NESTED MONTE CARLO STUDY OF RANDOM PACKING ON THE SPHERE

R. P. C. RODGERS,* *University of California, San Francisco*
A. J. BADDELEY,** *Centre for Mathematics and Computer Science, Amsterdam*

Abstract

We consider two random sequential packing processes in which spheres of unit radius are randomly attached to the surface of a fixed unit sphere. Independent random spheres are generated and added successively, provided there is no overlap with previous spheres. In model 1, the process stops when a trial sphere intersects one of the previously-accepted spheres. In model 2, *random sequential packing*, any such overlapping trial sphere is discarded and the next random sphere is tried, until it is impossible to add any further spheres.

Previous workers have conjectured convincingly that no exact analytical solution is possible for this type of problem. We use Monte Carlo simulation methods to estimate transition probabilities for the two models. Because some probabilities are extremely small, a simulation using independent repetitions of the model would be inefficient. We designed a branching process of conditionally binomial trials, and performed over 10^8 trials on a supercomputer.

BRANCHING PROCESS; BRANCHING SIMULATION; CAR PARKING PROBLEM; CRAY XMP; MONTE CARLO; RANDOM AGGREGATION; RANDOM NUMBER GENERATORS; RANDOM POINTS ON A SPHERE; SPHERE PACKING; SUPERCOMPUTING; TESTS OF UNIFORM DISTRIBUTION

1. Introduction

In the study of molecular phenomena such as particle aggregation and gelation one considers random packing models like the following. A sphere U of unit radius is fixed. Independent, uniformly distributed random points $x_1, x_2, \cdots$ are generated on the surface of U. At each point x_i we attach a unit sphere U_i to the surface of U (i.e. so that U_i is tangent to U at that point).

It is required that spheres $U_1, \cdots, U_k$ should not overlap, i.e. their interiors should be disjoint. Two ways of prohibiting overlap are the following.

Model 1. Independent spheres $U_1, U_2, \cdots$ are added successively until an overlap occurs. The process stops at the first index k for which U_{k+1} overlaps some $U_j, j \leq k$.

Model 2. If sphere U_n overlaps some previously-added sphere, we discard it and try the next sphere U_{n+1}. The process stops when it is impossible to add another non-overlapping sphere.

In both cases we are interested in the distribution of K, the number of spheres attached when the process stops. Model 1 is similar to certain coverage models and we dub it the 'accessible surface area problem' after Wodak and Janin [19]. Model 2 is usually called *random sequential packing* and is a spherical analogue of Rényi's car parking problem [13]. In honor of the well-known joke with the punch line 'Assume a spherical cow ...', we have dubbed it the 'cow parking problem'.

These models are clearly equivalent to random packing models for spherical caps. Two spheres U_i, U_j overlap iff the corresponding spherical caps C_i, C_j on U of radius $\frac{1}{6}\pi$ centred on points x_i, x_j overlap. In what follows we use the random cap formulation.

Historically most of the interest in random sequential packing has focused on the *average packing density c*. In our case this could be interpreted as the expected fraction of sphere area *occupied* by caps C_i when the model stops; it equals the fractional area of one cap, $\frac{1}{2}(1 - \frac{1}{2}\sqrt{3})$, times the expectation of K. Note this is not directly related to the *accessible* surface area, which is the area of the set of locations where a new sphere can legally be attached.

Rényi [13] analytically solved the sequential packing density problem for segments of fixed length on the real line. In two dimensions, random sequential packing of squares and discs has been studied by Monte Carlo simulation and physical experiment [1], [2], [4], [5], [6], [8] and the apparent general consensus (see e.g. [12], [4], [5], [16]) is that analytic solutions are impossible. Similarly for the random sequential packing of hypercubes and spheres in $\mathbb{R}^n$ [1], [3], [7], [8].

