Faster Valuation of Financial Derivatives

S. Paskov*
Department of Computer Science
Columbia University
New York, NY 10027
paskov@cs.columbia.edu

and

J. Traub
Department of Computer Science
Columbia University
New York, NY 10027
traub@cs.columbia.edu

*Currently at Union Bank of Switzerland, 299 Park Avenue, New York, NY 10171

1 Introduction

Monte Carlo simulation is widely used to value complex financial instruments. Vast sums are spent annually on these methods.

Monte Carlo methods use random (or more precisely, pseudo-random) points. If we plot a moderate number of pseudo-random points in two dimensions, we observe regions where there are no points, see e.g. [5]. Rather than using pseudo-random points, it seems attractive to choose points which are as uniformly distributed as possible. There is a notion in number theory called discrepancy which measures the deviation of a set of points in d dimensions from uniformity. Although the question of which point set in d dimensions have the lowest discrepancy is open, various *low discrepancy* point sets are known.

We compared the efficacy of low discrepancy methods with Monte Carlo methods on the valuation of financial derivatives. We use a Collateralized Mortgage Obligation (CMO), provided to us by Goldman Sachs, with ten bond classes (tranches) which is formulated as the computation of ten integrals of dimension up to 360. The reasons for choosing this CMO is that it has fairly high dimension and that each integrand evaluation is very expensive making it crucial to sample the integrand as few times as possible. We believe that our conclusions regarding this CMO will hold for many other financial derivatives.

The low discrepancy sample points chosen for our tests are Sobol and Halton points. We compared the methods based on these points with the classical Monte Carlo method and also with the classical Monte Carlo method combined with antithetic variables.

An explanation of our terminology is in order here. *Low discrepancy points* are sometimes referred to as *quasi-random* points. Although in widespread use, we believe the latter term to be misleading since there is nothing random about these deterministic points. We prefer to use the terminology *low discrepancy* or *deterministic*.

We assume the finance problem has been formulated as an integral over the unit cube in d dimensions. We have built a software system called FINDER for computing high dimensional integrals. FINDER runs on a heterogeneous network of workstations under PVM 3.2 (Parallel Virtual Machine). Since workstations are ubiquitous, this is a cost-effective way to do large computations fast. Of course, FINDER can also be used to compute high dimensional integrals on a single workstation.

A routine for generating Sobol points is given in [4]. *However, we incorporated major improvements in FINDER and we stress that the results reported in this paper were obtained using FINDER.* One of the improvements was developing the table of primitive polynomials and initial direction numbers for dimensions up to 360.

This paper is based on two years of software construction and testing. Preliminary results were presented to a number of New York City financial houses in the Fall of 1993 and the Spring of 1994. A January, 1994 article in Scientific American [5] discussed the theoretical issues and reported that "Preliminary results obtained by testing certain finance problems suggest the superiority of the deterministic methods in practice." Further results were reported at a number of conferences in the summer and fall of 1994. A June, 1994 article in Business Week [1] indicates the possible superiority of low discrepancy sequences.

Details on the CMO, the numerical methods, and the test results are presented in [3]. Here we limit ourselves to stating our main conclusions and indicating typical results. For brevity, we shall refer to the method which uses Sobol points as the Sobol method.

We summarize our main conclusions regarding the evaluation of this CMO. The conclusions may be divided into three groups.

A. Deterministic and Monte Carlo Methods

The Sobol method consistently outperforms the Monte Carlo method. The Sobol method consistently outperforms the Halton method. In particular,

- The Sobol method converges significantly faster than the Monte Carlo method;
- The convergence of the Sobol method is smoother than the convergence of the Monte Carlo method. This makes automatic termination easier for the Sobol method;
- Using our standard termination criterion the Sobol method terminates 2 to 5 times faster than the Monte Carlo method often with smaller error;
- The Monte Carlo method is sensitive to the initial seed.

