An Analysis of Gradient Estimates in Stochastic Network Optimization Problems

Nikolai Krivulin

Two classes of stochastic networks and their performance measures are considered. These performance measures are defined as the expected value of some random variables and cannot normally be obtained analytically as functions of network parameters. We give similar representations for the random variables to provide a useful way of analytical study of these functions and their gradients. The representations are used to obtain sufficient conditions for the gradient estimates to be unbiased. The conditions are rather general and usually met in simulation of the stochastic networks. Applications of the results are discussed and some practical algorithms of calculating unbiased estimates of the gradients are also presented.

1 Introduction

Stochastic network models are widely used in modern engineering, management, biology *etc* to investigate real systems. These models are usually so complicated that can hardly be studied with the help of the analytical methods only. A more fruitful way is to use computer simulation to analyze the networks [1,2,3]. By performing simulation experiments one may get a great amount of information about the network behaviour.

Usually, the main aim of the analysis is to improve a network performance. In order to optimize a performance criterion with respect to network parameters one needs to evaluate it. Simulation provides estimating the criterion as well as its sensitivity (or its gradient, when the parameters are continuous) in a simple way. It is not difficult to obtain estimates provided there exists a simulation model, however each simulation experiment may be very time consumming. Therefore, it is of great importance to develop efficient methods of simulation and estimation.

There are many stochastic optimization procedures which use the data obtained by simulation (see [1] and also a short survey in [4]). In many cases, the procedures that exploit gradient are prefered to those using estimates of the objective function only. The stochastic algorithms which apply unbiased estimates of gradient are often highly efficient. As an example, one can compare the Robbins–Monro procedure with the Kiefer–Wolfowitz one. It is well known [4] that the first procedure based on the unbiased estimates of gradient converges to the solution faster than the second one which approximates the gradient by the finite differences.

In this paper we analyse the problem of unbiased estimation of the gradient of stochastic network performance measures. The paper is based on the results of [5,6]. In Section 1 we describe two classes of stochastic networks and give some examples. We show that the sample performance functions of the networks of both classes may be represented in a similar way. In

fact, these functions are expressed through given ones by using the operations of maximum, minimum and addition.

Section 2 includes a technical result which provides a general representation for the sample performance functions of the networks.

In Section 3 we briefly discuss three methods of estimating gradients, based on simulation data.

The main results are presented in Section 4. Firstly, we introduce a set of functions for which one may obtain unbiased estimates of their gradients. We prove some technical lemmae to state properties of the set. In conclusion, we give the conditions that provide the gradient estimates to be unbiased. These conditions are rather general and usually fulfilled in simulation studies of the stochastic networks.

In Section 5 we show how the results may be applied in practice. Some algoritms of calculating the gradient estimates are described.

2 Stochastic networks and related optimization problems

In this section we present two classes of stochastic networks and discuss optimization problems related to the networks. The performance criterion of the network is normally defined as the expected value of a random variable, $f(\theta,\omega)$, *ie*

$$F(\theta) = E_\omega[f(\theta,\omega)] = E[f(\theta,\omega)],$$

where $\theta \in \Theta \subset R^n$ is a set of decision parameters and ω is a random vector representing the randomness of network behaviour. As a function of the parameters, $f(\theta,\omega)$ is often called sample performance function.

The problem is to optimize the performance measure $F(\theta)$ with respect to $\theta \in \Theta$. In practical problems it is very hard to evaluate the expectation analytically in closed form, even if there is an analytical formula available for $f(\theta,\omega)$. However, it is not difficult to obtain the value of $f(\theta,\omega)$ for any fixed $\theta \in \Theta$ and any realization of ω by using simulation. In that case, one normally use the Monte Carlo approach to estimate the objective function $F(\theta)$ or its gradient.

The main purpose of this section is to show that for many optimization problems, $f(\theta,\omega)$ may be represented in similar algebraic forms. In other words, $f(\theta,\omega)$ is expressed in terms of some given random variables by means of the operations *max*, *min* and $+$. This representation offers the potential for analytical study of the estimates of performance measure gradient. It also provides a theoretical background for efficient algorithms of calculating the estimates.

Activity network. We begin with stochastic activity network models widely used in corporate management in the scheduling of large projects. Consider a project consisting of some activities (or jobs) which must be done to complete it. Each activity is presumed to require a random amount of time for performing it. It is not permitted to begin each activity until some others ones preliminary to it have been completed. One is normally interested in reducing the expected completion time of the whole project.

In order to describe the project as a network, we define an oriented graph $(\mathbf{N},\mathbf{A})$, where $\mathbf{N}$ is the set of nodes and $\mathbf{A}$ is the set of arcs. Each node $i \in \mathbf{N}$ represents the corresponding

activity of the project. For some $i, j \in \mathbf{N}$, the arc (i,j) belongs to $\mathbf{A}$ if and only if the ith activity must precede the jth one directly.

To simplify further formulae we define the set of the father nodes as $\mathbf{N}_F(i) = \{j \in \mathbf{N} | (j,i) \in \mathbf{A}\}$, and the set of the daughter nodes as $\mathbf{N}_D(i) = \{j \in \mathbf{N} | (i,j) \in \mathbf{A}\}$ for every $i \in \mathbf{N}$. In addition, we introduce the set of starting nodes $\mathbf{N}_S = \{i \in \mathbf{N} | \mathbf{N}_F(i) = \emptyset\}$ and the set of the end nodes $\mathbf{N}_E = \{i \in \mathbf{N} | \mathbf{N}_D(i) = \emptyset\}$ of the graph.

Now we have to define the duration of the activities, so that the network would be described completely. Denote the duration of the ith activity by $\tau_i, i \in \mathbf{N}$. We assume τ_i to be a positive random variable, such that $\tau_i = \tau_i(\theta, \omega)$, where $\theta \in \Theta$ is a set of decision parameters and ω is a random vector which represents the random effects involved in realizing the project. The set $\mathbf{T} = \{\tau_i | i \in \mathbf{N}\}$ is presumed to be given.

