A CONTROL VARIATE APPROACH FOR EFFICIENT STOCHASTIC FINITE ELEMENT ANALYSIS OF GEOTECHNICAL PROBLEMS

T. S. Charlton,¹ M. Rouainia,² and R.J. Dawson³

ABSTRACT

Monte Carlo simulation is the most versatile solution method for problems in stochastic computational mechanics but suffers from a slow convergence rate. The number of simulations required to produce an acceptable accuracy is often impractical for complex and time-consuming numerical models. In this paper, an element-based control variate approach is developed to improve the efficiency of Monte Carlo simulation in stochastic finite element analysis, with particular reference to high-dimensional and nonlinear geotechnical problems. The method uses a low-order element to form an inexpensive approximation to the output of an expensive, high-order model. By keeping the mesh constant, a high correlation between low-order and high-order models is ensured, enabling a large variance reduction to be achieved. The approach is demonstrated by application to the bearing capacity of a strip footing on a spatially variable soil. The problem requires 300 input random variables to represent the spatial variability by random fields, and would be difficult to solve by methods other than Monte Carlo simulation. Using an element-based control variate reduces the standard deviation of the mean bearing capacity by approximately half. In addition, two methods for estimating the cumulative distribution function as a complement to the improved mean estimator are presented.

Keywords: Control variate, Stochastic finite element, Spatial variability, Random field, Bearing capacity.

¹Ph.D student, School of Engineering, Newcastle University, Newcastle upon Tyne, NE1 7RU, UK. Email: T.S.Charlton@newcastle.ac.uk.

²Reader in Computational Geomechanics, School of Engineering, Newcastle University, Newcastle upon Tyne, NE1 7RU, UK. (corresponding author). Email: M.Rouainia@ncl.ac.uk

³Professor of Earth Systems, School of Engineering, Newcastle University, Newcastle upon Tyne, NE1 7RU, UK. Email: Richard.Dawson@ncl.ac.uk

INTRODUCTION

The aim of this paper is to improve the efficiency of Monte Carlo simulation in the context of stochastic finite element analysis. Despite the development of intrusive stochastic finite element methods, which require reformulation of the governing equations (e.g. Ghanem and Spanos 2003), for many practical applications Monte Carlo simulation using an unmodified deterministic code remains the method of choice. It is relatively straightforward to implement with existing commercial finite element programs, allowing complex industrial models and nonlinear behaviour to be analysed in a probabilistic setting, and does not suffer from the ‘curse of dimensionality’, being able to handle any number of input random variables. On the other hand, the convergence of Monte Carlo simulation is well-known to be slow. For computationally demanding, time-consuming models the number of simulations required to achieve an acceptable accuracy may be unattainable.

The efficiency of a Monte Carlo estimator can be related to the variance of the estimator and the computational cost (Hammersley and Handscomb 1964). A lower variance results in narrower confidence intervals, potentially eliminating unnecessary conservatism in an engineering design, and it is desirable to minimise the time spent running computer simulations for financial and practical reasons. There are several well-developed methods for reducing the variance of a Monte Carlo estimate. The sampling schemes by which the random seeds of Monte Carlo simulation are chosen can be made efficient to suit a particular problem, for example by applying stratified or importance sampling (Fishman 1996). For small failure probabilities, subset simulation (Au and Beck 2001) has been used effectively in several engineering applications such as foundation design (Ahmed and Soubra 2012) and the strength of composite laminate joints (Li et al. 2014), while Schöbi et al. (2017) have recently developed a combined polynomial chaos-kriging method which uses an active learning algorithm to select evaluation points. Similarly, Quasi-Monte Carlo (QMC) reduces error by using function evaluations at a carefully selected, ‘low discrepancy’ sequence of points (Nie and Ellingwood 2004; Lemieux 2009). Attention here is limited to the classic Monte Carlo method using random sampling of the inputs for maximum generality.

An alternative approach to variance reduction of Monte Carlo simulation is to use a control

variate (Glasserman 2004). This method introduces a second random variable with known statistics that is correlated with the simulation output of interest (or quantity of interest, QoI) but cheaper to evaluate. The estimate of a statistic of the QoI, such as the mean, is corrected by exploiting the known error in the control variate estimate. The variance reduction is related to the strength of the correlation between the quantity of interest and control variate. In the context of stochastic finite element analysis, finding a control variate with known statistics is a difficult task due to the complexity of the governing equations, which may be nonlinear in practical scenarios.

Schmeiser et al. (2001), Pasupathy et al. (2012) and Ng and Willcox (2014) have extended the classical control variate framework to situations where analytical information is not available by estimating the statistics of the control variate through Monte Carlo simulation. A further generalisation is the Multilevel Monte Carlo (MLMC) method, developed by Giles (2008) for the solution of stochastic differential equations. The method uses a sequence of mesh densities to compute statistics of the QoI on the finest mesh through a telescoping sum of the coarsest estimate plus a series of correction terms from consecutive discretisation levels. Combinations of MLMC and the control variate technique have also been introduced for specific classes of problems governed by linear partial differential equations (PDEs). Nobile and Tesei (2015) constructed a control variate for each MLMC level in a stochastic diffusion problem characterised by a rough, high-dimensional random field. The control variate was the output from an auxiliary problem, employing a smoothed, low-dimensional random field, which could be solved by stochastic collocation, also enabling the control variate statistics to be evaluated. Vidal-Codina et al. (2015) considered linear, elliptic PDEs and used reduced basis approximations to create control variates at different levels, depending on the number of functions retained in the reduced basis. Since the reduced basis approximations are cheap to evaluate, statistics of the control variate can be inexpensively obtained.

Recently, Fairbanks et al. (2017) proposed a more general Multilevel Control Variates (MLCV) approach that is able to handle nonlinear equations and utilise standard deterministic solvers (e.g. finite element routines). The method is based on changing the mesh density from coarse to fine with a control variate constructed by a low-rank approximation, efficiently obtained in each level by

selecting evaluation points from realisations of the previous (coarser) mesh. The authors considered two examples, the natural frequency of a spatially variable linear elastic block and thermally driven flow in a square domain with a random temperature boundary, and showed that MLCV outperforms MLMC. The success of the method depends on the low-rank approximation being suitably accurate, such that a strong correlation is achieved between the QoI and control variate.

