gradient Technique for General Buffer Storage Design in a Serial Production Line," *International Journal of Production Research*, **17**, 557-580.

[14] Ho, Y.C., and S. Li (1988). "Extensions of Infinitesimal Perturbation Analysis," *IEEE Trans. Automatic Control*, **AC-33**, 827-838.

[15] Shi, L.Y. (1996). "Discontinuous Perturbation Analysis of Discrete Event Dynamic Systems," *IEEE Trans. Automatic Control*, **41**, 1676-1681.

[16] Suri, R. and M.A. Zazanis, M.A. (1988). "Perturbation Analysis Gives Strongly Consistent Sensitivity Estimates for the $M/G/1$ Queue," *Management Science*, **34**, 39-64.

[17] Vakili, P. (1991). "Using a Standard Clock Technique for Efficient Simulation," *Operations Research Letters*, 445-452.

Note: should probably have an entry
INFINITESIMAL PERTURBATION ANALYSIS
that refers to this entry.