B. Sobol, Monte Carlo, and Antithetic Variables Methods

The Sobol method consistently outperforms the antithetic variables method, which in turn, consistently outperforms the Monte Carlo method. In particular,

- These conclusions also hold when a rather small number of sample points are used, an important case in practice. For example, for 4000 sample points, the Sobol method running on a single Sun-4 workstation achieves accuracies within range from one part in a thousand to one part in a million, depending on the tranche, within a couple of minutes;
- Statistical analysis on the small sample case further strengthens the case for the Sobol method over the antithetic variables method. For example, to achieve similar performances with confidence level 95%, the antithetic variables method needs from 7 to 79 times more sample points than the Sobol method, depending on the tranche;

- The antithetic variables method is sensitive to the initial seed. However, convergence of the antithetic variables method is less jagged than convergence of the Monte Carlo method.

C. Network of Workstations *All the methods benefit by being run on a network of workstations.* In particular,

- For N workstations, the measured speedup is at least $0.9N$, where $N \leq 25$;
- A substantial computation which took seven hours on a Sun-4 workstation took twenty minutes on the network of 25 workstations.

We emphasize that we do not claim that the Sobol method is always superior to the Monte Carlo method. We do not even claim that it is always superior for financial derivatives. After all, the test results reported here are only for one particular CMO. However, we do believe it will be advantageous to use the Sobol method for many other types of financial derivatives.

2 Numerical Methods

The idea underlying the Monte Carlo method is to replace the integral of $f(x)$, which is a continuous average, by a discrete average over randomly chosen points. More precisely, let D denote the d dimensional unit cube. We approximate

$$\int_D f(x)dx$$

by

$$\frac{1}{n}\sum_{i=1}^{n} f(t_i).$$

It is well known that if one chooses n points from a flat distribution, then the expected error is

$$E_n(f) = \frac{\sigma(f)}{\sqrt{n}},$$

where $\sigma^2(f)$ denotes the variance of f.

The Monte Carlo method has the advantage that the expected error is independent of dimension but suffers from the disadvantage that the rate of convergence is only proportional to $n^{-1/2}$. This motivates the search for methods which converge faster.

Low discrepancy methods also approximate the integral of $f(x)$ by a discrete average. However, this time the average is taken over low discrepancy points. A number of low discrepancy point sets are known. Here we confine ourselves to Sobol or Halton points. Roughly speaking both have the property that the rate of convergence is proportional to $(\log n)^d/n$. See [2] for the theory of low discrepancy points and references to the literature.

The n^{-1} factor in the convergence formula for low discrepancy points may be contrasted with the $n^{-1/2}$ convergence of Monte Carlo and suggests that low discrepancy methods are sometimes superior to Monte Carlo methods. However, a number of researchers report that this advantage decreases with increasing dimension. Furthermore, they report that the theoretical advantage of low discrepancy methods disappear for rather modest values of the dimension, say, $d \leq 30$.

However, these conclusions are based on mathematical problems specifically constructed for testing purposes or for certain problems arising in physics. As we shall see, tests on 360 dimensional integrals arising from a CMO lead to very different conclusions.

3 The Finance Problem

We tested a Collateralized Mortgage Obligation (CMO) provided to us by Goldman Sachs. This CMO consists of ten tranches which derive their cash flows from an underlying pool of mortgages. The cash flows received from the pool of mortgages are divided and distributed to each of the tranches according to a set of prespecified rules. The cash flows consist of interest and repayment of the principal. The technique of distributing the cash flows transfers the prepayment risk among different tranches. We stress that the amount of obtained cash flows will depend upon the future level of interest rates. Our problem is to estimate the expected value of the sum of present values of future cash flows for each of the tranches.

The underlying pool of mortgages has a thirty-year maturity and cash flows are obtained monthly. This leads to 360 cash flows and hence to integration in 360 dimensions. The precise mathematical formulation for this CMO may be found in Section 5 of [3].

4 Software System for Computing High Dimensional Integrals

Theory suggests that the low discrepancy deterministic methods provide an interesting alternative to the Monte Carlo method for computing high dimensional integrals. We have developed and tested a distributed software system for computing multivariate integrals on a network of workstations. The software also runs on a single workstation. The software utilizes the following sequences of sample points:

- Halton points;
- Sobol points;
- Uniformly distributed random points.

The user can choose the sequence of sample points from a menu. The software is written in a modular way so other kinds of deterministic and random number generators can be easily added. One or several multivariate functions defined over the unit cube of up to 360 variables can be integrated simultaneously.