In geotechnical engineering, stochastic finite element analysis is increasingly used to investigate the effect of the natural variability of soils on the performance and behaviour of foundations and earth structures (e.g., Griffiths and Fenton 2001; Charlton and Rouainia 2016; Charlton and Rouainia 2017). Soil is a highly nonlinear material, meaning it is desirable to use existing finite element codes and material models coupled with random field representations of soil properties. The vertical autocorrelation distance within a soil deposit is usually short (Phoon and Kulhawy 1999), and so geotechnical problems are often high-dimensional and suited to Monte Carlo solutions. In this paper, an element-based control variate approach is introduced to improve the efficiency of stochastic finite element analysis for high-dimensional, nonlinear geotechnical problems. Our requirements for the control variate are that it is strongly correlated with the QoI and is straightforward to construct. The advantage of using element-order to form a control variate is that a high correlation is essentially guaranteed, with little extra effort required to set up the auxiliary finite element model, compared to the construction of a reduced basis approximation (e.g., Fairbanks et al. 2017). Using a single mesh also avoids convergence issues related to different spatial discretisations, which may be significant if the soil model is highly nonlinear. Mesh-dependency of the solution, for example if the soil model predicts softening behaviour, could also affect the correlation of the principal model and control variate if several meshes are used.

The method is demonstrated by application to the bearing capacity of a strip footing on spatially variable soil, which is modelled using a high-dimensional random field. The mean bearing capacity is calculated by a control variate estimator and compared with the results of standard Monte Carlo simulation. In engineering applications, obtaining the cumulative distribution function (CDF) or certain quantiles of the distribution may be important, for example in reliability analysis. CDF and

quantile estimation is discussed in the context of the control variate approach and two methods of obtaining an improved estimate of the CDF as a complement to the improved mean estimator are proposed.

CONTROL VARIATES

Consider a model output Y that is a function of ω , a vector of random variables. Assuming that we wish to estimate the expected value $E[Y(\omega)]$, the standard Monte Carlo estimator from n function evaluations is the sample mean $\bar{Y} = \frac{1}{n} \sum_{i=1}^n y_i$, where y_i is the output of the i th evaluation. The Monte Carlo estimator is unbiased and has a variance of $Var(\bar{Y}) = Var(Y)/n$.

The control variate method estimates $E[Y]$ by exploiting a function $X(\omega)$ that is correlated with $Y(\omega)$. In the classical method, $E[X]$ is known analytically. The control variate estimator of $E[Y]$, $\bar{Y}^{CV}$, is:

$$\bar{Y}^{CV} = \frac{1}{n} \sum_{i=1}^n y_i - \beta(x_i - E[X]) = \bar{Y} - \beta(\bar{X} - E[X]) \quad (1)$$

where β is a constant. The estimator $\bar{Y}^{CV}$ can be shown to be unbiased for any value of β (Glasserman 2004). The variance of $\bar{Y}^{CV}$ is:

$$Var(\bar{Y}^{CV}) = Var(\bar{Y}) + \frac{1}{n} (\beta^2 Var(X) - 2\beta Cov(X, Y)) \quad (2)$$

If Eq. 2 is differentiated with respect to β and the result set equal to zero, the optimal value of β that minimises the variance of $\bar{Y}^{CV}$ can be obtained as:

$$\beta^* = \frac{Cov(X, Y)}{Var(X)} \quad (3)$$

By substituting β^* into Eq. 2, it can be seen that:

$$Var(\bar{Y}^{CV}) = (1 - \rho_{XY}^2) Var(\bar{Y}) \quad (4)$$

The reduction in variance of $\bar{Y}^{CV}$ is therefore directly related to the strength of the correlation between Y and X : as ρ_{XY} approaches ± 1 , a perfect correlation, the variance of $\bar{Y}^{CV}$ goes to zero.

Unfortunately, certain practical difficulties are associated with the classical control variate technique. Firstly, the quantities required to calculate β^* are likely to be unknown. In this instance sample values are used instead (Lavenberg and Welch 1981), so that:

$$\hat{\beta}^* = \frac{\sum_{i=1}^n (x_i - \bar{X})(y_i - \bar{Y})}{\sum_{i=1}^n (x_i - \bar{X})^2} \quad (5)$$

Eq. 5 admits an alternative interpretation of control variates because $\hat{\beta}^*$ is the gradient of a linear regression on the points (x_i, y_i) by least squares (Weisberg 2005). As will be discussed later, this linear relationship between Y and X may be useful if ρ_{XY} is suitably high.

Secondly, if $Y(\omega)$ is sufficiently complex to require a numerical solution, such as by a finite element or finite difference method, then finding a control variate $X(\omega)$ to fit into the classical framework is a difficult task. On the other hand, assuming there is a ‘principal’ model that is deemed to provide an acceptably accurate solution then it may often be relatively simple to construct an auxiliary model that provides a less accurate solution but is computationally cheaper. In this paper, we use finite elements of different order to create principal (high-order) and auxiliary (low-order) models. For consistency, in the following the principal/high-order and auxiliary/low-order models are denoted Y and X respectively.

From this perspective, it is only possible to estimate the value of $E[X]$ by simulation. The control variate estimator is now:

$$\bar{Y}_n^{CV} = \bar{Y}_n - \beta(\bar{X}_n - \bar{X}_m) \quad (6)$$

where $\bar{X}_m$ is an estimate of $E[X]$, constructed by the standard Monte Carlo estimator. Since the auxiliary model is relatively cheap to evaluate, $\bar{X}_m$ is computed from a number of simulations m , with $m > n$ so that $\bar{X}_m$ is a more accurate estimate of $E[X]$ than $\bar{X}_n$.

Recently, Ng and Willcox (2014) developed an optimization algorithm incorporating a simplified auxiliary numerical model as a control variate. The authors consider the case where n principal and auxiliary simulations are run from the same random inputs, with an additional $m - n$ samples forming a control variate consisting of m simulations. Following this approach, the optimal value of β remains that of Eq. 3 and the variance of the estimator in Eq. 6 is:

$$Var(\bar{Y}_n^{CV}) = Var(\bar{Y}_n) \left[1 - \rho_{XY}^2 \left(1 - \frac{n}{m} \right) \right] \quad (7)$$

The variance reduction is diminished by the estimation of $E[X]$ as the correlation ρ_{XY}^2 is effectively reduced by the factor $(1 - n/m)$. It is clear that as $m \rightarrow \infty$, the variance reduction becomes that of the classical control variate estimator given in Eq. 5 while if $m = n$, the variance is equal to that of the Monte Carlo estimator.