The sample completion time of the ith activity may be expressed in the form

$$t_i(\theta,\omega) = \begin{cases} \max_{j \in \mathbf{N}_F(i)} t_j(\theta,\omega) + \tau_i(\theta,\omega) & \text{if } i \notin \mathbf{N}_S \\ \tau_i(\theta,\omega) & \text{if } i \in \mathbf{N}_S \end{cases} \tag{1}$$

For the sample completion time of the whole project, we have $t(\theta, w) = \max_{i \in \mathbf{N}_E} t_i(\theta, \omega)$.

In that case, the expected completion time is $T(\theta) = E[t(\theta,\omega)]$, and we wish to minimize $T(\theta)$ with respect to $\theta \in \Theta$.

It is easy to see from (1) that one can represent t as a function of $\tau \in \mathbf{T}$ by using the operations *max* and $+$. To illustrate, consider the simple network depicted in Figure 1.

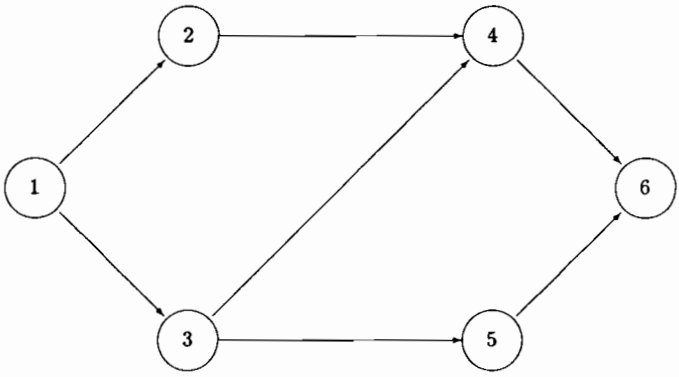

Figure 1. An activity network

For this network, applying (1) successively, we may write the sample completion time as

$$t = \tau_1 + \max\{\max\{\tau_2, \tau_3\} + \tau_4, \tau_3 + \tau_5\} + \tau_6.$$

We will exploit the possibility of t being expressed in such a form in the discussion below.

We conclude this example with the remark about the main difficulty of the activity network optimization problem. It is easy to understand that in the case of general random variables $\tau \in \mathbf{T}$ it is usually very difficult or even impossible to obtain the expected completion time analitically, even if the network is as simple as that in Figure 1. To apply an

optimization procedure in this situation one normally estimate this function or its gradient by using the Monte Carlo approach. Notice, however, that the simulation models of such networks are generally rather simple.

Reliability network. Another class of stochastic network models arises from the reliability investigation of complex interconnected systems in engineering, millitary research, biology etc. Consider a system of elements having bounded random lifetimes. Each element keeps in order until either this element has failed or all those supplying it directly have lost their working conditions. The whole system is presumed to be in order if at least one of the main elements that are supplied by some others but do not supply any element keeps working. An important problem in analyzing this system is to maximize its expected lifetime.

Let $(\mathbf{N}, \mathbf{A})$ be the directed graph discribing the relations between the system elements. In the graph the set of nodes $\mathbf{N}$ corresponds to the set of system elements. If for some $i, j \in \mathbf{N}$, the ith element supplies the jth one directly, then $(i,j) \in \mathbf{A}$. For the graph we retain the notations $\mathbf{N}_F(i), \mathbf{N}_D(i), \mathbf{N}_S$ and $\mathbf{N}_E$ introduced above.

For every element $i \in \mathbf{N}$, we define the lifetime as the random variable $\tau_i(\theta, \omega)$ which depends on the set of decision parameters $\theta \in \Theta$. Assume the set $\mathbf{T} = \{\tau_i\}$ to be given. Now, we may represent the time for the ith element to be in order as

$$t_i(\theta,\omega) = \begin{cases} \min\{\max_{j \in \mathbf{N}_F(i)} t_j(\theta,\omega), \tau_i(\theta,\omega)\} & \text{if } i \notin \mathbf{N}_S \\ \tau_i(\theta,\omega) & \text{if } i \in \mathbf{N}_S \end{cases} \tag{2}$$

The sample and expected lifetimes of the whole system may be written as

$$t(\theta,\omega) = \max_{i \in \mathbf{N}_E} t_i(\theta,\omega) \quad \text{and} \quad T(\theta) = E[t(\theta,\omega)],$$

respectively.

To illustrate this realibility network model consider that depicted in Figure 2 (Ermakov, [1]).

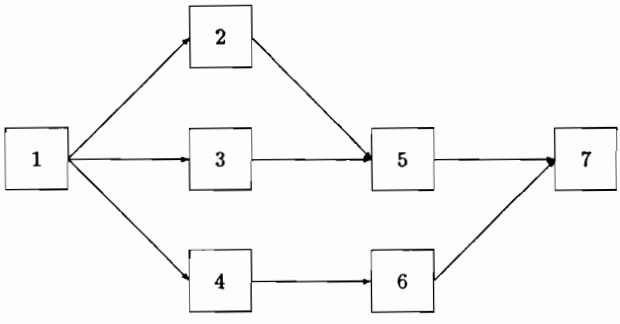

Figure 2. A reliability network

For the sample lifetime of the system, we have from (2)

$$t = \min\{\tau_1, \max\{\min\{\tau_4, \tau_6\}, \min\{\max\{\tau_2, \tau_3\}, \tau_5\}\}, \tau_7\}.$$

We can see that the sample lifetime of such a network has one important property: it may be represented as a function of all $\tau \in \mathbf{T}$ by using only the operations *max* and *min*. Note that the difficulties in solving the problem of expected lifetime maximization are the same as in activity network optimization.