A routine for generating Sobol points is given in [4]. However, we have made major improvements and we stress that the results reported in this paper were obtained using FINDER and not the routine in [4]. One of the improvements was developing the table of primitive polynomials and initial direction numbers for dimensions up to 360.

The software permits the use of various random number generators. In particular, RAN1 and RAN2 from [4] are used because of their wide availability and popularity.

5 Comparison of Deterministic and Monte Carlo Methods

We now present a selection of the results of extensive testing of the deterministic and Monte Carlo methods for the CMO. For the reader's convenience, the results are summarized in a number of graphs.

Figure 1 shows the results for one of the ten tranches (tranche A) of Sobol, Halton, and Monte Carlo runs with two randomly chosen initial seeds. Throughout this section, we describe results on tranche A in more details. Results for other tranches are similar unless stated explicitly. The pseudorandom generator RAN2 from [4] is used to generate random sample points for the Monte Carlo runs.

It is striking how typical this figure is of the vast amount of data we collected. We summarize our conlusions.

- The Monte Carlo method is sensitive to the initial seed;
- The deterministic methods, especially the Sobol method, converge significantly faster than the Monte Carlo method;

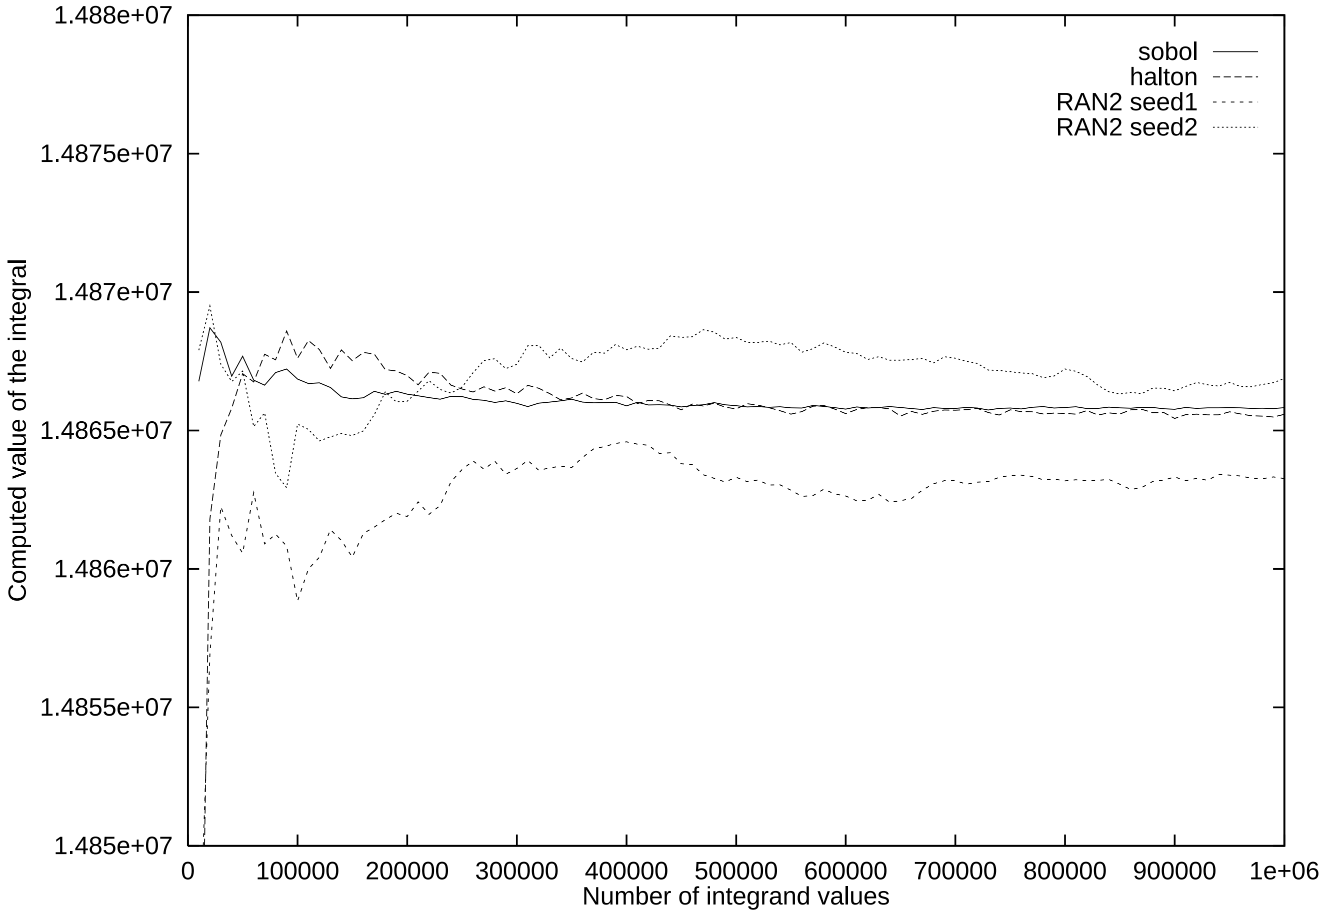

Figure 1: Sobol and Halton runs for tranche A and two Monte Carlo runs using RAN2

- The convergence of the deterministic methods, especially of the Sobol method, is smoother than the convergence of the Monte Carlo method. This makes automatic termination easier for the Sobol method; see the discussion below;
- The Sobol method outperforms the Halton method.

Figure 2 plots the same Sobol and Halton runs versus the arithmetic mean of twenty Monte Carlo runs. The twenty Monte Carlo runs use twenty different randomly chosen initial seeds. We stress that the number of sample points on the x-axis is correct only for the deterministic methods. The actual number of sample points for the averaged Monte Carlo graph is twenty times the number of sample points on the x-axis. The results of the deterministic methods and the averaged Monte Carlo result are approximately the same. After roughly the first 50,000 integrand evaluations, the behaviour of the deterministic methods and average Monte Carlo is roughly the same even though we are using 20 times more random than deterministic points.