CDF AND QUANTILE ESTIMATION

For engineering purposes, obtaining the CDF or specific quantiles of the model response can be an important objective of a probabilistic analysis. The direct control variate estimator for the α -quantile of $Y(\omega)$ is:

$$Y_\alpha^{CV} = Y_\alpha - \beta(X_\alpha - X_\alpha) \quad (8)$$

where X_α is the corresponding α -quantile of $X(\omega)$. The difficulty with this estimator is that the entire set of simulation outputs $\mathbf{y} = \{y_1, \dots, y_n\}$, and similarly for $\mathbf{x}$, is required for one quantile estimate. This means that the optimal value of β cannot be found because it is not possible to calculate $\hat{\beta}^*$ from the sample values (Hesterberg and Nelson 1998). In addition, in this case a fixed value of β is not sufficient to ensure an unbiased estimator (Hsu and Nelson 1990).

Instead, quantiles are generally estimated using the inverse cumulative distribution function. The CDF also allows probability estimation and provides a useful graphical representation of the performance of the physical system in terms of the quantity of interest, and so is far more versatile than a direct quantile estimate. The α -quantile is obtained by $Y_\alpha = F_Y^{-1}(\alpha)$ if F_Y is the CDF of Y .

To estimate Y_α by Monte Carlo simulation the empirical CDF, an estimator of F_Y , can be formed from the set of n independent evaluations of Y . The empirical CDF is defined by:

$$\hat{F}_{Y_n}(y) = \frac{1}{n} \sum_{i=1}^n I(y_i \leq y) \quad (9)$$

where I is the indicator function. The Monte Carlo estimator of Y_α is then $\hat{Y}_{\alpha,n} = \hat{F}_{Y_n}^{-1}(\alpha)$, which has the classic $n^{-1/2}$ convergence rate in standard deviation (David, 1981). A control variate can therefore be used to improve the estimate of the CDF and, indirectly, the quantile estimator. For generality, we now consider the control variate $C = g(X)$, with $g(X)$ being some transformation of the function $X(\omega)$. The estimator of the CDF can then be written as:

$$\hat{F}_{Y_n}^{CV}(y) = \frac{1}{n} \sum_{i=1}^n [I(y_i \leq y) - \beta(c_i - E[C])] = \hat{F}_{Y_n}(y) - \beta(\bar{C}_n - E[C]) \quad (10)$$

where $\hat{F}_{Y_n}$ is the empirical CDF of Y from n simulations, $\bar{C}_n = \frac{1}{n} \sum_{i=1}^n c_i$ and $E[C]$ is known analytically. The control variate estimator of F_Y , $\hat{F}_{Y_n}^{CV}$, can then be inverted to form the quantile estimator. However, assuming that $Cov(I(Y \leq y), C)$ is unknown, the optimal value of β must be estimated by (Chu and Nakayama, 2012):

$$\hat{\beta}^* = \frac{[\frac{1}{n} \sum_{i=1}^n I(y_i \leq y) c_i] - \hat{F}_{Y_n}(y) \bar{C}_n}{\frac{1}{n} \sum_{j=1}^n (c_j - \bar{C}_n)^2} \quad (11)$$

which is a function of y . Therefore a new β must be computed for each y value in order to construct the CDF. To overcome this difficulty, Hesterberg and Nelson (1998) note that there is an alternative way of writing $\hat{F}_{Y_n}^{CV}$:

$$\hat{F}_{Y_n}^{CV}(y) = \sum_{i=1}^n w_i I(y_i \leq y) \quad (12)$$

where:

$$w_i = \frac{1}{n} + \frac{(\bar{C}_n - C_i)(\bar{C}_n - E[C])}{\sum_{j=1}^n (c_j - \bar{C}_n)^2} \quad (13)$$

This is a weighted average interpretation of control variates with each y_i assigned a weight w_i . The weights sum to 1 and, importantly, are not dependent on y . The advantage is that the value of $\hat{F}_{Y_n}^{CV}$ can be obtained for any y without needing to recalculate $\hat{\beta}^*$. The empirical CDF is a step function, with the height of the steps at each data point equal to the weight of the sample, and with a limited number of samples it is generally desirable to smooth $\hat{F}_{Y_n}^{CV}$ to provide a better idea of the underlying function. The weighted samples can be used directly in a kernel density estimate (KDE) as follows (Bowman and Azzalini 1997):

$$\hat{f}(y) = \frac{1}{h} \sum_{i=1}^n w_i K\left(\frac{y - y_i}{h}\right) \quad (14)$$

where $\hat{f}(y)$ is the estimated probability density function (PDF), K is a kernel function with $\int K(y)dy = 1$ (usually chosen to be a suitable PDF), and h is the bandwidth controlling the amount of smoothing applied to the data (Wand and Jones 1995). The standard kernel density estimate is recovered by setting $w_i = 1/n$ (i.e. equal weights). The CDF is obtained as the integral of the density estimate.

Hesterberg and Nelson (1998) discuss several choices of control variate and show that $C = g(X) = I(x_i \leq x)$ has better variance reduction properties for extreme probabilities than using the mean-based control $g(X) = X$. Here, we limit attention to the mean-based control since, assuming the experiment will be run only once, this choice maintains the optimal variance reduction of the mean and the CDF is constructed as a useful complement.

IMPLEMENTATION

Computational cost

The optimal ratio of high-order and low-order simulations can be found by considering an equivalent number, p , of high-order model simulations (Ng and Willcox 2014); this also allows for

a direct comparison of computational cost with the standard Monte Carlo estimator. If the ratio of the time per high-order simulation to time per low-order simulation is denoted t , with $t > 1$, then:

$$p = n + \frac{m}{t} = n \left(1 + \frac{r}{t} \right) \quad (15)$$

where $r = m/n$ is the ratio of the number of low-order simulations to the number of high-order simulations. In stochastic computations t is a random quantity, and the ratio of the mean simulation times can be used.

The parameter r divides the computational budget p between the high-order and low-order models. By substituting for m and n in Eq. 9, the variance can be minimised with respect to r to obtain the optimal ratio r^* as:

$$r^* = \sqrt{\frac{t\rho_{XY}^2}{1 - \rho_{XY}^2}}. \quad (16)$$

The optimal value of β remains unchanged from Eq. 4 and the variance of the estimator in terms of the optimal parameters is:

$$Var(\bar{Y}_p^{CV}) = Var(\bar{Y}_p) \left[1 - \rho_{XY}^2 \left(1 - \frac{1}{r^*} \right) \right] \left(1 + \frac{r^*}{t} \right). \quad (17)$$

In practice r^* must be estimated from the available sample values because ρ_{XY}^2 and t are unknown. The sample estimates are denoted $\hat{r}^*$, $\hat{\rho}_{XY}$, and $\hat{t}$.