Naive simulations of random sequential packing, using a sequence of independent random discs or spheres, will be very inefficient [2], [5], [6], [8], [16]. In a detailed experimental study of disc packing in the plane, Feder [5] noted the difference between maximum packing density and that achieved after N independent random discs are generated, was of order $N^{-1/2}$, meaning that in a typical simulated realization there would still be unfilled holes after $N = 10^8$ random discs. In an attempt to accelerate this, Finegold and Donnell [6] proposed a discretized simulation method wherein the centres of random discs or squares lie on a coarse grid; Tory, Jodrey and Pickard [8], [16] showed this produces biased estimates. The latter authors introduced an efficient and unbiased Monte Carlo method based on constraining the rejection sampling, i.e. keeping a list of subregions where a trial disc will always be rejected.

In this paper we develop a new (in this context) Monte Carlo method for sequential simulations, in which each successful configuration of k objects is used to generate *several* random configurations of $k + 1$ objects, in a nested fashion. The method is used to estimate the distribution of K under Models 1 and 2 of random sphere packing.

The next section formulates our problem; Section 3 describes the Monte Carlo method for Model 1; Section 4 considers the efficiency of the simulation relative to naive methods; Section 5 describes the adaptation of the simulation method to Model 2; then we report on implementation and results.

2. Formulation

Let $C_1, C_2, \cdots$ be independent, uniformly distributed spherical caps of radius $\frac{1}{6}\pi$ on the surface of a unit sphere. Write E_k for the event that $C_1, \cdots, C_k$ do not overlap pairwise. In Model 1 the process stops at the first k for which E_{k+1} does not occur. In Model 2 we generate C_k conditional on E_k, stopping when $\mathbb{P}(E_{k+1} \mid C_1, \cdots, C_k) = 0$, i.e. when it is impossible to add further caps to the current configuration of k caps.

Our main objective is to estimate the conditional probability of success, $\alpha_k = \mathbb{P}(E_k \mid E_{k-1})$, and the conditional probability of complete obstruction $\pi_k = \mathbb{P}(\mathbb{P}(E_k \mid C_1, \cdots, C_{k-1}) = 0 \mid E_{k-1})$ with respect to which the distribution of stopping time K is

$$\mathbb{P}(K = k) = \alpha_1 \cdots \alpha_k (1 - \alpha_{k+1})$$

$$\mathbb{P}(K = k) = (1 - \pi_1) \cdots (1 - \pi_k)\pi_{k+1}$$

for Models 1 and 2 respectively.

Define the random variable $A_k = \mathbb{P}(E_k \mid C_1, \cdots, C_{k-1})$ called the *accessible surface area fraction* since $4\pi A_k$ is the area of the set of points x_k such that the associated cap C_k does not overlap any cap $C_1, \cdots, C_{k-1}$. Then we can rewrite $\alpha_k = \mathbb{E}(A_k \mid E_{k-1})$, $\pi_k = \mathbb{P}(A_k = 0 \mid E_{k-1})$, and we also have $\mathbb{E}(A_k) = \mathbb{P}(E_k) = \alpha_1 \cdots \alpha_k = \beta_k$, say.

Clearly $A_1 \equiv 1$. Since the area of a spherical cap of radius θ is $2\pi(1 - \cos\theta)$, elementary geometry shows that

$$A_2 \equiv 1 - 2\pi(1 - \cos(\tfrac{1}{3}\pi))/(4\pi) = \tfrac{3}{4} \qquad \text{(a.s.)}$$

For $2 < k < 12$ the distribution of A_k is non-degenerate. A non-overlapping configuration of 12 caps occurs only when the centre points x_i are the vertices of an inscribed regular icosahedron: clearly this has zero probability so $A_{12} \equiv 0$ (a.s.).

Note that the accessible surface area A_k is not directly related to the fraction of total area covered by the $k - 1$ caps. The latter is $(k - 1)b$ where $b = \frac{1}{2}(1 - \frac{1}{2}\sqrt{3})$. The mean packing density is $c = b\mathbb{E}(K)$.