3 An algebraic representation lemma

We have seen that the functions of network performance possesses some algebraic properties. The point is that they may be expressed as a function of given random variables by means of the operations *max*, *min* and $+$. For the networks, one can obtain such representations from recursive equations (1) and (2). In this section we present a general form of the representations, which provides a common way of examining analytical properties of the performance functions in both networks.

In order to simplify futher formulae we introduce the notations $\vee$ for maximum and $\wedge$ for minimum. In addition, we will use the sign $\bigvee$ ($\bigwedge$) to represent an iterated maximum (minimum), *ie*

$$\bigvee_{i=1}^{n} x_i = x_1 \vee \ldots \vee x_n \quad (\bigwedge_{i=1}^{n} x_i = x_1 \wedge \ldots \wedge x_n)$$

Let $\mathbf{X}$ be a set supplied with the operations $+$, $\vee$ and $\wedge$. Without loss of generality we may consider $\mathbf{X}$ to be a set of real numbers. It is easy to extend the result of this section to various sets of real-valued functions and random variables. We assume that the traditional algebraic axioms are fulfilled in $\mathbf{X}$. In particular, we will use the following axioms.

Axiom 1. Distributivity of maximum over minimum.

$\forall x, y, z \in \mathbf{X}$, $(x \wedge y) \vee z = (x \vee z) \wedge (y \vee z)$.

Axiom 2. Distributivity of minimum over maximum.

$\forall x, y, z \in \mathbf{X}$, $(x \vee y) \wedge z = (x \wedge z) \vee (y \wedge z)$.

Axiom 3. Distributivity of sum over maximum and minimum.

$\forall x, y, z \in \mathbf{X}$, $(x \vee y) + z = (x + z) \vee (y + z)$, $(x \wedge y) + z = (x + z) \wedge (y + z)$.

The general form of the representation is determined in the next technical lemma.

Lemma 1. *Let $\varphi(z_1, \ldots, z_p)$ be a function of the variables $z_1, \ldots, z_p$ taking their values in $\mathbf{X}$, φ is defined as a composition of the operations $\vee$, $\wedge$ and $+$. Then φ can be represented as*

$$\varphi(z_1, \ldots, z_p) = \bigvee_{i \in I} \bigwedge_{j \in J_i} \sum_{k=1}^{p} \alpha_{ij}^k z_k,$$

where I and J_i for all $i \in I$ are finite sets of indices, and all α_{ij}^k are integers.

Proof. Without loss of generality we suppose that there is no more than one entry of each variable $z_1, \ldots, z_p$ into the expression. If some variable has two or more entries, we introduce additional ones so that the above presupposition would be fulfilled. Let us prove the lemma by induction on the number of variables.

For $p = 1$, the statement of the lemma is obvious. If $p = 2$, there are three possibilities

$$z_1 \vee z_2, \quad z_1 \wedge z_2 \quad \text{and} \quad z_1 + z_2,$$

and it is clear that the statement is also true.

Assume that the statement of the lemma is true up to some value $p - 1$. Consider an expression φ of p variables. Clearly, there is an operation in the expression that should be performed after the other ones. Denote this operation by the asterisk $*$. In this case, we have $\varphi = \varphi_1 * \varphi_2$, where φ_1 and φ_2 are expressions such that each of them cannot include all the variables $z_1, \ldots, z_p$. By the assumption, the statement of the lemma holds for both φ_1 and φ_2. Now, we have three possibilities for the operation $*$.

1. $\vee$. This is obvious.
2. $\wedge$. It is sufficient to apply Axiom 1.
3. $+$. To obtain the representation in this case, one has to apply successively Axioms 1, 2 and 3.

Consequently, the statement of the lemma is true for $\varphi = \varphi_1 * \varphi_2$. □

4 Estimates of gradient

To optimize the network performance measure $F(\theta) = E[f(\theta, \omega)]$ one often needs information about the gradients $\partial F(\theta)/\partial\theta$. In the absence of analytical formulae for the gradient, Monte Carlo experiments may be performed to estimate its values. There are three general methods of estimating $\partial F(\theta)/\partial\theta$ based on data obtained by simulation [1,3,7]. In the first two methods the gradient is approximated by the finite differences and then estimated by using the Monte Carlo approach. To illustrate these two methods, assume θ to be a scalar and consider the following estimates:

The crude Monte Carlo (CMC) estimate:

$$G_{CMC} = \frac{1}{N\Delta\theta} \sum_{i=1}^{N} (f(\theta + \Delta\theta, \omega_i) - f(\theta, \omega_{N+i}))$$

The common random number (CRN) estimate:

$$G_{CRN} = \frac{1}{N\Delta\theta} \sum_{i=1}^{N} (f(\theta + \Delta\theta, \omega_i) - f(\theta, \omega_i))$$

where $\omega_i, i = 1, \ldots, 2N$ are independent realizations of the random vector ω. The second estimate differs from the first in one respect: in the CRN estimate the random variables ω_i are the same for both $\theta + \Delta\theta$ and θ, whereas in the CMC estimate they are different. Note that each of them requires $2 \times N$ simulation runs (N at the original value θ and N at $\theta + \Delta\theta$). Clearly, in the case of the vector $\theta \in R^n$, one must perform $(n + 1) \times N$ simulation experiments to get each estimate. In [1, pp. 153-154] Ermakov has shown that the finite difference estimates have the mean square error (MSE) which is of order $O(N^{-1/3})$ for G_{CRN} and $O(N^{-1/2})$ for G_{CRN}.

We may somewhat improve the MSE properties of the estimate by using more sophisticated difference formulae, however, the estimates become very expensive in terms of computer time because they require a large number of additional simulation experiments. For example, the following symmetric difference estimate

$$G_{CRN}^{SD} = \frac{1}{2N\Delta\theta}\sum_{i=1}^{N}(f(\theta+\Delta\theta,\omega_i) - f(\theta-\Delta\theta,\omega_i))$$

requires $2 \times N$ simulation runs ($2 \times n \times N$, when $\theta \in R^n$).