Algorithm

A procedure to construct a control variate estimator of $E[Y]$ with a fixed computational cost is outlined in Algorithm 1. To begin the experiment, a value for n must be chosen. A reasonable estimate of both $\hat{\rho}_{XY}^2$ and $\hat{t}$ can be obtained from a limited initial run (n_0) of each model (from the same random inputs) consisting of, for example, 10 simulations. An initial value of n (n_1) can then be obtained by rearranging Eq. 15.

To avoid inflation of the computational cost, a series of bootstrap samples is used to obtain a

95% confidence interval around $\hat{\rho}_{XY}^2$ and the upper bound used to calculate n_1 . The correlation between the models is likely to be of more use for manipulating the experiment than the ratio of computational time since t may be affected by factors such as background processes which might be difficult to control.

After evaluating both the high-order and low-order model $n_1 - n_0$ times, an updated and more accurate $\hat{\rho}_{XY}^2$ is obtained. The value of n_1 is then also updated (n_2); since the new estimate of $\hat{\rho}_{XY}^2$ will almost certainly fall inside the confidence interval, n is adjusted slightly upwards ($n_2 > n_1$). Subsequently the high-order and low-order models can be run $n_2 - n_1$ and $m - n_2$ times respectively and the control variate estimator constructed.

Algorithm 1 Element-based control variate algorithm for a fixed computational cost

1. Initial run of n_0 simulations of high-order and low-order models: $Y(\mathbf{\Omega}_{n_0})$ and $X(\mathbf{\Omega}_{n_0})$, where n_0 is a parameter and $\mathbf{\Omega}_{n_0} = \{\omega_1, \dots, \omega_{n_0}\}$ is a set of n_0 independent realisations of ω .
2. Use a series of bootstrap samples to construct a 95% confidence interval for $\hat{\rho}_{XY}^2$; take $\hat{\rho}_{XY}^2$ as the upper bound.
3. Estimate $\hat{t}$ and obtain $\hat{r}^*$ from Eq. 16 using sample values.
4. Calculate n_1 from Eq. 15 and evaluate $Y(\mathbf{\Omega}_{n_1-n_0})$ and $X(\mathbf{\Omega}_{n_1-n_0})$.
5. Update $\hat{\rho}_{XY}^2$ and $\hat{r}^*$.
6. Calculate n_2 . If $(n_2 < n_1)$ then $n_2 = n_1$.
7. $m = \text{round}(n_2 \hat{r}^*)$.
8. If $(n_2 > n_1)$ then evaluate $Y(\mathbf{\Omega}_{n_2-n_1})$ and $X(\mathbf{\Omega}_{n_2-n_1})$.
9. Calculate $\hat{\beta}^*$ from Eq. 5.
10. Evaluate $X(\mathbf{\Omega}_{m-n_2})$.
11. Obtain control variate estimator $\bar{Y}_{n_2}^{CV} = \bar{Y}_{n_2} - \hat{\beta}^*(\bar{X}_{n_2} - \bar{X}_m)$.

If the goal of the simulation is to reduce the variance to a certain level, the algorithm can run adaptively by iteratively increasing the computational cost p by a prescribed amount Δp (so that in the i th iteration $p_{i+1} = p_i + \Delta p$). The estimate of $\hat{\rho}_{XY}^2$ is updated at the end of each iteration and used to compute the values of n and m in the subsequent step. By checking the variance of the estimator at the end of each iteration the algorithm can be stopped once a specified variance is achieved.

EXAMPLE: BEARING CAPACITY OF A STRIP FOOTING ON SPATIALLY VARIABLE SOIL

Here, the element-based control variate methodology is implemented for the case of a strip footing on spatially variable soil using the geotechnical finite element software Plaxis (2012). This is representative of a typical geotechnical finite element analysis and it should be noted that the proposed approach can be implemented with any finite element solver where elements of different order are available.

Finite element model

In Plaxis 2D, two isoparametric triangular elements are implemented: a 6-node triangle and a 15-node triangle. Figure 1 shows the positions of nodes and integration points for each element. The 6-node triangle has a second-order interpolation function and three integration points whilst the 15-node triangle provides fourth-order interpolation and uses 12 integration points. The 15-node triangle is capable of producing highly accurate results (Sloan and Randolph 1982). However, the better prediction of the 15-node element comes at the cost of a more intensive computational process compared to the 6-node element.

A plane strain model is used in Plaxis 2D with a footing of width 2m on a soil domain of size 15m $\times$ 6m. The bearing capacity of the footing is denoted Q , with Q_{6n} and Q_{15n} the capacities evaluated using 6-node and 15-node elements respectively. The finite element mesh is shown in Fig. 2, where fixities are also indicated. The boundaries are chosen such that there is no influence on the bearing capacity of the footing. The mesh consists of 591 elements and remains the same whether 6-node or 15-node elements are used. When 15-node elements are used, the mesh contains 4974 nodes compared to 1307 nodes in the case of 6-node elements.

The soil is drained and behaves in a linear elastic-perfectly plastic manner with a Mohr-Coulomb failure criterion. Both the effective angle of friction, ϕ' , and cohesion, c' , are represented by random fields in order to simulate the natural spatial variability of a soil deposit. This problem has been studied by several authors, such as Fenton and Griffiths (2003), Popescu et al. (2005), and Al-Bittar and Soubra (2013) and consideration of spatial variability has been shown to have an important effect on bearing capacity. The footing itself is rough and is modelled as a block of stiff, linear

elastic material. The deterministic material properties are given in Table 1.

Representation of spatial variability

A non-intrusive version of stochastic finite element analysis is used in this study, meaning that the FE solver is not modified and the calculation proceeds in the usual deterministic manner. The Monte Carlo method involves repeatedly generating a random field and passing each realisation to the deterministic finite element program to solve the governing equations. The transfer of information from random field to finite element model is achieved by discretising the random field on a stochastic mesh which is completely separate from the finite element mesh but defined over the same geometry. The stochastic mesh consists of bilinear quadrilateral elements and the random field is mapped to the integration points of the FE mesh using interpolation functions.