Clearly the inaccessible area $1 - A_k$ is greater than the total area occupied by caps, $b(k - 1)$, and less than the total area of $(k - 1)$ caps of radius $\frac{1}{3}\pi$, i.e. we have the following coarse bounds on A_k:

$$(1) \qquad 1 - \tfrac{1}{4}(k - 1) \leqq A_k \leqq 1 - \tfrac{1}{2}(k - 1)(1 - \tfrac{1}{2}\sqrt{3}) \qquad \text{(a.s.).}$$

The upper bound is never sharp: it gives 0.933 for $k = 2$ (correct value 0.75) and 0.263 for $k = 12$ (correct value 0). The lower bound is sharp for $k = 2$, but useless for $k > 5$ when it is negative.

Wodak and Janin [19] proposed a coarse analytic approximation to accessible surface area for not-necessarily-equal spheres which does not depend on the positions of the spheres. In our case their approximation reduces to

$$(2) \qquad A_k \approx \left(\frac{3}{4}\right)^{k-1}$$

which is correct for $k = 2$ and gives 0.042 for $k = 12$. This is claimed to give errors averaging 20% in general [19].

3. Simulation of accessible surface area

A straightforward Monte Carlo simulation using independent repetitions of $\{C_1, \cdots, C_{11}\}$ would generate samples from the distribution of stopping time K. These probabilities for $k > 3$ are quite small: for example our pilot experiments gave $\beta_7 = \alpha_1 \cdots \alpha_7 \approx 10^{-4}$. Highly accurate estimates of α_k, π_k were required, particularly for larger values of k, to study the long-term behaviour of gelation models.

In order to estimate the conditional probabilities α_k, π_k it is inefficient to throw away each successful realization of a $(k-1)$-cap configuration after only one trial of cap k. Instead, we use a branching Monte Carlo scheme which generates t_k conditionally independent trials of cap C_k for each successful realization of $C_1, \cdots, C_{k-1}$.

The remainder of this section concentrates on the estimation of α_k. Our Monte Carlo simulation runs as follows.

Stage 1.   Generate a fixed number t_1 of independent realizations of the first cap C_1.

Stage 2.   For each realization of C_1, generate a fixed number t_2 of independent trials, each attempting to fit a random cap C_2. A trial is successful if C_1, C_2 do not overlap.

Stage 3.   For each *successful* configuration $\{C_1, C_2\}$, make a fixed number t_3 of independent trials, attempting to fit a cap C_3. A trial is successful if C_1, C_2, C_3 do not overlap.

Stage k.   For each *successful* configuration $\{C_1, \cdots, C_{k-1}\}$, make a fixed number t_k of independent trials, attempting to fit a cap C_k.

Let S_k be the total number of successful configurations at stage k. A simple estimate of α_k is

$$(3) \qquad \hat{\alpha}_k = \frac{S_k}{t_k S_{k-1}} \ ,$$

the proportion of successes in a random number of trials.

Lemma 1.   *Defining* $\beta_k = \alpha_1 \cdots \alpha_k$ *and* $m_k = t_1 \cdots t_k$ *we have*

$$(4) \qquad \mathbb{E}(S_k) = m_k \beta_k.$$

This is clear by linearity (a proof is subsumed in the next result.) Thus $\hat{\beta}_k = S_k / m_k$ is an unbiased estimator of β_k, and

$$\frac{\mathbb{E}(S_k)}{t_{k-1}\mathbb{E}(S_{k-1})} = \alpha_k.$$

However, (3) is a biased estimator of α_k.

Clearly the number of successes at stage k obtained from a given configuration $C_1, \cdots, C_{k-1}$ at stage $k-1$ is conditionally binomial (t_k, p_k) given $C_1, \cdots, C_{k-1}$.

