QUASI-MONTE CARLO METHODS

Phelim P. Boyle
Centre for Advanced Studies in Finance
University of Waterloo
Waterloo, Ontario, Canada N2L 3G1
Email: pboyle@uwaterloo.ca

and

Ken Seng Tan
Department of Statistics and Actuarial Science
University of Waterloo
Waterloo, Ontario, Canada N2L 3G1
Email: kstan@barrow.uwaterloo.ca

To be presented at the AFIR Conference in Cairns, Australia, August 1997

Abstract

The prices of complex derivative securities are often represented as high-dimensional integrals in modern finance. The basic Monte Carlo approach has proved useful in the evaluation of these integrals. The paper describes a recent development in this area that is generating considerable interest. Instead of using random points to evaluate the integrals as in standard Monte Carlo one can use a deterministic sequence that has suitable properties. These sequences are known as low discrepancy sequences and the method is known as quasi-Monte Carlo. The paper describes this approach and summarizes some of the applications to finance problems. We also describe some recent advances in this field.

Acknowledgements:
Ken Seng Tan is grateful for research support from the Society of Actuaries.

1 Introduction

In recent years, numerical methods have become very important in finance. There are several reasons why this has happened. First, the underlying models that describe the evolution of the prices of the basic securities and the relevant state variables have become more sophisticated. Second, the types of securities and their associated derivative offspring have become more complex. To compute the prices and risk sensitivities of these instruments one often has to evaluate high dimensional integrals. Third, advances in risk management technology and practice now mandate more comprehensive and intricate analyses at the portfolio level. For example due to regulatory requirements, many financial institutions are currently devoting considerable resources to the setting up of systems to compute Value at Risk.[1] These VaR calculations can involve thousands of different variables. Credit risk calculations can also involve extensive numerical work. Of course, the continuing improvement in computing power has made many calculations feasible that would have been impossible even a few years ago.

There is a wide range of numerical methods available for these purposes. We now discuss these in turn. First, there are problems which have closed form or analytical solutions but even in these cases we usually still require numerical work. Even when analytical expressions are available, one usually has to invoke some numerical algorithm to consummate the calculation. For example, the Black-Scholes formula furnishes us with a closed form expression for the price of a European call option on a non-dividend paying stock. However, to obtain numerical values we still need to evaluate two cumulative normal distribution functions. There are many derivative contracts for which analytical solutions do not appear to exist. However, for many cases of interest the price of the derivative is governed by a order partial differential equation (PDE): Schwartz (1975) and Brennan and Schwartz (1976) (1977) pioneered the use of finite difference solutions of PDE's in the finance literature in the late seventies. More recently Wilmott, Dewynne and Howison (1993) have espoused the PDE approach. Zvan, Forsyth and Vetzal (1996) have applied modern techniques from the PDE literature to derive efficient methods for the pricing of certain types of exotic options. The disadvan-

[1] For a comprehensive and accessible overview of Value at Risk(VaR) models and implementation issues, see Duffie and Pan (1997).

tage of PDE methods is that they can only deal with small number of state variables.

In 1979, the paper by Cox, Ross and Rubinstein introduced and popularized[2] the use of lattice (or tree) methods in computational finance. This is essentially a finite difference method but it resonates well with basic finance intuition. In the backward recursion, the value of the derivative at a given time step is equal to its discounted expected value at the next time step where the discounting is at the one period risk-free rate. Lattice methods are widely used for different types of derivatives including interest sensitive derivatives. This method works best if the tree is *recombining*. Thus it is not an efficient method for models involving non-Markovian state variables: for example, the general versions of the Heath, Jarrow and Morton (1992) model. Like PDE methods, lattice methods are most suitable for situations where the number of state variables is small.

Many problems of interest involve large numbers of state variables. Examples include mortgage-backed securities and other path-dependent securities. In these cases, the price of a derivative security whose payoff is a function of these state variables can be written as a multi-dimensional integral. For such problems, the Monte Carlo(MC) method provides a powerful and flexible tool. It was first introduced to the finance literature by Boyle (1977). The Monte Carlo method is usually straightforward to implement and because of the unprecedented improvements in computer efficiency, it can now be applied to a wide range of complex problems. Even so there are still problems for which the standard Monte Carlo approach is computationally burdensome. Various techniques have been proposed to speed up the convergence.

A recent innovation in this area has been the introduction of quasi-Monte Carlo methods. In this case the integral is evaluated using deterministic sequences rather than random sequences. These deterministic sequences have the property that they are well dispersed throughout the domain of integration. Sequences with this property are known as *low discrepancy sequences* (LD sequences). The monograph by Niederreiter (1992) provides a compre-

[2]Most textbooks credit the inspiration for the binomial tree method to Bill Sharpe. However Rendleman and Bartter (1979) also used a binomial tree to value options. David Emanuel had also sketched out this approach when he was a doctoral student in the late 1970's.

hensive overview of LD sequences and their properties. For general[3] integrals (not related to finance applications), the conventional wisdom[4] is that LD sequences outperform standard Monte Carlo for moderately high dimensions (5–25) but that standard Monte Carlo is more efficient for higher dimensional problems unless one uses a very large number of points. Several authors have used LD sequences in finance applications. These include Boyle, Broadie and Glasserman (1996), Caflisch and Morokoff (1996), Tan and Boyle (1997), Joy Boyle and Tan (1996), Ninomiya and Tezuka (1996) and Paskov and Traub (1995). The consensus seems to be that for most finance applications the superiority of the LD methods over standard MC methods persists for a much higher range of dimensions than is the case for more general integrals.