For simplicity, lognormal random fields are adopted for both c' and ϕ' . If $\mathbf{z}$ is a spatial location, the random fields are generated by:

$$c'(\mathbf{z}) = \exp [\mu_{L,c} + \sigma_{L,c} G_c(\mathbf{z})] \quad (18)$$

$$\phi'(\mathbf{z}) = \exp [\mu_{L,\phi} + \sigma_{L,\phi} G_\phi(\mathbf{z})] \quad (19)$$

where $\mu_{L,c}$ and $\sigma_{L,c}$ are respectively the mean and standard deviation of $\ln(c')$, and equivalently for ϕ' , and $G_c(\mathbf{z})$ and $G_\phi(\mathbf{z})$ are standard Gaussian random fields of zero mean and unit variance. The Karhunen-Loeve (KL) expansion is used to discretise the standard Gaussian random fields as follows:

$$G(\mathbf{z}) = \sum_{i=1}^{\infty} \sqrt{\lambda_i} \xi_i \Omega_i(\mathbf{z}) \quad (20)$$

where Ω and λ are the eigenfunctions and eigenvalues from a spectral decomposition of a prescribed autocorrelation function and $\{\xi_i\}_{i=1}^{\infty}$ is a set of independent standard Gaussian random variables. In practice the expansion is truncated to include only a finite number of terms and the resulting error depends on the decay of the eigenvalues, which is influenced by the chosen autocorrelation function

and the correlation length. In this study an anisotropic exponential autocorrelation function is used, which is of the form:

$$\rho_G(\mathbf{z}, \mathbf{z}') = \prod_{i=1}^d \exp\left(-\frac{|\mathbf{z} - \mathbf{z}'|}{\theta_i}\right) \quad (21)$$

where d is the number of spatial dimensions. For this autocorrelation function, analytical solutions are available for the eigenfunctions and eigenvalues (Ghanem and Spanos 2003).

Note that the stochastic mesh is defined at 0.5m intervals in the x -direction and 0.2m intervals in the y -direction, reflecting the anisotropic correlation structure. The error in the KL expansion increases rapidly at the boundaries of the domain on which the random field is defined (Sudret and Der Kiureghian 2000) and so the stochastic mesh is extended 1m beyond the boundaries of the finite element mesh.

The random fields of c' and ϕ' are also cross-correlated in the standard Gaussian space following the method of Vořechovský (2008). Previous studies have reported a negative correlation between c' and ϕ' , and Lumb (2015) suggests a range for $\rho_{c,\phi}$ of between -0.7 and -0.37. In view of this, $\rho_{c,\phi}$ is taken to be -0.5 to represent a typical scenario.

The random field properties are summarised in Table 2. Typical values for the coefficient of variation (COV) are used; the friction angle has been shown to have less variability and uncertainty associated with its value compared with the cohesion (Phoon and Kulhawy 1999). In a soil deposit, the correlation of parameters is usually an order of magnitude greater in the horizontal direction than in the vertical direction. The values of $\theta_{\text{horizontal}}$ and θ_{vertical} of $\ln(c')$ and $\ln(\phi')$ are taken to be 10m and 1m respectively, consistent with reported values in literature (e.g. Lacasse and Nadim 1996; Phoon and Kulhawy 1999). The slow decay of the eigenvalues in the KL expansion due to the short vertical autocorrelation distance means that 150 terms must be retained in the expansion in order to capture at least 90% of the variance of the random field. A total of 300 random variables is therefore required in the analysis, illustrating the high-dimensional nature of the problem.

Note that the stochastic mesh is defined at 0.5m intervals in the x -direction and 0.2m intervals in the y -direction, reflecting the anisotropic correlation structure. The error in the KL expansion

increases rapidly at the boundaries of the domain on which the random field is defined (Sudret and Der Kiureghian 2000) and so the stochastic mesh is extended 1m beyond the boundaries of the finite element mesh.

RESULTS AND DISCUSSION

Mean bearing capacity

The mean bearing capacity of the strip footing is computed using both the element-based control variate approach and standard Monte Carlo simulation. The 6-node model takes around 40s to run on a PC equipped with a 3.2GHz Intel Core i5-3470 processor and 8GB of RAM, whilst the 15-node model requires 240s. The ratio of computational times, t , is therefore approximately 6. In the commercial software environment used for this study there is essentially no extra effort necessary to setup the auxiliary model, as the element type is simply a user option. Additional time costs not considered in the algorithm are therefore likely to be minimal.

A realisation of the random fields of cohesion and friction angle is shown in Fig. 3. The anisotropic correlation function produces a spatial distribution of the cohesion and friction angle in horizontal bands, representative of the layered structure common to many soil deposits. An additional factor to note is that with fewer integration points in the 6-node element model, a coarser representation of the random field is used in the 6-node finite element analysis.

Fig. 4 shows the standard deviation of both the control variate and Monte Carlo estimators as the computational cost is increased. The standard deviation is obtained from the square root of the estimator variance. The Monte Carlo estimator uses p 15-node simulations and the control variate estimator requires n 15-node and m 6-node simulations, the total of which in terms of computational time is equal to the p simulations of the Monte Carlo estimator (e.g. when $p = 100$, the computational time is 24×10^3 seconds). For comparison, the p Monte Carlo simulations include the n 15-node simulations.

The control variate estimator offers a clear improvement, with the standard deviation being approximately half that of the Monte Carlo estimator. The reduction occurs at a similar rate for both estimators, controlled by the classic $n^{-1/2}$ convergence rate. After 950 simulations the

respective standard deviations for the control variate and Monte Carlo estimators are 2.6kPa and 5.6kPa. At this computational cost, the control variate algorithm requires 173 evaluations of the principal 15-node model and 4673 of the auxiliary 6-node model. The division of computing power between the two models is controlled by the parameter r , a function of the correlation between the principal and auxiliary models. In this experiment the correlation remained above 0.995 so the ratio of 15-node to 6-node simulations was relatively constant throughout, as shown in Fig. 5.

The variance reduction can be visualised by observing the PDF of the individual samples adjusted by the control variate technique, as in Eq. 1. The PDFs of the n unadjusted Monte Carlo samples and the n samples adjusted using the control variate are shown in Fig. 6. It can be seen that the PDF of the control variate-adjusted values becomes concentrated around the mean capacity, and the variance is correspondingly reduced.