However, the success probability $p_k = A_k$ is the accessible surface area fraction, a function of $C_1, \cdots, C_{k-1}$. Hence the variance of S_k depends on geometrical properties of the model.

Lemma 2. Defining

$$n_{jk} = m_k(t_{j+1} - 1)t_{j+2} \cdots t_k, \qquad r_{jkl} = m_k(t_{j+1} - 1)t_{j+2} \cdots t_l$$

we have

(5)
$$\operatorname{var}(S_k) = m_k(\beta_k - \beta_k^2) + \sum_{j=1}^{k-1} n_{jk} \operatorname{var}(A_{jk})$$

and for $k < l$

(6)
$$\operatorname{cov}(S_l, S_k) = \sum_{j=0}^{k} r_{jkl} \operatorname{cov}(A_{jk}, A_{jl})$$

where $A_{jk} = \mathbb{P}(E_k \mid C_1, \cdots, C_j)$.

Proof. Embed the Monte Carlo experiment in a larger model, in which t_k random caps are generated for *each* configuration of caps 1 to $k - 1$, regardless of whether these caps are overlapping. Thus there are exactly $t_1 \cdots t_k = m_k$ different, not necessarily successful, realizations of a set of k caps $C_1, \cdots, C_k$.

Denote strings of integer indices by $I = (i_1, \cdots, i_k)$. Write $|I|$ for the length k of string I. Define $I \leqq J$ iff I is an initial substring of J, that is $i_l = j_l$ for all $l = 1, \cdots, |I|$ (implying $|I| \leqq |J|$). The empty string $\varnothing$ is allowed. The greatest lower bound $I \wedge J$ of two strings I and J is the longest initial substring common to I and J. We also write I^- for the initial substring $(i_1, \cdots, i_{k-1})$ of length $|I| - 1$.

The experiment is described by a tree, with one node for each string I, and edges joining I to I^- for all I. This corresponds to joining each realization of k caps to the realization of $k - 1$ caps from which it was derived. Define the 'state' at node I as

$$\mathscr{C}(I) = \mathscr{C}(i_1, \cdots, i_k) = (C_1^{(i_1)}, C_2^{(i_1, i_2)}, \cdots, C_k^{(i_1, \cdots, i_k)})$$

where $C_1^{(i_1)}$ is the i_1th realization of cap C_1, and in general $C_k^{(i_1, \cdots, i_k)}$ is the i_kth *not necessarily successful* realization of C_k amongst the trials using $\mathscr{C}(I^-) = (C_1^{(i_1)}, C_2^{(i_1, i_2)}, \cdots, C_{k-1}^{(i_1, \cdots, i_{k-1})})$. Then the states at two nodes $\mathscr{C}(I)$ and $\mathscr{C}(J)$ are conditionally independent given $\mathscr{C}(I \wedge J)$.

The simulation could otherwise be described as a multitype branching process where the type of each individual incorporates all its previous ancestry (since $\mathscr{C}(I)$ contains $\mathscr{C}(J)$ for all $J \leqq I$). It is also formally a Markov random field on the given graph structure.

Let $Z_l(I)$ for $l \leqq |I|$ be the indicator of the event E_l for configuration $\mathscr{C}(I)$, i.e. $Z_l(I) = 1$ if $C_1^{(i_1)}, \cdots, C_l^{(i_1, \cdots, i_l)}$ are pairwise non-overlapping, and 0 otherwise. The relationship between this model and the real Monte Carlo experiment is that $C(I)$ is actually generated only when $Z_{|I|}(I) = 1$. Note that $\mathbb{E}(Z_l(I)) = \mathbb{P}(E_l) = \beta_l$. Since