The advantages of LD methods for finance applications appears to stem from the smoothness of the integrand in many applications and the fact that the effective dimension in finance applications is sometimes lower than the nominal dimension. However, it is not the case that quasi-Monte Carlo outperforms standard Monte Carlo for all finance applications. The relative efficiency depends on the nature of the problem and the properties of the LD sequence used. Until recently a drawback of the classical LD approach has been the absence of a workable error bound. There is a deterministic upper bound for the error in quasi-Monte Carlo integration but it is invariably of no practical value in applications. In contrast the standard error of the estimate is easy to compute in the case of the crude Monte Carlo method.

The layout of the remainder of the paper is as follows. In the next section, we briefly review the standard Monte Carlo approach. In Section 3, we discuss the use of deterministic sets of points to evaluate integrals. We discuss low discrepancy sequences and related concepts in Section 4. These sequences can be used to tackle problems in computational finance and we review some of the results in Section 5. In Section 6, we discuss some recent research advances in this field.

---

[3]Numerical analysts use test integrals which are often badly behaved in that they are not very smooth and/or are highly periodic. This makes sense in their context. They want to ensure that the methods can handle *troublesome* integrals and not just well-behaved ones.

[4]For a comparison of the two approaches in the case of more general integrands see Bratley, Fox and Niederreiter (1992) and Morokoff and Caflisch (1994),(1995).

2 Standard Monte Carlo

In this section, we give a brief outline of the standard Monte Carlo method. Consider first the one-dimensional case. Suppose we wish to evaluate the integral of the function $h(x)$ over the range from 0 to 1.

$$\int_0^1 h(x)dx. \tag{1}$$

The standard (crude) Monte Carlo (MC) estimate of the integral is equal to the average value of the function taken over n points selected randomly[5] from the unit interval. From the strong law of large numbers, this estimate tends to the true value of the integral as n tends to infinity. From the central limit theorem, the standard error tends to zero as $\frac{1}{\sqrt{n}}$. The only restriction on the function h is that it be square integrable. The same results obtain for the multi-dimensional case.

The multi-dimensional case is a more natural candidate for the application of the Monte Carlo method. Assume we wish to integrate the function $f(\boldsymbol{x})$ over the d-dimensional unit hypercube, where $\boldsymbol{x}$ now represents a point in d-dimensional space. The integral in question is

$$\int_{[0,1)^d} f(\boldsymbol{x})d\boldsymbol{x} \tag{2}$$

We pick n random points that are uniformly distributed throughout the hypercube. Let f_i be the value of the function corresponding to the i-th point. The unbiased Monte Carlo estimates of the mean $\bar{f}$ and variance s_f^2 are

$$\begin{aligned} \bar{f} &= \frac{1}{n}\sum_{i=1}^{n} f_i \\ s_f^2 &= \frac{1}{n-1}\sum_{i=1}^{n}(f_i - \bar{f})^2 \end{aligned}$$

where n is the number of simulations. As in the one dimensional case the error is proportional to $\frac{1}{\sqrt{n}}$ and is thus independent of the dimension. Furthermore,

[5] In practice, the numbers selected by computer algorithms are are pseudo-random rather than random. For a discussion of this issue see Knuth (1981).

the error term is probabilistic and so it is easy to assess the accuracy of the MC estimates by constructing the confidence intervals.

To see why the MC method is of interest in finance, we recall that the the price of a (European) derivative security can often be expressed as the discounted expected value of its payoff. The expectation is taken with respect to the so called risk neutral measure. For a discussion of this issue, see the book by Duffie (1996). We state this important result in more detail below. Suppose we have M assets whose prices at time t are

$$\{S_1(t), S_2(t), ...S_M(t)\}.$$

Suppose the short term interest rate at time u is $r(u)$. Assume that there is a money market account which corresponds to rolling over funds at the short rate of interest. Let the value of the money market account at time t be $B(t)$. Hence

$$B(t) = e^{\int_0^t r(u)du}. \tag{3}$$

It is often convenient to take this money market account as the numeraire. The basic valuation result is that the absence of arbitrage implies the existence of a probability measure $\boldsymbol{Q}$ under which the deflated price of any derivative is a martingale. The value of a European derivative security V which expires at time $(t+T)$ with a payout that depends explicitly on the prices of the underlying securities at that time is given by

$$\frac{V(t)}{B(t)} = \mathrm{E}^{\boldsymbol{Q}} \left\{ \frac{V(t+T)}{B(t+T)} \mid \Omega_t \right\}. \tag{4}$$

If the short term interest rate is constant, say r, then the valuation equation becomes

$$V(t) = \mathrm{E}^{\boldsymbol{Q}} \{e^{-\int_t^{t+T} r(u)du} V(t+T) | \Omega_t\} = e^{-rT} \mathrm{E}^{\boldsymbol{Q}} \{V(t+T) | \Omega_t\}. \tag{5}$$

The expectation on the right hand side of this last equation can be written as an M-dimensional integral and the Monte Carlo method can be used to value such integrals.