The convergence of the mean bearing capacity is shown in Fig. 7. Again, the Monte Carlo estimator uses p runs of the 15-node model to enable a fair comparison. A reference case of 10,000 15-node simulations is used to bracket the expected value, $E[Q_{15n}]$. The figure presents a single run of the algorithm, meaning statistical error must be considered. To this end, 95% confidence intervals are presented to bound the estimates. The control variate estimator has a much narrower confidence interval than that obtained by a Monte Carlo estimator of equivalent computational cost. Indeed, the width of the interval is comparable with that of the 10,000-simulation reference case even with a computational cost of less than 500 15-node model evaluations. The control variate estimator appears to be marginally lower than the reference value, however the confidence interval always intersects the reference confidence interval, indicating that the control variate estimator is unbiased.

CDF and quantile estimation

For the same computational cost, the control variate algorithm uses fewer evaluations of the principal model than standard Monte Carlo simulation. While the variance of the mean can be effectively reduced, there are also fewer samples available to form an estimate of the CDF. This is likely to be of particular relevance when the principal model is extremely expensive, so a

tolerable computational cost may be only a few hundred simulations. In this section, we introduce two methods of constructing an improved CDF estimate in a post-processing step, i.e. from the collection of high order and low order model outputs obtained from the control variate algorithm demonstrated previously.

Firstly, weights can be assigned to each sample of Q_{15n} according to the formulation in Section 2. Fig. 8 illustrates the improvement enabled by using weighted samples for a computational budget of $p = 100$. In this case, the control variate algorithm uses just 14 simulations of the 15-node model and 502 of the 6-node model. Both the empirical and KDE-based CDFs of the weighted samples and the unweighted samples of Q_{15n} are shown alongside the CDF from the 10,000 reference simulations. Note that a Gaussian kernel is adopted, with the bandwidth taken to be the optimal value for estimating normal densities (Wand and Jones 1995). The figure presents an extreme case, with a very low computational budget, but serves to illustrate the advantages of the proposed methodology. It is clear that the empirical CDF does not produce a particularly useful function for quantile or probability estimation. The smooth KDE-based CDFs offer a better approximation of the underlying distribution function.

Secondly, the CDF estimate can be improved by exploiting the high correlation between the high-order and low-order simulations. The control variate method is based on the linear regression of Q_{15n} on Q_{6n} , with $\hat{\beta}^*$ the gradient of the regression line from a least squares fit. The linear regression therefore provides a model with which we can estimate the value of Q_{15n} from a given Q_{6n} . Denoting the estimated samples $\tilde{Q}_{15n}$ this gives a set $\hat{Q}_{15n} = \{Q_{15n}^{(1)}, \dots, Q_{15n}^{(n)}, \tilde{Q}_{15n}^{(n+1)}, \dots, \tilde{Q}_{15n}^{(m)}\}$ which can be used to form the CDF. As before, KDE can be applied to produce a smooth function. The approximation by the linear model introduces a bias into the CDF estimator, but the high correlation between high-order and low-order models ($\hat{\rho}_{XY} > 0.995$ in the presented example) ensures that the bias is small enough for this approach to remain practically useful. Note that the mean of $\hat{Q}_{15n}$ is just the control variate estimate of the mean, as in Section 5.

An example is shown in Fig. 9, for a computational cost of $p = 200$. A least squares regression of Q_{15n} on Q_{6n} from 27 evaluations of the principal model with corresponding auxiliary simulations

gives $\hat{\beta}^* = 0.917$, with an intercept of 8.349. Each $\tilde{Q}_{15n}$ sample is the estimate of $E(Q_{15n}|Q_{6n})$ from the linear regression model; in this case there are 1032 ($= m - n$) additional 6-node simulations that can be used to predict values of Q_{15n} . Fig. 10 shows the CDF constructed using the 27 15-node simulations together with the 1032 estimated samples $\tilde{Q}_{15n}$. There is an excellent fit with the reference CDF and it can be seen that exploiting the auxiliary samples produces a more accurate CDF compared to both the (mean) weighted-KDE method and from p Monte Carlo simulations (again including the n 15-node simulations).

To assess the variance of the proposed CDF estimators, multiple runs of the algorithm are required. In order to keep the computational cost reasonable, 100 experiments with a computational budget of $p = 200$ are undertaken by resampling from reference populations of 10,000 15-node and 6-node model evaluations. For each case, the CDF is constructed by using: n control variate-weighted samples of Q_{15n} ; the set $\hat{Q}_{15n}$; n unweighted samples, using the same n samples as in the weighted case; and p unweighted samples, including the n used in the control variate algorithm. The last two cases are simply standard Monte Carlo estimators.

The results are presented in Fig. 11. Comparing the weighted and unweighted samples of Q_{15n} (Fig. 11a and c), it can be seen that the variance of the weighted CDF estimator is reduced, with the weighted CDFs more tightly banded around the reference function. However, improvement in the tail regions ($F(Q) < 0.1$ and $F(Q) > 0.9$) is limited. The concentration of variance reduction around the centre of the distribution ($F(Q) \approx 0.5$) is due to the choice of control variate, $C = g(Q_{6n}) = Q_{6n}$. With this choice, the variance reduction is optimal for mean estimation but performs less well in the tails of the distribution (Hesterberg and Nelson 1998).

It is also evident that, at least away from the centre of the distribution, the weighted CDF estimate does not improve upon the standard Monte Carlo estimator (Fig. 11d), where the entire computational budget p is devoted to evaluating the 15-node model. Although a quantile-based control variate might locally improve the weighted CDF estimate, the advantage of choosing $C = Q_{6n}$ is that $\hat{\beta}^*$ is then taken from the regression of Q_{15n} on Q_{6n} and the set $\hat{Q}_{15n}$ can be obtained, as detailed earlier. It should also be remembered that other choices of the control variate

will not provide optimal variance reduction of the mean. The CDFs constructed from $\hat{Q}_{15n}$ are shown in Fig. 11b. The use of the estimated samples increases the accuracy of the CDF estimate substantially, and a clear improvement is observed when compared with the estimator from p Monte Carlo simulations.