(7)
$$S_k = \sum_{|I| = k} Z_k(I)$$

linearity gives $\mathbb{E}(S_k) = m_k \beta_k$. Squaring (7),

$$\mathbb{E}(S_k^2) = \mathbb{E}\left(\left[\sum_{|I|=k} Z_k(I)\right]^2\right)$$

$$= \mathbb{E}\left(\sum_{|I|=k} Z_k(I)^2\right) + \mathbb{E}\left(\sum_{I \neq J} \sum Z_k(I) Z_k(J)\right)$$

$$= m_k \beta_k + \mathbb{E}\left(\sum_{I \neq J} \sum Z_k(I) Z_k(J)\right)$$

where $|I| = |J| = k$ in the double sum. Using the conditional independence property,

$$\mathbb{E}(Z_k(I) Z_k(J)) = \mathbb{E}(\mathbb{E}(Z_k(I) \mid \mathscr{F}_l)^2)$$

$$= \mathbb{E}(A_{l,k}^2)$$

$$= \operatorname{var}(A_{l,k}) + \beta_k^2$$

where $l = |I \wedge J|$. The number of pairs (I, J) such that $|I| = |J| = k$, $|I \wedge J| = l < k$ is

$$m_l t_{l+1}(t_{l+1} - 1)(t_{l+2} \cdots t_k)^2 = n_{k,l}$$

so

$$\mathbb{E}(S_k^2) = m_k \beta_k + \sum_{l=0}^{k-1} n_{k,l} \operatorname{var}(A_{l,k}^2) + \beta_k^2 \sum_{l=0}^{k-1} n_{k,l}.$$

Subtracting $[\mathbb{E}(S_k)]^2$ and evaluating $\sum_{l=0}^{k-1} n_{k,l} = m_k$ gives the result. The cross moment is calculated in a similar way.

Lemma 3. *To first order*,

(8) $$\mathbb{E}(\hat{\alpha}_k) \approx \alpha_k(1 - \operatorname{cov}(R_k, R_{k-1}))$$

where $R_k = S_k/m_k \beta_k$.

This is an application of the delta method (i.e. expand X/Y in a Taylor series).

4. Accuracy

No explicit results have been obtained for the variance of $\hat{\alpha}_k$. Instead, we use empirical variance estimates as follows.

As a special case of the conditional independence property, configurations $\mathscr{C}(I)$ with different realizations of the *first* cap (different values of i_1) are independent. For large values of t_1, $\hat{\alpha}_k$ is asymptotically normal with mean and variance given by (8), (5).

Let $S_k(i)$ for $1 \leqq i \leqq t_1$ denote the total number of successful configurations at stage k starting from the ith (automatically successful) realization of the first cap, $C_1^{(i)}$. Then $\{(S_k(i), S_{k-1}(i)): 1 \leqq i \leqq t_1\}$ is an i.i.d sample from the joint distribution of $(S_k(1), S_{k-1}(1))$, and S_k, S_{k-1} are the marginal sample totals. So the sample moment estimator

(9) $$\frac{1}{t_1(t_1 - 1)t_k^2}\left[\sum_{i=1}^{t_1} \frac{S_k(i)^2}{S_{k-1}(i)^2} - \frac{1}{t_1}\left(\sum_{i=1}^{t_1} \frac{S_k(i)}{S_{k-1}(i)}\right)^2\right]$$

is a consistent unbiased estimator of var($\hat{\alpha}_k$), and

$$(10) \qquad \frac{1}{t_1 - 1}\left[\frac{t_1}{S_k S_{k-1}} \sum_{i=1}^{t_1} S_k(i)S_{k-1}(i) - 1\right]$$

is a consistent estimator of the fractional bias term cov(R_k, R_{k-1}) in (8).

Lemma 2 allows us to speculate about the relative efficiency of our Monte Carlo experiment compared with a simple i.i.d. sequence of simulations. Assuming that the total computational effort is proportional to the number of random caps generated, the expected effort in generating n i.i.d. simulations of model 1 up to stage k is

$$\mathbb{E}(S_1 + S_2 + \cdots + S_k) = n(\beta_1 + \cdots + \beta_k)$$

which is approximately n for $k > 4$. The variance of the associated S_k is the binomial variance $n(\beta_k - \beta_k^2)$.