We now illustrate the basic MC approach. We consider a standard European call option in the Black-Scholes framework. In this case the return on

the risky asset follows geometric Brownian motion with constant mean and diffusion parameters. We assume that the asset price is $S(t)$, the strike price is K, the volatility is σ. The current call price, C_t is given by

$$C_t = e^{-r(T-t)} \int_{-\infty}^{\infty} \max[S(T) - K, 0]\phi(z)dz \tag{6}$$

where

$$S(T) = S(t)e^{(r-\frac{1}{2}\sigma^2)(T-t)+\sigma z\sqrt{T-t}}$$

and z is a standard normal variate and $\phi(z)$ is its distribution function. For our numerical example we use following input parameters

$$\begin{aligned} S(t) &= 100 \\ K &= 100 \\ r &= 10\% \text{ p.a.} \\ T-t &= 1 \text{ year} \\ \sigma &= 20\% \text{ p.a.} \end{aligned}$$

The accurate Black-Scholes call option value for these input values is 13.270. This provides a benchmark against which we can assess the performance of the Monte Carlo method. We estimate the call price by using standard Monte Carlo to evaluate the integral. The results are given in Table 1 for different values of n. The numbers in brackets are the standard errors.

n	C_t
100	14.74(1.632)
10,000	13.202(0.160)
1,000,000	13.273(0.016)

Table 1: MONTE CARLO ESTIMATES(AND STANDARD ERRORS) OF THE CALL PRICE FOR DIFFERENT VALUES OF n

We note that the accuracy is proportional to $\frac{1}{\sqrt{n}}$. Hence one way to improve the accuracy is increase the number of simulations. However, several

other methods are available for improving the convergence without increasing the number of simulation trials. These methods are described under the rubric of *variance reduction methods*. The survey paper by Boyle, Broadie and Glasserman (1996) discusses many of these methods, especially those that have proved effective in finance applications. The use of LD sequences provides another method for speeding up the computations for certain classes of problems and this brings us to the topic of our next section.

3 Using Deterministic Sequences for Integration

In this section we give the basic intuition behind the use of deterministic sequences to value multi-dimensional integrals. In the one dimensional case, this procedure in fact is the classical approach and predates[6] the Monte Carlo approach. The basic idea extends back to Newton. The method is very simple and appealing in the one dimensional case. We now show how the method works in this case using an example from Joy, Boyle and Tan (1996).

We use the basic Black-Scholes formula for a European call option for this illustration. The call price can be written as a one dimensional integral. We transform the range of integration to the unit interval and then perform the integration using the midpoint rule. It will be no surprise to learn that the midpoint method is vastly more efficient than standard Monte Carlo in this case. From (6), we can change the integration domain to $[0,1]$ as follows:

$$\begin{aligned} C_t &= e^{-r(T-t)} \int_{-\infty}^{\infty} [S(T) - K]^+ \phi(z) dz \\ &= \int_{-\infty}^{\infty} h(z)\phi(z) dz \\ &= \int_0^1 h[\Phi^{-1}(x)] dx \\ &= \int_0^1 f(x) dx. \end{aligned}$$

[6]For a discussion of numerical integration formulae see Numerical Recipes by Press et. al. (1992). As the authors dryly remark " Where would any book on numerical analysis be without Mr Simpson and his *rule*".

To value this integral it is natural to take n equally spaced points. Suppose $n = 4$ and we select the following points

$$\frac{1}{8}, \frac{3}{8}, \frac{5}{8}, \frac{7}{8}.$$

The corresponding values of z are

$$\begin{aligned} & \{\Phi^{-1}(\frac{1}{8}), \Phi^{-1}(\frac{3}{8}), \Phi^{-1}(\frac{5}{8}), \Phi^{-1}(\frac{7}{8})\} \\ = \; & \{-1.15035, -.31864, .31864, 1.15035\}. \end{aligned}$$

If we use these four points and the numerical parameters we had earlier, we obtain an estimate of 13.363 for the integral and this is our estimate of the call price. As the accurate value in this case is 13.270, the difference is only .093 and the agreement is quite good. If we were to use the standard Monte Carlo approach to compute the same integral, we would need 117,000 points to ensure that the 95% confidence interval was .093.

Given the success of the midpoint method in the one dimensional case, it is instructive to see why this approach breaks down in high dimensions. First note that we can transform the region of integration of the d-dimensional integral to the unit hypercube. The analogue of the midpoint method is to construct a regular grid throughout the hypercube. Suppose we do so by picking two points along each axis that divide the interval into three equal parts. We use these points to construct a regular structure. In the two dimensional case we have the points: $\{(\frac{1}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3}), (\frac{2}{3}, \frac{1}{3}), (\frac{2}{3}, \frac{2}{3})\}$. For a three dimensional hypercube we have 8 points and for a d-dimensional hypercube we have 2^d points. For a thirty dimensional integral this would lead to over one billion points: clearly too large a number to handle. It is also clear that with the regular grid we have to decide in advance the number of points and that we cannot preserve the uniformity by adding points sequentially.

There are are other reasons why the regular grid approach is not a good method for high dimensions. It can be shown that the error bound for the d-dimensional trapezoidal rule is $\mathcal{O}(n^{-2/d})$. In addition the points on a regular grid, somewhat surprisingly, do not achieve as uniform dispersion throughout the hypercube as other sets of deterministic points. Fox (1987) discusses this issue. These other sets come from so-called low discrepancy sequences. These sequences have the property that they are evenly dispersed throughout the unit hypercube and that points can be added one at a time while preserving

this property. These sequences and related constructs play a key role in the evaluation of multi-dimensional integrals using the quasi-Monte Carlo approach.