In Figure 3, an automatic termination criterion is applied to Sobol, Halton, and three Monte Carlo runs. We choose a standard automatic termination criterion. Namely, when two consecutive differences between consecutive approximations using $10,000\,i,\ i = 1, 2, \ldots, 100$, sample points become less than some threshold value for all of the tranches of the CMO, the computational process is terminated. With the threshold value set at 250, the Sobol run terminates at 160,000 sample points, the Halton run terminates at 700,000 sample points, and the three Monte Carlo runs terminate at 410,000, 430,000, and 780,000 sample points, respectively. Hence, the Sobol run terminates 2 to 5 times faster than the Monte Carlo runs.

We stress that even though the Sobol method terminates faster, it is often more accurate than the Monte Carlo method. Details may be found in [3].

6 Antithetic Variables

An important advantage of Monte Carlo and deterministic methods is that they can be utilized very generally. This is important in a number of situations:

- If a financial house has a book with a wide variety of derivatives, it is advantageous to use methods which do not need to be tuned to a particular derivative;

- If a new derivative has to be priced, then there is no immediate opportunity to tailor a variance reduction technique to a particular integrand.

Variance reduction techniques are commonly used in conjuction with Monte Carlo methods. Although variance reduction techniques can be very powerful, they can require considerable analysis before being applied. We will therefore limit ourselves here to just one variance reduction technique; antithetic variables. The advantage of antithetic variables is that it can be easily utilized. Tests reveal that it is superior to the Monte Carlo method for our CMO problem. We emphasize that antithetic variables is not a palliative; it can be inferior to the Monte Carlo method.

Figure 4 is analogous to Figure 1. It compares the results of Sobol, Halton, and antithetic variables runs with two randomly chosen initial seeds. The data graphed in Figure 4 is typical of our results. From these results we conclude that for this CMO:

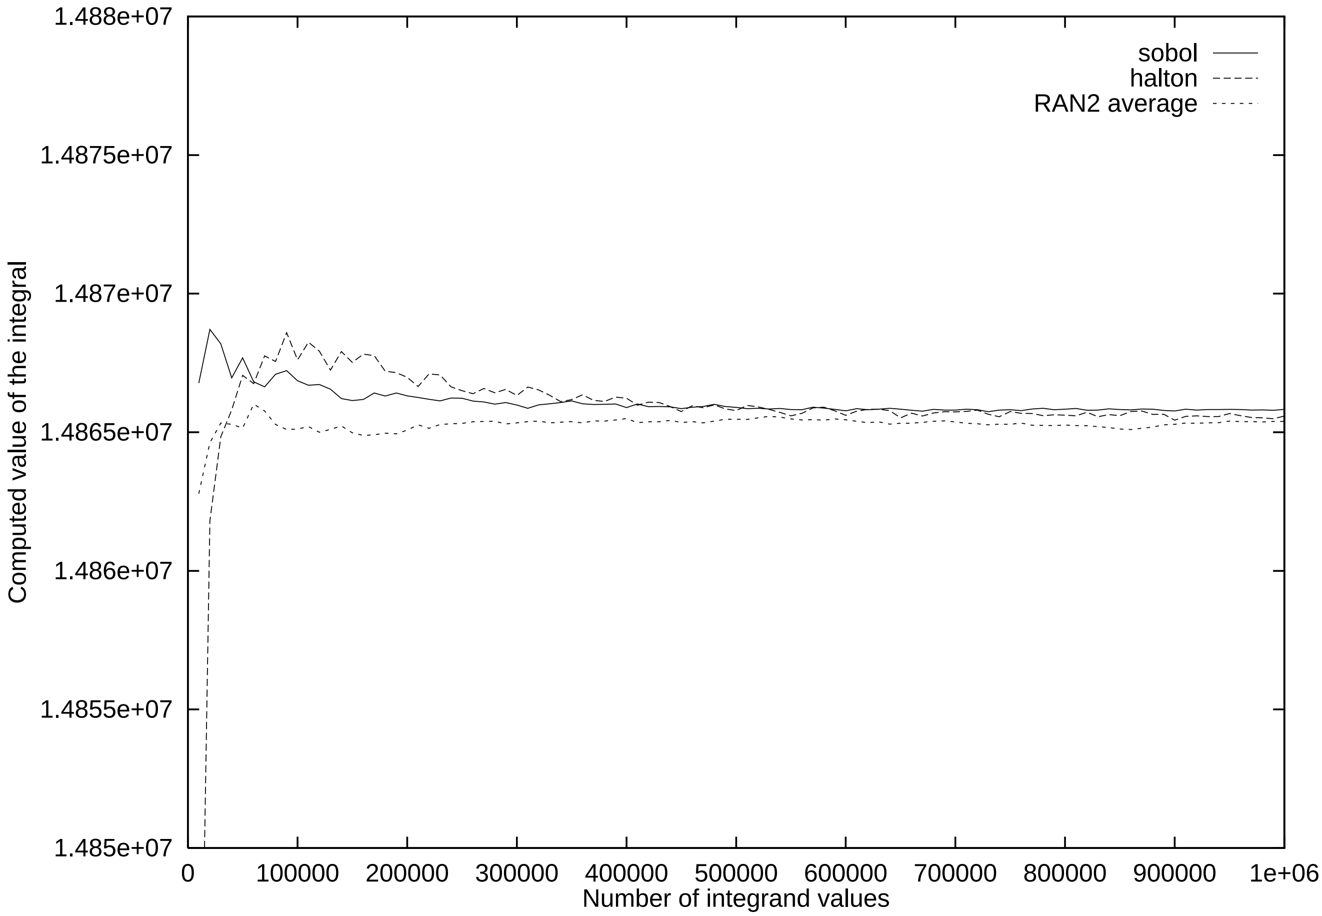

Figure 2: Sobol and Halton runs for tranche A and an average of twenty Monte Carlo runs using RAN2