In terms of quantile estimation, the standard deviations of the various estimators are shown in Fig. 12, calculated from the same 100 experiments used in Fig. 11. The shape of the error curves reflects the fact that estimation of tail quantiles is more challenging than quantiles near to the median. Use of the linear regression model in the set $\hat{Q}_{15n}$ almost halves the standard deviation when compared against the Monte Carlo estimator of equivalent computational budget, with the reduction being consistent across the distribution. The direct control variate quantile estimator, using Eq. 10 with β chosen to be 1 (due to the same units being used in both high-order and low-order models), achieves a slightly reduced standard deviation in the tails. However, construction of the CDF is more difficult using the direct estimator and its variance reduction performance is not as impressive as that of $\hat{Q}_{15n}$, using the linear regression model.

CONCLUSION

In this paper, an element-based control variate approach has been developed to improve the efficiency of Monte Carlo simulation in stochastic finite element analysis. An inexpensive auxiliary model consisting of low-order elements is used as a control variate and the variance of statistics of the expensive, high-order model reduced. The mesh is kept constant in both high-order and low-order models so a high correlation, and hence a large variance reduction, would always be expected.

The methodology is demonstrated by application to the bearing capacity of a strip footing on a spatially variable soil, modelled by high-dimensional cross-correlated random fields. Using an element-based control variate reduces the standard deviation of the mean bearing capacity by approximately half. Two methods have also been given to construct a CDF when the experiment is controlled by a mean-based control variate: firstly, by weighting each high-order sample, and secondly, by exploiting the linear regression interpretation of the control variate method and using

an enlarged set of high-order samples estimated by a linear regression model. The second approach is particularly successful in improving the CDF and quantile estimates.

The methodology proposed in this paper is straightforward to implement, requiring only a change in the element type, and has been demonstrated to effectively reduce the variance of the mean and CDF estimate in a typical stochastic finite element analysis. Further investigation of the correlation between different element types, and in different problems, is necessary to confirm the generality of the findings but the similarity of elements of the same basic shape should ensure a suitable correlation exists if the mesh is kept constant.

ACKNOWLEDGEMENTS

The first author is funded by a studentship grant from the Engineering and Physical Sciences Research Council (EPSRC) and Atkins. This support is gratefully received and acknowledged.

REFERENCES

- Ahmed, A. and Soubra, A. H. (2012). "Probabilistic analysis of strip footings resting on a spatially random soil using subest simulation approach." *Georisk*, 6(3), 188–201.
- Al-Bittar, T. and Soubra, A. H. (2013). "Bearing capacity of strip footings on spatially random soils using sparse polynomial chaos expansion." *International Journal for Numerical and Analytical Methods in Geomechanics*, 37(31), 2039–2060.
- Au, S.-K. and Beck, J. L. (2001). "Estimation of small failure probabilities in high dimensions by subset simulation." *Probabilistic Engineering Mechanics*, 16(4), 263–277.
- Bowman, A. W. and Azzalini, A. (1997). *Applied Smoothing Techniques for Data Analysis: The Kernel Approach with S-Plus Illustrations*. Clarendon Press, Oxford.
- Charlton, T. S. and Rouainia, M. (2016). "Probabilistic capacity analysis of suction caissons in spatially variable clay." *Computers and Geotechnics*, 80, 226–236.
- Charlton, T. S. and Rouainia, M. (2017). "A probabilistic approach to the ultimate capacity of skirted foundations in spatially variable clay." *Structural Safety*, 65, 126–136.
- Fairbanks, H. R., Doostan, A., Ketelsen, C., and Iaccarino, G. (2017). "A low-rank control vari-

- ate for multilevel monte carlo simulation of high-dimensional uncertain systems.” *Journal of Computational Physics*, 341, 121–139.
- Fenton, G. A. and Griffiths, D. V. (2003). “Bearing-capacity prediction of spatially random $c - \phi$ soils.” *Canadian Geotechnical Journal*, 40(1), 54–65.
- Fishman, G. S. (1996). *Monte Carlo: Concepts, Algorithms and Applications*. Springer-Verlag, New York.
- Ghanem, R. G. and Spanos, P. D. (2003). *Stochastic Finite Elements: A Spectral Approach*. Dover Publications, New York, 2nd edn. edition.
- Giles, M. (2008). “Multilevel monte carlo path simulation.” *Operations Research*, 56(3), 607–617.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer, New York.
- Griffiths, D. V. and Fenton, G. A. (2001). “Bearing capacity of spatially random soil: the undrained clay prandtl problem revisited.” *Géotechnique*, 51(4), 351–359.
- Hammersley, J. M. and Handscomb, D. C. (1964). *Monte Carlo Methods*. Chapman and Hall, London.
- Hesterberg, T. C. and Nelson, B. L. (1998). “Control variates for probability and quantile estimation.” *Management Science*, 44, 1295–1312.
- Hsu, J. C. and Nelson, B. L. (1990). “Control variates for quantile estimation.” *Management Science*, 36(7), 835–851.
- Lacasse, S. and Nadim, F. (1996). “Uncertainties in characterising soil propertiest.” *Report No. 49-75*, ASCE, New York.
- Lavenberg, S. S. and Welch, P. D. (1981). “Control variates for probability and quantile estimation.” *Management Science*, 27(3), 322–335.
- Lemieux, C. (2009). *Monte Carlo and Quasi-Monte Carlo Sampling*. Springer, New York.
- Li, H.-S., Xia, S., and Luo, D.-M. (2014). “A probabilistic analysis for pin joint bearing strength in composite laminates using subset simulation.” *Composites: Part B*, 56, 780–789.
- Lumb, P. (2015). “Safety factors and the probability distribution of soil strength.” *Canadian Geotechnical Journal*, 7(3), 225–242.