In our simulations of Model 1, the expected effort is

$$\mathbb{E}(S_1 + \cdots + S_k) = m_1\beta_1 + \cdots + m_k\beta_k$$

and the variance of S_k is given in (5). If the variances of the conditional expectations $A_{l,k}$ are of the same order as β_k, then (5) is of order $m_k\beta_k$, and the branching simulations will be vastly more efficient.

5. Cow parking model

To simulate Model 2 we run the same branching design as for Model 1, with the following addition. For each successful configuration $(C_1, \cdots, C_{k-1})$ at stage $k - 1$, if *none* of the t_k trials at stage k is successful, we test whether the configuration is completely blocked ($A_k = 0$). Ignoring cases of measure zero, this occurs iff we cannot fit a cap centred at any of the $2k - 6$ vertices of the Dirichlet tessellation generated by the centres of $C_1, \cdots, C_{k-1}$.

Let N_k be the number of successful configurations at stage $k - 1$ which are found to be blocked. Using the same argument as for Lemma 2 we have $\mathbb{E}(N_k) = m_{k-1}\beta_{k-1}\pi_k$.

As with $\hat{\alpha}_k$, we use a ratio estimator

$$(11) \qquad \hat{\pi}_k = \frac{N_k}{S_{k-1}}$$

which is biased for π_k,

$$(12) \qquad \mathbb{E}(\hat{\pi}_k) \approx \pi_k(1 - \text{cov}(T_k, R_{k-1}))$$

where $R_k = S_k/(m_k\beta_k)$ and $T_k = N_k/m_{k-1}\beta_{k-1}$.

Let $N_k(i)$ for $1 \leq i \leq t_1$ be the contribution to N_k originating from the ith realization of the first cap, $C_1^{(i)}$. Then the sample moment estimator

$$(13) \qquad \frac{1}{t_1(t_1 - 1)}\left[\sum_{i=1}^{t_1} \frac{N_k(i)^2}{S_{k-1}(i)^2} - \frac{1}{t_1}\left(\sum_{i=1}^{t_1} \frac{N_k(i)}{S_{k-1}(i)}\right)^2\right]$$

is a consistent unbiased estimator of var($\hat{\pi}_k$), and

(14)
$$\frac{1}{t_1 - 1}\left[\frac{t_1}{N_k S_{k-1}} \sum_{i=1}^{t_1} N_k(i) S_{k-1}(i) - 1\right]$$

is a consistent estimator of the fractional bias term $\text{cov}(T_k, R_{k-1})$ in (12).

6. Variable sphere radius

It was scientifically interesting to know whether variations in the sizes of the interacting spheres would affect the predictions of this model. Thus the simulations described above were further elaborated by making the sphere radii follow a uniform distribution on $[1 - a, 1 + a]$ for some $a \geqq 0$. The definitions and methods used above continue to hold, except for the bounds (1). No exact geometrical test for blocked configurations was available without difficult three-dimensional geometrical calculations, so we chose an empirical approach as follows. For each successful configuration $(C_1, \cdots, C_{k-1})$ at stage $k - 1$, if *none* of the t_k trials at stage k is successful, we generated an additional u_k random caps (not used in any further stages).

Let N'_k be the number of succesful configurations at stage $k - 1$ for which *all* of the $v_k = t_k + u_k$ attempts to fit an additional cap are *unsuccessful*. Then $\mathbb{E}(N'_k) = m_{k-1}\beta_{k-1}\zeta_k$ where $\zeta_k = \mathbb{E}((1 - A_k)^{v_k} \mid E_{k-1}) > \pi_k$. By dominated convergence $\zeta_k \to \pi_k$ as $v_k \to \infty$, but the convergence depends on distributional properties. In practice we have to choose v'_k sufficiently large to make $f(x) = (1 - x)^{v_k}$ an acceptable approximation to $f(x) = 1\{x = 0\}$. We took $v_k = 10^4$. The method for estimating ζ_k was otherwise identical to the estimation of π_k in the previous section, with N_k replaced by N'_k throughout.