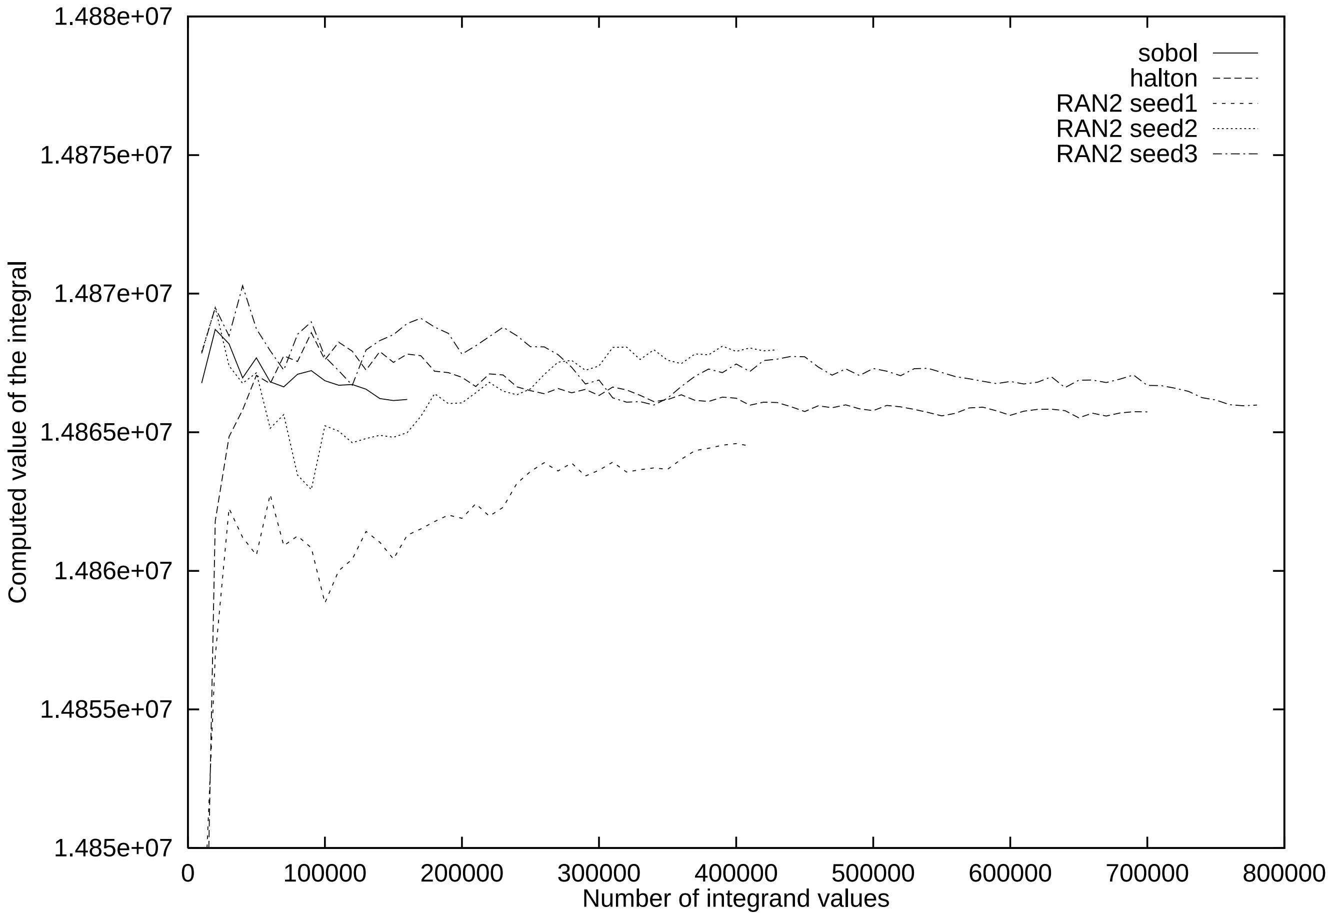

Figure 3: Automatic termination criterion applied to Sobol, Halton, and three Monte Carlo runs using RAN2 for tranche A