- Ng, L. W. T. and Willcox, K. E. (2014). “Multifidelity approaches for optimization under uncertainty.” *International Journal for Numerical and Analytical Methods in Engineering*, 100(10), 746–772.
- Nie, J. and Ellingwood, B. R. (2004). “A new directional simulation method for system reliability. part i: application of deterministic point sets.” *Probabilistic Engineering Mechanics*, 19(34), 425–436.
- Nobile, F. and Tesei, F. (2015). “A multi level monte carlo method with control variate for elliptic pdes with log-normal coefficients.” *Stochastic Partial Differential Equations: Analysis and Computations*, 3(3), 398–444.
- Pasupathy, R., Schmeiser, B. W., Taafe, M. R., and Wang, J. (2012). “Control-variate estimation using estimated control means.” *IIE Transactions*, 44, 381–385.
- Phoon, K. K. and Kulhawy, F. H. (1999). “Characterization of geotechnical variability.” *Canadian Geotechnical Journal*, 36(4), 612–624.
- Plaxis (2012). *PLAXIS 2D 2012 - Reference Manual*. Delft.
- Popescu, R., Deodatis, G., and Nobahar, A. (2005). “Effects of random heterogeneity of soil properties on bearing capacity.” *Probabilistic Engineering Mechanics*, 20(4), 324–341.
- Schöbi, R., Sudret, B., and Marelli, S. (2017). “Rare event estimation using polynomial-chaos kriging.” *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 3(2).
- Schmeiser, B. W., Taafe, M. R., and Wang, J. (2001). “Biased control-variate estimation.” *IIE Transactions*, 33, 219–228.
- Sloan, S. W. and Randolph, M. F. (1982). “Numerical prediction of collapse loads using finite element methods.” *International Journal for Numerical and Analytical Methods in Geomechanics*, 6(1), 47–76.
- Sudret, B. and Der Kiureghian, A. (2000). “Stochastic finite elements and reliability: a state-of-the-art report.” *Report no.*, School of Civ. Engrg., Purdue Univ., Berkeley: University of California.

- 519 Vidal-Codina, F., Nguyen, N. C., Giles, M. B., and Peraire, J. (2015). “A model and variance reduc-
520 tion method for computing statistical outputs of stochastic elliptic partial differential equations.”
521 *Journal of Computational Physics*, 197, 700–720.
- 522 Vořechovský, M. (2008). “Simulation of simply cross correlated random fields by series expansion
523 methods.” *Structural Safety*, 30(4), 337–363.
- 524 Wand, P. and Jones, C. (1995). *Kernel Smoothing*. Chapman & Hall, London.
- 525 Weisberg, S. (2005). *Applied Linear Regression*. Hoboken, N. J., Wiley-Interscience.

526

List of Tables

527

1 Deterministic material properties. 24

528

2 Random field properties. 25

TABLE 1. Deterministic material properties.

Parameter	Soil	Footing
Dilatancy angle ψ [$^{\circ}$]	20	-
Elastic modulus, E [MPa]	60	20×10^3
Poisson's ratio, μ	0.3	0.2
Unit weight, γ [kN/m ³]	18	0

TABLE 2. Random field properties.

	Cohesion, c'	Friction angle, ϕ'
Distribution	Lognormal	Lognormal
μ	20kPa	30°
COV	0.25	0.1
$\theta_{\text{horizontal}}$	10m	10m
θ_{vertical}	1m	1m
Expansion order	150	150

List of Figures

1	Position and local numbering of nodes and integration points for (a) a 6-node triangular element, and (b) a 15-node triangular element.	27
2	Finite element mesh.	28
3	Realisation of random field of (a) cohesion and (b) friction angle.	29
4	Standard deviation of Monte Carlo (MC) and control variate (CV) mean estimators.	30
5	Number of principal (15-node) $[n]$ and auxiliary (6-node) $[m]$ simulations.	31
6	PDFs of Monte Carlo (MC) and control variate (CV) samples.	32
7	Convergence of Monte Carlo (MC) and control variate (CV) estimators of the mean. Dotted lines are 95% confidence intervals.	33
8	CDFs of bearing capacity using weighted and unweighted samples. Both empirical and smooth KDE-based CDFs are shown. The computational budget is $p = 100$, with 14 15-node simulations ($n = 19$) and 502 6-node simulations ($m = 502$).	34
9	Evaluated (Q_{15n}) and approximated ($\tilde{Q}_{15n}$) samples for $p = 200$ ($n = 27$, $m = 1059$). Linear regression line also shown.	35
10	Comparison of CDFs from control variate-weighted samples, Monte Carlo simulations ($p = 200$), and from the set $\hat{Q}_{15n}$ consisting of the n 15-node simulations (Q_{15n}) and the estimated ($\tilde{Q}_{15n}$) samples from the linear regression model.	36
11	Variability of CDFs from 100 experiments with $p = 200$, resampled from the reference population. CDFs estimated by (a) control variate weights, (b) $\hat{Q}_{15n}$, (c) Monte Carlo from n (unweighted) samples, and (d) Monte Carlo from p samples.	37
12	Standard deviation of quantile estimators.	38

Fig. 1. Position and local numbering of nodes and integration points for (a) a 6-node triangular element, and (b) a 15-node triangular element.

Fig. 2. Finite element mesh.

Fig. 3. Realisation of random field of (a) cohesion and (b) friction angle.

Fig. 4. Standard deviation of Monte Carlo (MC) and control variate (CV) mean estimators.

Fig. 5. Number of principal (15-node) $[n]$ and auxiliary (6-node) $[m]$ simulations.

Fig. 6. PDFs of Monte Carlo (MC) and control variate (CV) samples.

Fig. 7. Convergence of Monte Carlo (MC) and control variate (CV) estimators of the mean. Dotted lines are 95% confidence intervals.

Fig. 8. CDFs of bearing capacity using weighted and unweighted samples. Both empirical and smooth KDE-based CDFs are shown. The computational budget is $p = 100$, with 14 15-node simulations ($n = 19$) and 502 6-node simulations ($m = 502$).

Fig. 9. Evaluated (Q_{15n}) and approximated ($\tilde{Q}_{15n}$) samples for $p = 200$ ($n = 27, m = 1059$). Linear regression line also shown.

Fig. 10. Comparison of CDFs from control variate-weighted samples, Monte Carlo simulations ($p = 200$), and from the set $\hat{Q}_{15n}$ consisting of the n 15-node simulations (Q_{15n}) and the estimated ($\tilde{Q}_{15n}$) samples from the linear regression model.

Fig. 11. Variability of CDFs from 100 experiments with $p = 200$, resampled from the reference population. CDFs estimated by (a) control variate weights, (b) $\hat{Q}_{15n}$, (c) Monte Carlo from n (unweighted) samples, and (d) Monte Carlo from p samples.

Fig. 12. Standard deviation of quantile estimators.

Figure [Click here to download Figure fig1.pdf](#) ■

■ [Click here to download Figure fig2.pdf](#)

Figure

Figure

Figure 30 [Click here to download Figure fig4.pdf](#) ■

[Click here to download Figure fig5.pdf](#)

Figure [Click here to download Figure fig6.pdf](#) ■

[Click here to download Figure fig7.pdf](#)

Figure [Click here to download Figure fig9.pdf](#) ■

Figure [Click here to download Figure fig10.pdf](#) ■

Figure

[Click here to download Figure fig12.pdf](#)