4 Nets, Sequences and Discrepancy

In this section we review a number of concepts that are useful in discussing the theory behind quasi-Monte Carlo. First we define the concept of discrepancy. This is precise measure of the uniformity or rather the lack of uniformity of a set of points in the d-dimensional hypercube. The smaller the discrepancy of a set of points the more uniformly they are distributed throughout the hypercube. We might expect that the integration error would be related to the magnitude of the discrepancy and indeed there is an error bound which makes this idea precise. We also discuss the concepts of nets and sequences which provide useful tools on this area.

The concept of discrepancy is useful in discussing how well points are dispersed throughout a region. We now give a formal definition of this concept. Suppose we have an d-dimensional unit hypercube and a set of points scattered throughout this region. The discrepancy can be thought of as the greatest absolute difference between the continuous uniform probability distribution and the discrete uniform probability distribution taken over all possible subcubes of $[0,1)^d$ containing the origin. Let V be a subcube of $[0,1)^d$ containing the origin and let $v(V)$ be its volume. The discrepancy D_N^* of the sequence $\{\boldsymbol{X}_i\}$ of N points is defined as

$$D_N^* = \sup_{V \in [0,1)^s} \left| \frac{\#\text{ of points in } V}{N} - v(V) \right|.$$

The expectation of the discrepancy of a random sequence is bounded by $(\log(\log N))N^{-1/2}$. One can construct sequences whose discrepancies are smaller than this bound. Such sequences are known as low discrepancy sequences. Examples of these sequences are given by Halton (1960), Sobol' (1967), Faure (1982) and Niederreiter (1987). For many low discrepancy sequences the asymptotic form of the discrepancy satisfies

$$D_N^* = \mathcal{O}\left(\frac{(\log N)^d}{N}\right). \tag{7}$$

This bound for the discrepancy involves a constant that depends only on the dimension d. These constants are hard to estimate accurately in high dimensions and for large values of d, they "are often ridiculously large for reasonable values of n" according to Spanier and Maize (1994). It may take very high values of n before the discrepancy reaches its asymptotic level. Morokoff and Caflisch (1994) discuss this point and argue that the values of n at which this happens grow exponentially with the dimensions.

The importance of discrepancy can be seen from the integration error bound given by the Koksma-Hlawka inequality, which states that

$$\left| \int_{[0,1)^d} f(\boldsymbol{X}) d\boldsymbol{X} - \frac{1}{N} \sum_{i=1}^{N} \hat{f}(\boldsymbol{X}_i) \right| < V(f) \cdot D_N^* \tag{8}$$

where f is a function of bounded variation in the sense of Hardy and Krause, $\{\boldsymbol{X}_i\}$ is a sequence of N points in $[0,1)^d$ with discrepancy D_N^*, and $\hat{f}(\boldsymbol{X}_i)$ is the simulated function value. The Koksma-Hlawka inequality separates the integration error bound into two components. These components correspond to the roughness (smoothness) of the integrand and the uniformity of the sequence used in evaluating the function. The inequality suggests we should use a sequence with the smallest possible discrepancy in evaluating the function in order to achieve the smallest possible error bound.

The Koksma-Hlawka inequality represents an interesting result from a theoretical perspective. It furnishes us with a deterministic error bound that uses intuitive inputs. In practical applications there are three important disadvantages. First, it is very difficult to compute numerical values for the discrepancy D_N^*. Second, it is also very difficult to compute numerical values for the variation $V(f)$. Third, even when these two components are available the bound is generally extremely loose .

An important concept underlies the low discrepancy sequences is the theory of nets ((t,m,s)-nets and (t,s)-sequences) in which most of the sequences cited above can be considered as special cases. We now give a brief introduction to the net theory. See Niederreiter (1992, Chapter 4) for detailed discussion.

An *elementary interval* in base b is an interval E in $[0,1)^s$ of the form

$$E = \prod_{j=1}^{s} \left[\frac{a_j}{b^{d_j}}, \frac{(a_j+1)}{b^{d_j}} \right)$$

with $d_j \geq 0, 0 \leq a_j < b^{d_j}$ and a_j, d_j are integers. An elementary interval E is thus a subinterval of the unit-cube $[0,1)^s$ whose j-th axis has length $1/b^{d_j}$. When we divide the j-th axis into b^{d_j} equal slices and repeat the division for other axes, the subinterval obtained is the elementary interval having volume $b^{-\sum_{j=1}^{s} d_j}$.

Let $0 \leq t \leq m$ be integers. A (t, m, s)-net in base b is a finite sequence with b^m points from $[0,1)^s$ such that every elementary interval in base b of volume b^{t-m} contains exactly b^t points of the sequence. The intuition here is that the proportion of points in each sufficiently large box is equal to the volume of the box. If we hold s and b fixed then larger values of $(m-t)$ imply better uniformity properties as the volume of the elementary intervals decreases. If we fix m,s and b the best uniformity properties of a (t, m, s)-net are obtained when t =0.

An infinite sequence of points $\{\boldsymbol{X}_i\} \in [0,1)^s$ is a (t, s)-sequence in base b if for all $k \geq 0$ and $m > t$, the finite sequence $\boldsymbol{X}_{kb^m}, \ldots, \boldsymbol{X}_{(k+1)b^m-1}$ forms a (t, m, s)-net in base b.