An estimate of the third method can be written in the form

$$G = \frac{1}{N}\sum_{i=1}^{N}\frac{\partial}{\partial\theta}f(\theta,\omega_i), \tag{3}$$

provided that the gradient of the sample performance function (sample gradient) exists. It should be noted that, although we may obtain values of the sample performance function by simulation, it can be rather difficult to evaluate its gradient.

Recently, a new technique, called infinitesimal perturbation analysis (IPA), has been developed (Ho *et al.* [2]) as an efficient method of obtaining gradient information. The IPA method yields the exact values of the sample gradient $\partial f(\theta,\omega)/\partial\theta$ by performing one simulation run. The method is based on the analysis of the dynamics of the network and closely connected with the simulation technique. Therefore, one can easily combine an IPA procedure for calculating the sample gradient with a suitable algorithm of network simulation. Such a procedure provides all the partial derivatives of the sample gradient simultaneously during one simulation run. Furthermore, it needs an additional computation cost which is usually very small compared with that required for the simulation run alone.

The key question concerning the IPA method is whether it produces an unbiased estimate of the performance measure gradient. It can easily be shown that if $\partial f(\theta,\omega)/\partial\theta$ is an unbiased estimator of $\partial F(\theta)/\partial\theta$ then estimate (3) has MSE which is of order $O(N^{-1})$. In short, in the case of unbiasedness, this is a very efficient estimate, that provides considerable savings in computation.

In the next section, using the algebraic representation of Section 3, we will examine properties of the network performance functions so as to derive the conditions for estimate (3) of the performance measure gradient to be unbiased.

5 A theoretical background of unbiased estimation

A sufficient condition for the estimate (3) of $\partial E[f(\theta,\omega)]/\partial\theta$ at some $\theta \in \Theta$ to be unbiased is

$$\frac{\partial}{\partial\theta}E[f(\theta,\omega)] = E[\frac{\partial}{\partial\theta}f(\theta,\omega)]. \tag{4}$$

Cao in [7] showed that (4) holds in the case of $f(\theta,\omega)$ being uniformly differentiable at θ w.p.1. Note that such a differentiability property is not easy to verify and hard to interpret for practical systems. A useful way to prove the interchange in (4) is to apply the Lebesgue dominated convergence theorem (Loève [8]). We use this theorem in the following form.

Theorem 2. *Let $(\Omega, \mathcal{F}, P)$ be a probability space, $\Theta \subset R^n$ and $f : \Theta \times \Omega \longrightarrow R$ be a $\mathcal{F}$-measurable function for any $\theta \in \Theta$ and such that the following conditions hold:*

(i) *for every $\theta \in \Theta$, there exists $\partial f(\theta, \omega)/\partial\theta$ at θ w.p.1,*

(ii) *for all $\theta_1, \theta_2 \in \Theta$, there is a random variable $\lambda(\omega)$ defined on the same probability space, with $E\lambda < \infty$ and such that*

$$|f(\theta_1, \omega) - f(\theta_2, \omega)| \leq \lambda(\omega)\|\theta_1 - \theta_2\| \quad w.p.1. \tag{5}$$

Then equation (4) holds on Θ.

As an important consequence, we may state that the function $F(\theta) = E[f(\theta, \omega)]$ is a Lipschitz one with a constant $L = E\lambda$ and continuously differentiable on Θ, provided f satisfies the theorem conditions.

Definition 1. A function $f(\theta, \omega)$ defined on the probability space $(\Omega, \mathcal{F}, P)$ at every $\theta \in \Theta$ belongs to the set $\mathcal{D}_{\Theta,\Omega}$ (or simply $\mathcal{D}$) if and only if it satisfies the conditions of Theorem 2.

Example 1. Random variables which arise from simulation study of networks, can be treated as members of a family of random variables [1]. There are few families one usually applies, namely the Exponential family, the Gaussian family *etc.* Various random variables of a family may be obtained from the standard variable by using a suitable transformation. An ordinary way to transform random variables is based on changing location and scale parameters.

Let $\xi(\omega)$ be the standard random variable of a family. Define

$$f(\theta, \omega) = \theta_1 \xi(\omega) + \theta_2,$$

where $\theta = (\theta_1, \theta_2)^\top \in \Theta \subset R^2$. Let us check whether it holds that $f \in \mathcal{D}$. Obviously, the partial derivatives of f with respect to θ_1 and θ_2 exist for almost all ω and equal

$$\frac{\partial}{\partial\theta_1} f(\theta, \omega) = \xi(\omega) \quad \text{and} \quad \frac{\partial}{\partial\theta_2} f(\theta, \omega) = 1.$$

In addition, it is easy to verify that f satisfies Condition (ii) of Theorem 2 with $\lambda = |\xi| + 1$. If $E|\xi| < \infty$, as is usually the case, then the conditions of Theorem 2 are fulfilled for f and we have $f \in \mathcal{D}$.

The next technical lemmae give the sufficient conditions for the arithmetic operations and the operation *max* and *min* not to break the main properties of the functions from $\mathcal{D}$.

Lemma 3. *Let $f, g \in \mathcal{D}$ and let λ_1 and λ_2 be the random variables that provide Condition (ii) of Theorem 2 for f and g, respectively. Let μ_1, μ_2 and ν be positive random variables. Then the following are satisfied.*

(i) $f + g \in \mathcal{D}$.

(ii) *If α is a bounded random variable, then $\alpha f \in \mathcal{D}$.*

(iii) *If $|f| \leq \mu_1$ and $|g| \leq \mu_2$ hold w.p.1 for any $\theta \in \Theta$ and $E[\lambda_1\mu_2 + \lambda_2\mu_1] < \infty$, then $fg \in \mathcal{D}$.*

(iv) *If $|f| \leq \mu_1$ and $|g| \geq \nu$ hold w.p.1 for any $\theta \in \Theta$ and $E[\frac{\mu_1\lambda_2}{\nu^2} + \frac{\lambda_1}{\nu}] < \infty$, then $\frac{f}{g} \in \mathcal{D}$.*

Proof. Clearly, $f+g$, αf, fg and f/g are measurable functoins of ω and differentiable ones on Θ w.p.1. Since for all of these functions the proofs of inequality (5) are quite similar, we verify it for one of them only. For instance, we examine $h = fg$.

