Adaptive Simulations Using Perfect Control Variates

Shane Henderson, Cornell University
Burt Simon, University of Colorado at Denver

ABSTRACT

Consider a (Monte Carlo) simulation designed to estimate μ, the expected value of a random variable, X. Suppose there is a family of random variables, $Y(u),\ u \in U$, which can be evaluated during the simulation of X, with the following properties: $\nu(u^*) = 0$ for some $u^* \in U$, and $\nu'''(u)$ is continuous in a neighborhood of u^*, where $\nu(u) = \mathrm{var}(X + Y(u))$. We call $Y(u)$ a *perfect control variate*. Although they may appear rare, perfect control variates can be constructed in many simulation settings, via "approximating martingales".

We analyze various adaptive estimators for μ based on $\{X_n + Y_n(u_{n-1})\},\ n > 0$, where $u_n \to u^*$. When u_n is a sample mean of independent and identically distributed random variables we construct an estimator $\mathcal{A}_n$ which converges at rate $|\mathcal{A}_n - \mu| \sim O(n^{-1}\sqrt{\ln n})$, and prove a Central Limit Theorem for $\mathcal{A}_n$. A similar estimator, $\tilde{\mathcal{A}}_n$, converges at rate $O(n^{-1})$. (The "canonical" convergence rate for Monte Carlo simulations is $O(n^{-1/2})$.) In certain cases we can construct an estimator $\mathcal{E}_n$ that converges exponentially fast, i.e., $\limsup_n n^{-1} \ln(\mathrm{var}\mathcal{E}_n) \leq c$.