Tan and Boyle (1997) give a more detailed discussion of the of the construction of nets and sequences using a graphical approach. Sobol' (1967) describes how to construct (t, m, s)-nets and (t, s)-sequences in base 2. Faure (1982) provides a construction of the $(0, m, s)$-net and $(0, s)$-sequence in a prime base greater than or equal to s. Niederreiter (1987) generalizes the construction of the Sobol' sequence to arbitrary bases and the Faure sequence to bases that are of prime power greater than or equal to s.

4.1 Example: van der Corput Sequences

There are many ways of generating deterministic sequences that yield low discrepancy. Among these techniques, the radical inverse method has generated the most interest. Many of the sequences are directly or indirectly constructed using this approach. Each integer is first expressed as an unique expansion for a given prime base b and the required low discrepancy point is obtained by a sysmmetric reflection of the expansion in the "decimal point". More precisely, suppose n is an arbitrary integer. We can express n in terms of the prime base b as follows:

$$n = \sum_{j=0}^{m} a_j(n) b^j \tag{9}$$

where the coefficients $a_j(n)$ take the values $\{0, 1, \ldots, b-1\}$. We use these coefficients to obtain a unique number, $\phi_b(n)$, in the unit interval as follows:

$$\phi_b(n) = \sum_{j=0}^{m} a_j(n) \frac{1}{b^{j+1}}. \tag{10}$$

It can be shown that the resulting sequence $\{\phi_b(n)\}$ has low discrepancy and is suitable for quasi-Monte Carlo applications. The sequence also has a nice property that it is uniformly distributed no matter when we stop generating the points.

We now give some numerical values to illustrate how it works. Suppose $n = 7$ and the prime base $b = 3$. We can write 7 in base 3 as follows

$$7 = 2(3^1) + 1(3^0) = \mathbf{21} \text{ (in base 3)}.$$

We perform the radical inversion about the "decimal point" to obtain

$$\phi_3(7) = \frac{1}{3} + \frac{2}{3^2} = \frac{5}{9}.$$

So corresponding to 7 we obtain the number $\frac{5}{9}$ in the unit interval. The next two numbers in this sequence are

$$\begin{aligned} \phi_3(8) &= \frac{2}{3} + \frac{2}{3^2} = \frac{8}{9}, \\ \phi_3(9) &= \frac{0}{3} + \frac{0}{3^2} + \frac{1}{3^3} = \frac{1}{27}. \end{aligned}$$

The first nine numbers in this sequence, excluding zero, are

$$\{ \frac{9}{27}, \frac{18}{27}, \frac{3}{27}, \frac{12}{27}, \frac{21}{27}, \frac{6}{27}, \frac{15}{27}, \frac{24}{27}, \frac{1}{27} \}.$$

Notice that as the new points are added to the sequence, they seem to *know* how to fill the gaps.

The one-dimensional low discrepancy sequences described above is known as van der Corput sequences. They have been extended to multi-dimension by Halton (1960), Sobol' (1967), Faure (1982) and Niederreiter (1987).

5 Applications in Finance

The basic idea of using low discrepancy sequences in quasi-Monte Carlo methods for numerical integration has been known for thirty years. Given the prevalence of high dimensional problems in finance it is not surprising that quasi-Monte Carlo methods have attracted the attention of research workers and practitioners in this field. In this section we review some of the results to date and indicate areas of further research.

There have been several studies that used low discrepancy sequences in finance applications. Birge (1994) found that low discrepancy sequences gave very accurate results for pricing simple option contracts. Joy, Boyle and Tan (1996) found that Faure sequences outperformed standard Monte Carlo for a range of complex derivatives with dimensions as high as 50. Paskov (1994) used Sobol' sequences to value mortgage-backed securities and concluded that at least for this security the use of Sobol' sequences significantly outperformed the standard Monte Carlo. Since 360-dimensional integrals were involved, this result was surprising given the prevailing view that the superiority of low discrepancy sequences did not extend to very high dimensions. For example in the case of more general (non-finance related) integrals, numerical experiments by Bratley, Fox and Niederreiter (1992) and van Rensburg and Torrie (1993) indicated that the superiority of low discrepancy algorithms decreases with the dimensions of the problem. These authors suggest that for general integrals the superiority of the low discrepancy algorithms for dimensions around 30. Against this background, Paskov's results were surprisingly good.

Boyle, Broadie and Glasserman (1996) compared the performance of both Sobol' sequences and Faure sequences against plain Monte Carlo. They used a discretely sampled Asian option based on the geometric average of the realized prices to conduct this comparison. The advantages of this security are that there is a simple formula for the exact price and it is easy to alter the dimension of the problem by including more points in the average. These authors used 500 different options with parameters selected at random to reduce reliance on any one contract. By plotting the root mean square against the dimension of the problem, they were able to analyse the relative performance of the three methods on a simple graph. In general, the two low discrepancy methods outperformed the standard Monte Carlo approach. However the relative advantage of the low discrepancy sequences (over plain

Monte Carlo) declines with increases in dimension but that the crossover point for this case is beyond dimension 100. However the crossover point depends on the sequence, on the number of points used and also on the nature of the integrand.