For all $\theta_1, \theta_2 \in \Theta$ we have

$$\begin{aligned} |h(\theta_1,\omega) - h(\theta_2,\omega)| &= |f(\theta_1,\omega)g(\theta_1,\omega) - f(\theta_2,\omega)g(\theta_2,\omega)| = \\ |f(\theta_1,\omega)g(\theta_1,\omega) - f(\theta_2,\omega)g(\theta_1,\omega) + f(\theta_2,\omega)g(\theta_1,\omega) - f(\theta_2,\omega)g(\theta_2,\omega)| &\le \\ |g(\theta_1,\omega)||f(\theta_1,\omega) - f(\theta_2,\omega)| + |f(\theta_2,\omega)||g(\theta_1,\omega) - g(\theta_2,\omega)| &\le \\ (\lambda_1(\omega)\mu_2(\omega) + \lambda_2(\omega)\mu_1(\omega))\|\theta_1 - \theta_2\| \quad \text{w.p.1.} \end{aligned}$$

In short, $|h(\theta_1,\omega) - h(\theta_2,\omega)| \le \lambda(\omega)\|\theta_1 - \theta_2\|$ w.p.1, where $\lambda = \lambda_1\mu_2 + \lambda_2\mu_1$, $E\lambda = E[\lambda_1\mu_2 + \lambda_2\mu_1] < \infty$. By Theorem 2, we conclude $fg \in \mathcal{D}$. □

Notice, from Lemma 3 (i) and (ii) it follows that being closed for the operations of addition and multiplication by bounded random variables, $\mathcal{D}$ is a linear space of functions with these two operations.

Lemma 4. *Let $f, g \in \mathcal{D}$. Suppose that for any $\theta_0 \in \theta$, there exists a neighbourhood $\mathbf{U}_\omega(\theta_0)$ of θ_0 w.p.1 such that one and only one of the following conditions*

(i) $f(\theta,\omega) = g(\theta,\omega)$,

(ii) $f(\theta,\omega) < g(\theta,\omega)$,

(iii) $f(\theta,\omega) > g(\theta,\omega)$

is satisfied for all $\theta \in \mathbf{U}_\omega(\theta_0)$. Then $f \vee g \in \mathcal{D}$ and $f \wedge g \in \mathcal{D}$.

Proof. Consider $h(\theta,\omega) = f(\theta,\omega) \vee g(\theta,\omega)$. It is clear that h is measurable with respect to ω. In order to prove differentiability of h w.p.1 on Θ, we examine an arbitrary $\theta \in \Theta$. There are only two possibility for h not to be differentiable. Firstly, it is possible that the derivative of h at θ does not exist if at least one of the derivatives $\partial f(\theta,\omega)/\partial\theta|_{\theta=\theta_0}$ and $\partial g(\theta,\omega)/\partial\theta|_{\theta=\theta_0}$ does not. In addition, h may not be differentiable at θ if the maximum of the functions f and g changes over from f to g at this point or vice versa. The last case is equivalent to that there exists $\omega \in \Omega$ such that all the neighbourhoods $\mathbf{U}_\omega(\theta_0) \subset \Theta$ contain both points at which $f(\theta,\omega) = g(\theta,\omega)$ and $f(\theta,\omega) \neq g(\theta,\omega)$. By the assumption of the lemma, both of these cases may occur only with zero probability. Therefore, there exists $\partial h(\theta,\omega)/\partial\theta|_{\theta=\theta_0}$ at all $\theta \in \Theta$ w.p.1.

For the function h, the proof will be completed if we show that h satisfies Condition (ii) of Theorem 2. Since $f, g \in \mathcal{D}$, there are random variables λ_1 and λ_2 with $E\lambda_1 < \infty$ and $E\lambda_2 < \infty$ such that the inequalities

$$\begin{aligned} |f(\theta_1,\omega) - f(\theta_2,\omega)| &\le \lambda_1(\omega)\|\theta_1 - \theta_2\| \quad \text{w.p.1} \\ |g(\theta_1,\omega) - g(\theta_2,\omega)| &\le \lambda_2(\omega)\|\theta_1 - \theta_2\| \quad \text{w.p.1} \end{aligned}$$

hold for all $\theta_1, \theta_2 \in \Theta$. Let ω be an arbitrary element of Ω at which both these inequalities hold. Devide Θ into two subsets:

$$\begin{aligned} \mathbf{X}_\omega &= \{\theta \in \Theta | f(\theta,\omega) \ge g(\theta,\omega)\}, \\ \mathbf{Y}_\omega &= \{\theta \in \Theta | f(\theta,\omega) < g(\theta,\omega)\}. \end{aligned}$$

Obviously, it holds

$$|h(\theta_1,\omega) - h(\theta_2,\omega)| \le \lambda_1(\omega)\|\theta_1 - \theta_2\|$$

for all $\theta_1, \theta_2 \in \mathbf{X}_\omega$ and

$$|h(\theta_1,\omega) - h(\theta_2,\omega)| \le \lambda_2(\omega)\|\theta_1 - \theta_2\|$$

for all $\theta_1, \theta_2 \in \mathbf{Y}_\omega$. Assume $\theta_1 \in \mathbf{X}_\omega, \theta_2 \in \mathbf{Y}_\omega$. If $h(\theta_1,\omega) \ge h(\theta_2,\omega)$, we deduce

$$|h(\theta_1,\omega) - h(\theta_2,\omega)| = |f(\theta_1,\omega) - g(\theta_2,\omega)| < |f(\theta_1,\omega) - f(\theta_2,\omega)| \le \lambda_1(\omega)\|\theta_1 - \theta_2\|.$$