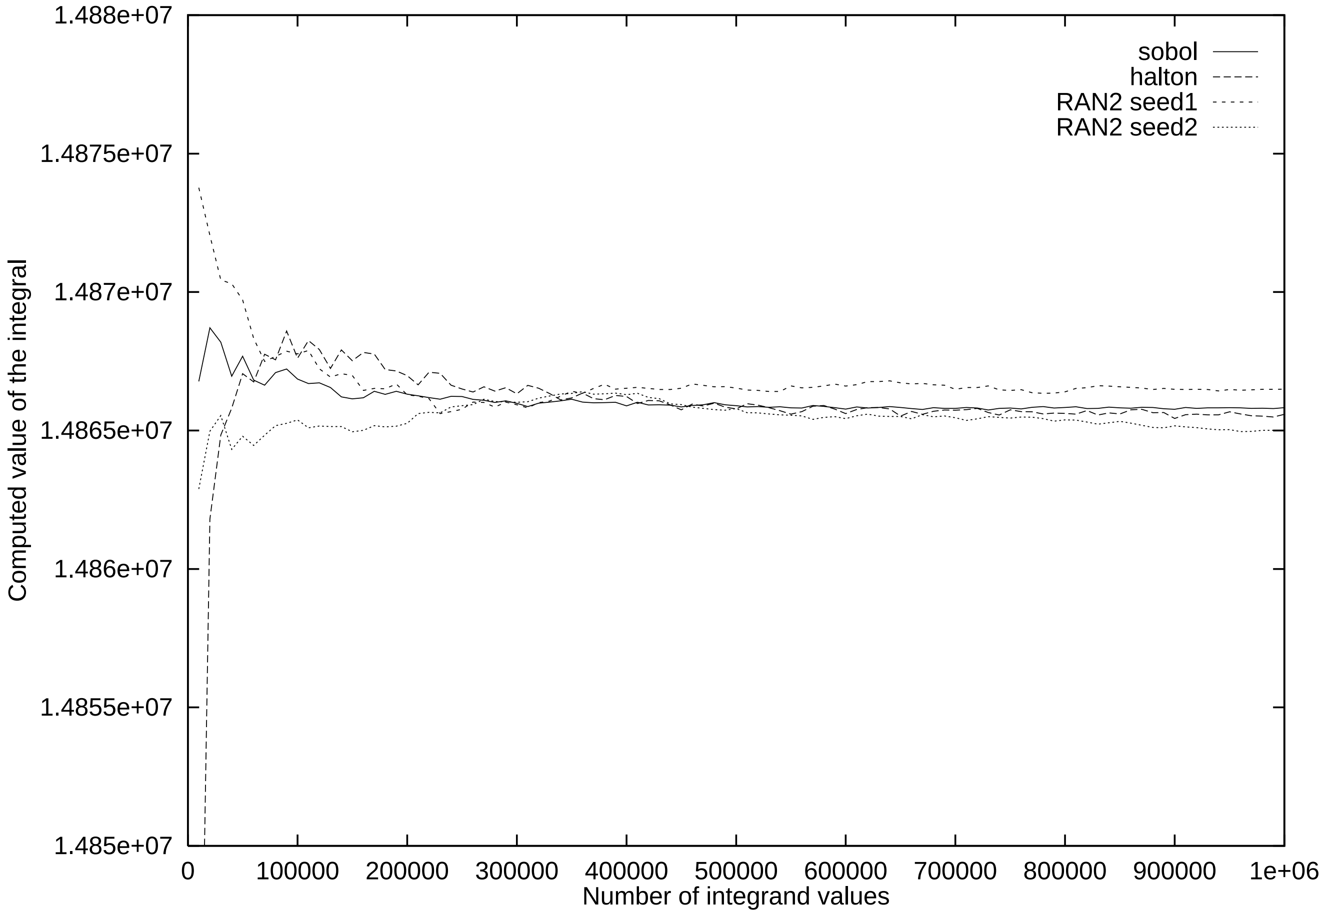

Figure 4: Sobol and Halton runs for tranche A and two antithetic variables runs using RAN2

- The Sobol method consistently outperforms the antithetic variables method;
- Convergence of the antithetic variables method is less jagged than the convergence of the Monte Carlo;
- The antithetic variables method consistently outperforms the Monte Carlo method.

Further results regarding antithetic variables may be found in [3].

7 Small Number of Sample Points

Results for a small number of points are sometimes of special importance for people who evaluate CMOs and other derivative products. They need methods which can evaluate a derivative in a matter of minutes. Rather low accuracy, on the order of 10^{-2} to 10^{-4}, is often sufficient. The integrands are complicated and computationally expensive. Furthermore, many may have to be evaluated on a daily basis with limited computational resources, such as workstations.