Ninomiya and Tezuka (1996) examined the performance of a number of low discrepancy sequences relative to each other and to standard Monte Carlo for different types of complex derivatives. They used three different securities: a pure discount bond, a swap and a mortgaged-backed security. The corresponding dimensions were 1439, 1727 and 360: high dimensions indeed!. The following low discrepancy sequences were compared: Sobol', Halton, Faure, generalized Faure and generalized Sobol'. Generalized low discrepancy sequences (see Tezuka (1993)) are based on a construction of (t, s)- sequences due to Niederreiter (1992). Ninomiya and Tezuka conclude that the generalized low discrepancy sequences perform better then both the classical low discrepancy sequences and the standard Monte Carlo method at these rarefied dimensions. They also note that the classical low discrepancy sequences underperform the standard Monte Carlo method for these problems.

Berman (1996) investigates the relative performance of Sobol' sequences against standard Monte Carlo and Monte Carlo with variance reduction. He uses a range of popular exotic options for this comparison. The dimensions considered are 8, 32 and 128. He concludes that the Sobol' sequences outperform the use of pseudo-random sequences but that the advantage decreases with dimension. When a particular type of variance reduction based on stratified sampling is introduced, the performance of the two approaches is similar.

There seems to be general agreement on the following point in the application of low discrepancy sequences to finance problems. The range of dimensions over which low discrepancy sequences outperform standard Monte Carlo is much wider in finance applications than for general integrals. For general integrals the cross over point occurs at moderate dimensions, say 10–30, while for finance applications it may be at around 100. There are a number of reasons why this is so. In finance applications the integrands are often quite regular and this favours the LD sequences. For some derivatives the effective dimension is lower than the actual dimension. Tan and Boyle (1997) explain why this is the case for Asian options. In the case of some LD sequences, the beginning coordinates are more uniformly distributed than the later coordinates and in the case of some derivatives, the earlier values have

a higher relative weight in the value. This favours the use of LD sequences.

Indeed this feature of the dependence of a derivative's price on the earlier coordinates can be exploited in the construction of an LD sequence. Suppose we wish to sample the path of Brownian motion at regular fixed time points. The conventional approach is to start at time zero and simulate the path one period at a time. For some applications, especially in finance, it is more helpful to use a Brownian bridge construction. For details see Caflisch and Morokoff (1996). The Brownian bridge construction can be used to generate a sample path in a way that puts more reliance on the initial coordinates of an LD sequence.

If we fix the start point and the end point of Brownian path, then the (conditional) distribution of in-between points is normal with known mean and variance. Let T be the total time interval and $m = 2^n$. We can generate the terminal point W_T by

$$W_T = \sqrt{T} z_1$$

where z_1 is a unit normal variate. The midpoint is generated by

$$W_{\frac{T}{2}} = \frac{1}{2} W_T + \frac{\sqrt{T}}{2} z_2$$

where z_2 is a another independent unit normal variate. Using a similar construction we can obtain values at times

$$T, \frac{T}{2}, \frac{T}{4}, \frac{3T}{4}, \cdots, \frac{T}{m}, \cdots, \frac{(m-1)T}{m}.$$

By construction, the earlier coordinates have a greater influence on the path than the later ones. This plays into the hands of LD sequences which have the feature that the initial coordinates have better equidistribution properties than the later coordinates. The power of this approach is documented by Acworth, Broadie and Glasserman (1996). These authors extend the idea to principal components and relate it to orthogonal expansions of Brownian motion.

6 Recent Advances and Future Research Directions

There has been considerable research interest in the use of LD sequences in finance applications in the last few years. In this section we summarize

two promising research directions. Owen (1995) and others have investigated the possibility of randomization(scrambling) low discrepancy sequences. The idea here is re-introducing randomness into the deterministic points. This can assist in bridging the gulf between quasi-Monte Carlo and plain Monte Carlo. This procedure has the capacity to provide standard errors and useful termination criteria and thus help overcome some drawbacks of traditional quasi-Monte Carlo. The other development comes from recent advances in the construction of (t, m, s)-nets. These new constructions may prove helpful in developing nets and sequences which have good properties for usage in quasi-Monte Carlo applications. We discuss these two developments in turn.

The idea of combining Monte Carlo and quasi-Monte Carlo goes back to Cranley and Patterson (1976) and Braaten and Weller (1979). The objective is to combine some of the best features of both methods. Standard Monte Carlo has the advantage that the probabilistic error term furnishes confidence intervals and a reliable termination criterion. On the other hand classical quasi-Monte Carlo does not have a well established termination criterion and one is never quite sure how accurate the estimate is. Different randomization techniques have been proposed to introduce enough randomness into the low discrepancy sequences to generate different estimates while still retaining good convergence properties.

In this connection the randomization technique proposed by Owen (1995) has attractive properties. The basic idea is to permute the digits that represent the coefficients of the terms in the expansion of different coordinates. This randomization technique preserves the low discrepancy property. If the original sequence is a (t, m, s)-net or a (t, s)-sequence the scrambled sequence is also a (t, m, s)-net or a (t, s)-sequence. Owen's randomization approach generates points that are more uniformly distributed than those in the original sequence. Numerical experiments indicate that the permuted sequence gives more accurate answers than the original low discrepancy sequence. Owen's procedure yields batches of random sequences and thus permits the computation of standard errors. However this method imposes a heavy computational burden and the existing published empirical studies by Owen (1995) and Owen and Tavella (1996) use low dimensional problems.

Tan and Boyle (1997) introduce a modification of Owen's randomization approach which can be applied to high dimensional problems. They have used this approach in large scale problems to value options on the geometric average of n ($n = 100, 365$) stocks. They find that the randomization im-

proves the accuracy of the estimates and that in addition it provides standard errors which can be used to establish reliable termination criteria. These results were obtained using a single type of derivative and one type of low discrepancy sequence (Faure). It will be interesting to see if similar results are obtained for a range of derivatives and different sequences.