Similarly, if $h(\theta_1,\omega) < h(\theta_2,\omega)$, we have $|h(\theta_1,\omega) - h(\theta_2,\omega)| \le \lambda_2(\omega)\|\theta_1 - \theta_2\|$. It follows that $|h(\theta_1,\omega) - h(\theta_2,\omega)| \le \lambda(\omega)\|\theta_1 - \theta_2\|$, $\lambda(\omega) = \lambda_1(\omega) \vee \lambda_2(\omega)$, for all $\theta_1, \theta_2 \in \Theta$. Since this inequality holds for almost all $\omega \in \Omega$, we conclude that

$$|h(\theta_1,\omega) - h(\theta_2,\omega)| \le \lambda(\omega)\|\theta_1 - \theta_2\| \quad \text{w.p.1},$$

and $E\lambda = E[\lambda_1 \vee \lambda_2] \le E\lambda_1 + E\lambda_2 < \infty$.

In other words, h satisfies the conditions of Theorem 2. Consequently, $f \vee g \in \mathcal{D}$. The proof of the statement $f \wedge g \in \mathcal{D}$, is analogous. □

It should be noted that the condition of Lemma 4 is not necessary, as the next example shows.

Example 2. Let $\Theta = [-1,1]$, $(\Omega, \mathcal{F}, P)$ be a probability space, where $\Omega = [0,1]$, $\mathcal{F}$ is the σ–field of Borel sets of Ω and P is the Lebesgue measure on Ω. Consider the following functions:

$$f(\theta,\omega) = -\theta^3 + \omega, \quad g(\theta,\omega) = \theta^2 + \omega$$

and

$$h(\theta,\omega) = f(\theta,\omega) \vee g(\theta,\omega) = \begin{cases} -\theta^3 + \omega & \text{if} \quad -1 \le \theta \le 0 \\ \theta^2 + \omega & \text{if} \quad 0 < \theta \le 1 \end{cases}$$

One can easily verify that for any neighbourhood of $\theta = 0$, there exist both points with $f > g$ and $f < g$ w.p.1. The conditions of Lemma 4 are therefore violated. Nevertheless, h is differentiable at 0 for all $\omega \in \Omega$. In addition, it holds that $h \in \mathcal{D}$.

Corollary 5. *Let $f, g \in \mathcal{D}$. If for every $\theta \in \Theta$ it holds that $f \ne g$ w.p.1, then $f \vee g \in \mathcal{D}$ and $f \wedge g \in \mathcal{D}$.*

Proof. Clearly, the condition of the corollary implies that either $f - g > 0$ or $f - g < 0$ holds at every $\theta \in \Theta$ w.p.1. Since $f, g \in \mathcal{D}$, these two functions are continuous ones of θ w.p.1 as well as $f - g$. Because of continuity, $f - g > 0$ $(f - g < 0)$ holds w.p.1 not only at θ, but also at every points of a neighbourhood of θ. It remains to apply Lemma 4. □

Using Corollary 5 we give the following general conditions for $\mathcal{D}$ to provide closeness with respect to the operations $\vee$ and $\wedge$.

Lemma 6. *Let $f, g \in \mathcal{D}$. If for any $\theta \in \Theta$ it holds that the random variables $f(\theta,\omega)$ and $g(\theta,\omega)$*

(i) *are independent, and*

(ii) *at least one of them is continuous*

then $f \vee g \in \mathcal{D}$ and $f \wedge g \in \mathcal{D}$.

To prove the lemma it is sufficient to see that its conditions lead to that of Corollary 5.

The next two examples show that both conditions of Lemma 6 are essential.

Example 3. Let $(\Omega, \mathcal{F}, P)$ and Θ be defined as in Example 2. Also define

$$f(\theta,\omega) = -\theta + \omega \quad \text{and} \quad g(\theta,\omega) = \theta + \omega.$$

Let us consider the function

$$h(\theta,\omega) = f(\theta,\omega) \vee g(\theta,\omega) = \begin{cases} -\theta+\omega & \text{if } -1 \le \theta \le 0 \\ \theta+\omega & \text{if } 0 < \theta \le 1 \end{cases}$$

It is clear that $f, g \in \mathcal{D}$ and for every $\theta \in \Theta$, the random variables $f(\theta,\omega)$ and $g(\theta,\omega)$ are continuous. Although inequality (5) holds with $\lambda = 1$ for h, this function is not differentiable at $\theta = 0$ for all $\omega \in \Omega$. Therefore, $h \notin \mathcal{D}$.

Example 4. Let $\Theta = [0,1]$, $\Omega_1 = \Omega_2 = [0,1]$ and P be the Lebesgue measure on $\Omega = \Omega_1 \times \Omega_2$. Denote $\omega = (\omega_1, \omega_2)^\top$ and consider the following functions:

$$f(\theta,\omega) = \begin{cases} \frac{1}{2}\theta & \text{if } \omega_1 \le \frac{1}{2} \\ 1 & \text{if } \omega_1 > \frac{1}{2} \end{cases}$$

$$g(\theta,\omega) = \begin{cases} \theta^2 & \text{if } \omega_2 \le \frac{1}{2} \\ 1 & \text{if } \omega_2 > \frac{1}{2} \end{cases}$$

and

$$h(\theta,\omega) = f(\theta,\omega) \vee g(\theta,\omega) = \begin{cases} \max\{\frac{1}{2}\theta, \theta^2\} & \text{if } \omega_1 \le \frac{1}{2} \text{ and } \omega_2 \le \frac{1}{2} \\ 1 & \text{otherwise} \end{cases}$$

One can see that $f, g \in \mathcal{D}$ and for every $\theta \in \Theta$, the random variables $f(\theta,\omega)$ and $g(\theta,\omega)$ are independent. In addition, the condition (ii) of Theorem 2 holds for h with $\lambda = 2$. Nevertheless, $h = \max\{\frac{1}{2}\theta, \theta^2\}$ with probability $\frac{1}{4}$, that is not a differentiable function at $\theta = \frac{1}{2}$. In that case, $h \notin \mathcal{D}$.