We therefore compare the performance of the Sobol method with Monte Carlo and antithetic variables for 4000 sample points. This leads to reasonable results and takes less than a couple of minutes of workstation CPU time. We believe that comparable results may hold for other mortgage-backed securities and interest rate derivatives. We drop the Halton method from consideration in this section since it is outperformed by both the Monte Carlo and antithetic variables methods for 4000 sample points. Sometimes computational speed is paramount. It would therefore also be of interest to study smaller number of points.

Our methodology was as follows. For each of the ten tranches we computed 20 approximate answers using the Monte Carlo method with 20 random initial seeds. For each tranche we also computed an approximation using Sobol points. We also computed the relative errors of all these approximations. To compute the relative errors we needed estimates of the true answers. We obtained these using antithetic variables with 20,000,000 points.

The results are summarized in Table 1. We say a method wins if it has a smaller relative error. (Recall we are fixing the number of samples at 4000.) Sobol points win for every tranche. In total, the Sobol method wins 177 times out of 200 cases; that is almost 90% of the time.

Table 2 exhibits the result of comparing the Sobol method with the antithetic variables method. The Sobol method wins for 8 of the tranches, ties for 1, and loses for 1. In total, Sobol wins almost 70% of the time.

The Sobol method achieves accuracies ranging from one part in a thousand to one part in a million, depending on the tranche. It takes about 103 seconds to compute the Sobol results and about 113 seconds to compute the antithetic variables results for all ten tranches running on a Sun-4 workstation.

8 Closing Remarks

We performed statistical analysis for the case of a small number of sample points. Methodology and results are reported in Section 9 of [3]. Here we confine ourselves to mentioning

Tranche	Monte Carlo	Sobol
A	3	17
B	0	20
C	3	17
D	3	17
E	2	18
G	0	20
H	0	20
J	0	20
R	8	12
Z	4	16

Table 1: Number of “wins” of the Monte Carlo method and the Sobol method

Tranche	Antithetic variables	Sobol
A	9	11
B	1	19
C	6	14
D	10	10
E	11	9
G	2	18
H	3	17
J	2	18
R	8	12
Z	9	11

Table 2: Number of “wins” of the antithetic variables method and the Sobol method

just one conclusion.

- Statistical analysis on the small sample case further strengthens the case for the Sobol method over the antithetic variables method. For example, to achieve similar performances with confidence level 95%, the antithetic variables method needs from 7 to 79 times more sample points than the Sobol method, depending on the tranche.

In closing, we suggest some directions for future work:

- Compare the performance of low discrepancy and Monte Carlo methods on other financial derivatives;
- Test the performance of other known low discrepancy sequences on various derivatives;
- As mentioned in Section 7, results for a small number of samples are often of special interest in finance. It would be attractive to design new deterministic sequences which are very uniformly distributed for a small number of points;
- Characterize analytic properties of classes of financial derivatives and design new methods tuned to these classes;
- Study error reduction techniques for deterministic methods;
- There are numerous open theoretical problems concerning high dimensional integration and low discrepancy sequences. We believe that their solution will aid in the design of better methods for finance problems.

Acknowledgments

We are grateful to Goldman Sachs for providing the finance problem. We appreciate help and discussions with S. Baker, A. Belur, B. Bojanov, T. Boult, P. Cheah, J. Davis, S. Heinrich, P. Karasinski, J. Langsam, N. Marinovich, P. Niculescu, E. Novak, D. Schutzer, B. Shukhman, I. M. Sobol, V. Temlyakov, S. Tezuka, J. Tilley, I. Vanderhoof, T. Warnock, A. Werschulz, R. Wertz, and H. Woźniakowski.

References

[1] Suddenly, Number Theory Makes Sense to Industry, *Business Week*, June 20, 172-174, 1994.

[2] Niederreiter, H., Random Number Generation and Quasi-Monte Carlo Methods, *CBMS-NSF,63, SIAM*, Philadelphia, 1992.

[3] Paskov, S. H., New Methodologies for Valuing Derivatives, *Technical Report, Comp. Sci. Dept., Columbia University,* October, 1994.

[4] Press, W., Teukolsky S., Vetterling, W., and B. Flannery, Numerical Recipes in C, Second Edition, *Cambridge University Press*, 1992.

[5] Traub, J. F. and Woźniakowski, H., Breaking Intractability, *Scientific American*, January 1994.