Another other area of active research is to provide explicit construction of the (t, m, s)-nets and (t, s)-sequences. Net theory does not give any explicit construction of the nets nor does it guarantee their existence. The theory merely states that when a sequence satisfies a net property, low discrepancy is ensured. Mullen, Mahalanabis and Niederreiter (1995) and Clayman, Lawrence, Mullen, Niederreiter and Sloane (1996) provide surveys of the known constructions to date.

One of the recent constructions of nets is based on the relationship between combinatorial structures such as latin squares, orthogonal hypercubes or orthogonal arrays and the (t, m, s)-nets. Niederreiter (1987) was the first to consider the connection between nets and the combinatorial theory of orthogonal squares. This idea was later extended by Mullen and Whittle (1992). In 1992, Niederreiter showed the equivalence between $(t, t+2, s)$-nets and orthogonal arrays of strength 2. A more general equivalence between (t, m, s)-nets and orthogonal arrays was independently developed by Mullen and Schmid (1996) and Lawrence (1996).

In addition to constructing the (t, m, s)-nets in base b, extensive effort has also been devoted to constructing nets using optimal parameter values t, m, b for each dimension s. In particular, since the asymptotic behaviour for the discrepancy of nets begins at $N \approx b^{t+s}$, nets with smaller b and t are preferred. Faure sequences provide nets with $t = 0$ and b is the smallest prime base $\geq s$. We use a numerical example that will help illustrate some aspects of the differences in efficiency between these new constructions and the classical sequences. For $s = 365$, the smallest prime base b for Faure sequences is 367, hence the asymptotic behaviour becomes relevant around $N = 367^{365} > 10^{936}$. Niederreiter's construction method provides (t, m, s)-net with $t = \mathcal{O}(s \log s)$. Similarly, the minimum number of points for which the the asymptotic behaviour becomes relevant also grow exponentially.

Recently, Niederreiter and Xing (1995) developed a new construction method which is based on the theory of algebraic function over finite fields. Niederreiter and Xing (1996) provide a method which dominates all other existing explicit construction methods. They show how to construct (t, m, s)-

nets such that $t = s$ for any b. This is an impressive improvement over the former explicit constructions. For instance, for $s = 365$ and $b = 2$, the number of points at which the asymptotic behaviour becomes relevant is $N = 2^{730} > 10^{219}$. This value is considerably lower than for other constructions such as Faure sequences with $N > 10^{936}$ for the same dimension.

These construction techniques which draw on ideas from number theory, finite field theory, combinatorics and coding theory provide new approaches to the construction of sequences and nets with good uniformity properties. We are currently investigating the application of these techniques to problems in computational finance.

References

[1] Acworth, P., Broadie M. and P. Glasserman "Applications of Quasi Random Methods," *Working paper Columbia University*, (1996).

[2] Berman L., "Comparison of Path Generation Methods for Monte Carlo Valuation of Single Underlying Derivative Securities: *Working Paper* ,IBM Research Division, T. J. Watson Research Center, Yorktown Heigths, (1996) New York.

[3] Birge J. R., "Quasi-Monte Carlo Approaches to Option Pricing," *Technical Report 94-19* , Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI 48109 (1994).

[4] Boyle, P. "Options: A Monte Carlo Approach," *Journal of Financial Economics*, 4 (1977), 323–338.

[5] Boyle, P., M. Broadie and P. Glasserman "Monte Carlo Methods for Security Pricing" forthcoming, *Journal of Economics Dynamics and Control* (1996)

[6] Braaten E. and G. Weller "An Improved Low Discrepancy Sequence for Multidimensional Quasi-Monte Carlo Integration," *Journal of Computational Physics*, 33 (1979), 249–258.

[7] Bratley, P., B.L. Fox and H. Niederreiter "Implementation and Tests of Low-discrepancy sequences", *ACM Transactions on Modeling and Computer Simulation*, 2 (1992), 195–213.

[8] Brennan M. J., and E. S. Schwartz " The Pricing of Equity-Linked Life Insurance Policies with an Asset Value Guarantee," *Journal of Financial Economics*, 3 (1976), 195–213.

[9] Brennan M. J., and E. S. Schwartz " The Valuation of American Put Options," *Journal of Finance*, 32 (1977), 449–462

[10] Caflisch, R.E. and W.J. Morokoff "Valuation of Mortgage-Backed Securities Using the Quasi-Monte Carlo Method", International Association of Financial Engineers First Annual Computational Finance Conference, Stanford University, (1996).

[11] Cranley, R. and T.N.L. Patterson "Randomization of Number Theoretic Methods for Multiple Integration," *SIAM Journal of Numerical Analysis*, 23 (1976), 904–14.

[12] Clayman, A.T., K.M. Lawrence, G.L. Mullen, H. Niederreiter and N.J.A. Sloane, "Updated Tables of Parameters of (t, m, s)-Nets" preprint 1996.

[13] Cox, J. C., S. A. Ross, and M. Rubinstein, "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, 7 (1979), 229–264.

[14] Duffie, D., "Dynamic Asset Pricing Theory", 2nd ed. (1996), Princeton University Press, Princeton, NJ.

[15] Faure, H. "Discrépance de suites associées à un système de numération (en dimension s)," *Acta Arithmetica*, XLI (1982), 337–351.