Lemma 7. *Let $\mathcal{M}$ be a set of functions from $\mathcal{D}$ such that for all $f, g \in \mathcal{M}$, the conditions of Lemma 4 are fulfilled. Then $\mathcal{M}$ is closed for the operations* max *and* min.

Proof. Let $f, g \in \mathcal{M}$ and let us define $h = f \vee g$. Note that $h \in \mathcal{D}$ by Lemma 4. We have to prove the conditions of Lemma 4 to be satisfied for h and any $u \in \mathcal{M}$.

If u is either f or g, say $u \equiv f$, we may write

$$h - u = f \vee g - f = \begin{cases} g - f & \text{if } f < g \\ 0 & \text{if } f \ge g \end{cases}$$

Since $f, g \in \mathcal{M}$, for any point of Θ, there is a neighbourhood on which only one of the conditions $f - g < 0, f - g = 0$ or $f - g > 0$ holds w.p.1. From the above identity this also holds for $h - u$ on the neighbourhood. Consequently, in this case the conditions of Lemma 4 are fulfilled.

Now we assume $u \in \mathcal{M} \setminus \{f, g\}$. We have

$$h - u = f \vee g - u = \begin{cases} g - f & \text{if } f < g \\ f - u & \text{if } f \ge g \end{cases}$$

Let us examine any $\theta \in \Theta$. Suppose that $f < g$ w.p.1 at θ. Since f, g and u belong to $\mathcal{M}$, there are neighborhoods $\mathbf{U}_\omega(\theta)$ and $\mathbf{V}_\omega(\theta)$ where the conditions of Lemma 4 are fulfilled

for each pairs of functons (f,g) and (g,u), respectively. It follows from the expression of h that the neighborhood $\mathbf{U}_\omega \bigcup \mathbf{V}_\omega(\theta)$ is that Lemma 4 requires for h and u. If it holds that $f \geq g$ or $f = g$ at θ, the reasoning is the same.

In short, we have shown that the conditions of Lemma 4 are satisfied for h and any $u \in \mathcal{M}$, and therefore, $h = f \vee g \in \mathcal{M}$. In the case of minimum the proof is analogous. □

Corollary 8. *If $f_j \in \mathcal{M}$ for every $j = 1,\ldots,N$, then it holds*

$$\bigvee_{i\in I} \bigwedge_{j\in J_i} f_j \in \mathcal{M},$$

where I is a finite set of indices and $J_i \subset \{1,\ldots,N\}$ for every $i \in I$.

This is an immediate consequence of the previous lemma.

The next example is of importance to the main result of the section.

Example 5. Let $f_j \in \mathcal{D}$ for all $j = 1,\ldots,N$. Suppose that at every $\theta \in \Theta$, all the random variables $f_j(\theta,\omega)$ are continuous and independent. Define $\mathcal{L}$ to be a set of linear combinations $\sum_{i\in I} a_i f_i$ with integer coefficients a_i, $i \in I \subset \{1,\ldots,N\}$. Obviously, $\mathcal{L}$ is stable for addition. For all functions $u = \sum_{i\in I} a_i f_i$ and $v = \sum_{j\in J} b_j f_j$ we have $u - v = \sum_{k\in K} c_k f_k$. It is clear that for every $\theta \in \Theta$, $u - v$ is a continuous random variable because of the properties of f (except for the case of all $c_k = 0$ which is obvious). Therefore, it holds that $u - v \neq 0$ w.p.1 at every $\theta \in \Theta$. Similarly as in Corollary 5, one can deduce that u and v satisfy the conditions of Lemma 4. From this we conclude that $\mathcal{L}$ may be treated as an example of $\mathcal{M}$.

One can easily see that the condition of continuity is essential to this reasoning. To illustrate the important role of independence, consider the following functions

$$f(\theta,\omega) = -2\theta + 2\omega, \quad g(\theta,\omega) = \theta - \omega \quad \text{and} \quad u(\theta,\omega) = \theta + \omega$$

under the same assumption as in Example 3. It is easy to verify that the conditions of Lemma 4 are fulfilled for any two functions of them. Nevertheless, the functions u and $v = f + g$ do not satisfy them, as Example 3 has shown.

Now, we may formulate the main result of the section. We first introduce some definitions. Let $\mathcal{A}$ be the algebra of all functions $f : \Theta \times \Omega \longrightarrow R$ being defined on the probability space $(\Omega, \mathcal{F}, P)$ at every $\theta \in \Theta$ with the operations $\vee, \wedge$ and $+$. In other words, this is a closed system of the functions for these operations.

Definition 2. Let $\mathbf{T}$ be a finite subset of functions of $\mathcal{A}$. We define $[\mathbf{T}]_{\mathcal{A}}$ to be the set generated by $\mathbf{T}$ in $\mathcal{A}$, that is the set of all functions being obtained from ones of $\mathbf{T}$ by means of the operations $\vee, \wedge$ and $+$.

Theorem 9. *Let $\mathbf{T} \in \mathcal{D}$. Suppose that for all $\tau \in \mathbf{T}$, $\tau(\theta,\omega)$ are continuous and independent random variables at any $\theta \in \Theta$. Then it holds $[\mathbf{T}]_{\mathcal{A}} \subset \mathcal{D}$.*

Proof. It results from Lemma 1 that every $f \in [\mathbf{T}]_{\mathcal{A}}$ can be represented as

$$f = \bigvee_{i\in I} \bigwedge_{j\in J_i} \sum_{\tau\in\mathbf{T}} a_{ij}^{\tau} \tau,$$

where all a_{ij}^{τ} are integers. It has been shown in Example 5 that the functions of the family $\{\sum_{\tau \in \mathbf{T}} a_k^{\tau} \tau\}_{k=1,2,\ldots}$ satisfy the conditions of Lemma 4. Applying Corollary 8, we conclude that the statement of the theorem is true. □

It is important to note that the conditions of Theorem 9 are rather general and usually fulfilled in the network simulation. In particular, in contrast with the traditional approaches (cp, for example, existing results on the unbiasedness of IPA estimates in [2,3]), we may not restrict ourself to the exponential distribution.

In short, to satisfy the theorem only the following are required for the functions of the set $\mathbf{T}$:

(i) for any $\theta \in \Theta$, all $\tau \in \mathbf{T}$ are continuous and independent random variables;

(ii) each $\tau \in \mathbf{T}$ as a function of θ is differentiable w.p.1 and Lipschitz one with an integrable random variable as a Lipschitz constant.

In the next section we will show how these results can be applied to some problems to verify the unbiasedness of gradient estimates.

6 Applications

Now we discuss the applications of the previous results to optimizing the networks. In particular, we describe algorithms of obtaining sample gradients, based on the algebraic representation of the networks. In this section we keep using the notations $(\Omega, \mathcal{F}, P)$ and Θ for the underlying probability space and the parameter space, respectively.

We begin with the stochastic activity network. Let the duration of the jth activity be represented by the function $\tau_j(\theta, \omega)$. Denote the set of all such functions of the network by $\mathbf{T}$. As we have seen, a sample completion time of the network $t(\theta, \omega)$ may be expressed by functions of $\mathbf{T}$ by using only the operations *max* and $+$. This implies $t \in [\mathbf{T}]_{\mathcal{A}}$.

Suppose that $\mathbf{T} \in \mathcal{D}$ and all $\tau \in \mathbf{T}$ are continuous, and they are independent random variables at every $\theta \in \Theta$. For the mean completion time $T(\theta) = E[t(\theta, \omega)]$, it follows from Theorem 2 that $\frac{1}{N} \sum_{i=1}^{N} \partial t(\theta, \omega_i)/\partial\theta$, where $\omega_i \in \Omega$, is an unbiased estimate of the gradient $\partial T(\theta)/\partial\theta$.

As an example, suppose $\tau(\theta, \omega) = -\theta \ln(1-\omega)$, where $\theta \in R$ and the random variable ω is uniformly distrubuted on $[0,1]$. It is well known [1] that $-\ln(1-\omega)$ has an exponential distribution with mean 1. Similarly as in Example 1, we have $\tau \in \mathcal{D}$. In addition, durations of the activities are normally considered as independent in the probabilistic sense. Our result is, therefore, applicatable in this case.

Now, suppose that there is a simulation procedure for the activity network with L nodes to provide a simulation experiment for any fixed $\theta \in \Theta$ and a realization of ω. One can easily combine it with the following algorithm.

Algorithm 1.

Step (i). At the initial time, fix values of θ and ω; set $g_j = 0$ for $j = 1, \ldots, L$, and set $c = 0$.

Step (ii). At the completion of any activity i, add the value of $\partial\tau_i(\theta, \omega)/\partial\theta$ to g_i and add 1 to c;
if $c = L$, then save g_i as the value of $\partial t(\theta, \omega)/\partial\theta$ and stop; otherwise go to Step (iii).

Step (iii). Determine the set $\mathbf{N}_D(i)$. For every $j \in \mathbf{N}_D(i)$, if all activities of the set $\mathbf{N}_F(j)$ have been completed, then set $g_j = g_i$.

To verify the correctness of Algorithm 1 it sufficies to see that it is simply based on recursive equation (1). Note that Algorithm 1 is similar to those based on the IPA method in [3].

For a reliability network, one can apply Theorem 2 in a similar way. As in Section 2, denote the sample lifetime of a system by $t(\theta,\omega)$. It is not difficult to construct the next algorithm that calculates the sample gradient $\partial t(\theta,\omega)/\partial\theta$.

Algorithm 2.

Step (i). At the initial time, fix values of θ and ω.

Step (ii). At the failure of element i, exclude all nodes representing the elements that are now not able to keep working from the set $\mathbf{N}$ as well as the corresponding arcs from the set $\mathbf{A}$.

Step (iii). If for the reduced set $\mathbf{N}$ it holds $\mathbf{N} \cap \mathbf{N}_E = \emptyset$, then save $\partial\tau_i(\theta,\omega)/\partial\theta$ as the value of $\partial t(\theta,\omega)/\partial\theta$ and stop; otherwise go to Step (ii).

In conclusion, note that both algorithms are rather simple. In fact, they only require calculating gradients of given functions and performing some trivial operations to produce values of the sample gradients. Using these values, one can easily estimate the gradients of the system performance measures so as to apply efficient optimization procedures.

References

[1] Ermakov, S.M. (1975). Die Monte-Carlo-Methode und verwandte Fragen. VEB Deutscher Verlag der Wissenschaften.

[2] Ho, Y.C. (1987). Performance evaluation and perturbation analysis of discrete event dynamic systems. IEEE Trans. Automat. Contr., AC-32, no. 7, 563–572.

[3] Suri, R. (1989). Perturbation analysis: The state of the art and research issues explained via the GI/G/1 queue. In Proc. of IEEE, 77, no. 1, 114–137.

[4] Glynn, P.W. (1986). Optimization of stochastic systems. In Proc. of 1986 Winter Simulation Conf., 52–58.

[5] Krivulin, N.K. (1990). On complex system optimization using simulation. Vestnik Leningrad. Univ. Mat. Meh. Astronom., 2(8), 100–102. (in Russian)

[6] Krivulin, N.K. (1990). Optimization of discrete event dynamic systems by using simulation. Ph.D. Dissertation. St Petersburg University. (in Russian)

[7] Cao, X.R. (1985). Convergence of parameter sensitivity estimates in a stochastic experiment. IEEE Trans. Automat. Contr., AC-30, no. 9, 845–853.

[8] Loève, M. (1960). Probability Theory. Van Nostrand.