[16] Fox, P. "ALGORITHM 647: Implementation and Relative Efficiency of Quasi-random Sequence Generators," *ACM Trans. Math Software*, 12 (1986), 362–376.

[17] Halton, J.H. "On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals," *Numer. Math.* 2 (1960), 84–90.

[18] Heath D., Jarrow R., and A. Morton, " Bond Pricing and the Term Structure of Interest Rates: A new Methodology for Contingent Claims Valuation," *Econometrica* 60:1 (1992), 77–106.

[19] Joe S. "Randomization of Lattice Rules for Numerical Multiple Integration", *Journal of Computational and Applied Mathematics*, 31 (1990), 299–304.

[20] Joy C., P. Boyle and K.S. Tan "Quasi-Monte Carlo Methods in Numerical Finance," *Management Science*, 42:6 (1996), 926–38.

[21] Knuth D.E., "Seminumerical Algorithms," *The Art of Computer Programming*, Vol 2, 2nd ed. (1981), Reading Mass, Addison-Wesley.

[22] Lawrence K.M., "A Combinatorial Characterization of (t, m, s)-Nets in Base b.", *J. Combinatorial Designs*, 4 (1996), 275–293.

[23] Morokoff, W.J. and R.E. Caflisch "Quasi-Monte Carlo Integration", *Journal of Computational Physics*, 122 (1995), 128–30.

[24] Morokoff, W.J. and R.E. Caflisch "Quasi-Random Sequences and their Discrepancies," *SIAM Journal Scientific Computing*, 15 (1994), 1251–79.

[25] Mullen G.L. and W. C. Schmid, "An Equivalence Between (t, m, s)-Nets and Strongly Orthogonal Hypercubes", *Journal of Combin. Theory Ser. A*, 76 (1996), 164–174.

[26] Ninomiya, S. and S. Tezuka "Toward Real-Time Pricing of Complex Financial Derivatives," *Applied Mathematical Finance* , 3 (1996), 1–20.

[27] Niederreiter, H., "Factorization of Polynomials and some Linear-algebra Problems over Finite Fields," *Linear Algebra Appl.*, 192 (1993), 301–328.

[28] Niederreiter, H., *Random Number Generation and Quasi-Monte Carlo Methods.* (1992), Philadelphia, PA: SIAM.

[29] Niederreiter H., "Point Sets and Sequences with Small Discrepancy", *Monatsh. Math.*, 104 (1987), 273–337.

[30] Niederreiter H. and C.P. Xing, "Low-discrepancy Sequences Obtained from Algebraic Function Fields Over Finite Fields", *Acta Arith.*, 72 (1995), 281–298.

[31] Niederreiter, H. and C.P. Xing "Quasi-random Points and Global Function Fields", In *Finite Fields and Applications*, (S.D. Cohen and H. Niederreiter, Eds), Lecture Note Series of the London Math Soc, Vol 233, Camb. Univ. Press, Cambridge, (1996), 269–296

[32] Niederreiter H. and C.P. Xing, "Low-discrepancy Sequences and Global Function Fields with Many Rational Places", *Finite Fields and Their Applications*, 2 (1996), 241–273.

[33] Niederreiter H., G.L. Mullen and A. Mahalanabis, "Tables of (t, m, s)-Net and (t, s)-sequence Parameters", in *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing*, (H. Niederreiter and P.J.-S. Shiue, Eds.), Lecture Notes in Statistics, 106 (1995), 58–86, Springer-Verlag, New York.

[34] Owen, A.B. "Randomly permuted (t, m, s)-Nets and (t, s)-sequences" in *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing*, (H. Niederreiter and P.J.-S. Shiue, Eds.), Lecture Notes in Statistics, 106 (1995), 299–317, Springer-Verlag, New York.

[35] Owen, A.B. and D.A. Tavella "Scrambled Nets for Value at Risk Calculations", International Association of Financial Engineers First Annual Computational Finance Conference, Stanford University, (1996)

[36] Paskov, S.H. and J.F. Traub "Faster Valuation of Financial Derivatives", *Journal of Portfolio Management*, 22:1 (1995), 113–120.

[37] Press, W., Teukolsky, S. et al. *Numerical Recipes in C*, 2nd ed. (1992) New York, NY: Cambridge University Press.

[38] Rendleman R.J., and B. J. Bartter "Two State Option Pricing " *Journal of Finance*, 15:1 (1979), 1093–1110.

[39] Sobol', I.M. "The Distribution of Points in a Cube and the Approximate Evaluation of Integrals," *U.S.S.R. Computational Math. and Math. Phys.*, 7:4 (1967), 86–112.

[40] Spanier J. and Maize E.H., "Quasi-Random Methods for Estimating Integrals using Relatively Small Samples," *SIAM Review*, 36 (1994), 18–44.

[41] Tan, K.S. and P. Boyle, "Applications of Scrambled Low Discrepancy Sequences to Exotic Options", Working Paper (1997), University of Waterloo.

[42] Tezuka S. "Polynomial Arithmetic Analogue of Halton Sequences", *ACM Transactions on Modeling and Computer Simulation*, 3 (1993), 99–107.

[43] Zvan R., Forsyth P. A., and Vetzal K., "Robust Numerical Methods for PDE Models of Asian Options," *Technical Report CS-96 28*, 1996, Department of Computer Science, University of Waterloo.