7. Implementation

The main computations were performed on a Cray XMP supercomputer, in integer and single precision arithmetic, employing the CFT Fortran compiler. The Fortran program was written and tested on DEC-20, Sun 3/160 and Sun-4/60 computers, using double precision arithmetic (employing the TOPS-10/TOPS-20 and f77 Fortran compilers).

Rough estimates of β_k, π_k from pilot runs were used to design the main simulation. The value of t_1 should be set as high as computer resources will allow (about 10^8 in this study). The values t_k were adjusted on the basis of the pilot runs so that the values of $S_{k-1}t_k$ (i.e. the number of trials at stage k) were approximately equal. The t_k values used for Model 1 were

$$(t_1, \cdots, t_{12}) = (10^8, 1, 1, 2, 4, 5, 10, 20, 100, 500, 1000, 1)$$

and for Model 2 we reset $t_1 = 1.1 \times 10^7$.

The simulation was carried out as a nested set of loops in indices $i_1, \cdots, i_{12}$. For each value of i_1 (representing a realization of the first cap) the values of $S_k(i_1)$ and $N_k(i_1)$ were calculated (for each $k \geqq 2$) by zeroing the registers at the beginning of the i_1 loop and incrementing them at the end of the i_k loop. At the end of the i_1 loop we then formed S_k^2, S_k/S_{k-1}, etc. and added these to running totals.

The test for blocked configurations was implemented simply by testing whether a cap could be placed at any of the $2\binom{k-1}{3}$ circumcentres of triples of points drawn from the $k-1$ centres of caps in the current configuration. This is a sufficient condition, since every vertex of the Dirichlet tessellation is the circumcentre of three points. This test requires approximately the same computational effort as first determining the Dirichlet tesselation and subsequently testing each of its vertices.

In a second simulation, the radii of the interacting spheres followed a uniform distribution on $[1 - a, 1 + a]$ with $a = 0.18$. The t_k values used were

$$(t_1, \cdots, t_{12}) = (10^8, 1, 1, 2, 4, 4, 8, 12, 25, 60, 210, 1)$$

and for Model 2 we took $u_k = 10^4$.

Random points uniformly distributed on the surface of a unit sphere were generated using the method of Marsaglia [10]. The underlying random number generator in the testing stages was a shuffled linear congruent generator [9] and later the Wichmann–Hill generator [17], [18], [11]. The final computations employed the standard linear congruent random number generator of CFT Fortran running under the CRAY CTSS operating system.

8. Results

Table 1 shows estimates of the average accessible surface area fraction α_k for the kth

TABLE 1

Estimates of mean accessible surface area α_k from 10^8 trials. Spheres with fixed radius

k	α_k	S.D. (est.)	Bias (est.)
1	1 (exact)	0	0
2	0.750001	4.3×10^{-5}	-5×10^{-23}
3	0.531087	4.8×10^{-5}	-2×10^{-9}
4	0.348677	2.7×10^{-5}	-5×10^{-9}
5	0.206223	1.3×10^{-5}	-8×10^{-9}
6	0.105552	6.0×10^{-6}	-8×10^{-9}
7	0.043246	2.0×10^{-6}	-7×10^{-9}
8	0.0133193	4.9×10^{-7}	-6×10^{-9}
9	0.00247792	5.8×10^{-8}	-5×10^{-9}
10	1.199292×10^{-4}	2.7×10^{-9}	-1.8×10^{-9}
11	3.38180×10^{-7}	1.1×10^{-11}	-1.35×10^{-10}
12	0 (exact)	0	0

cap (or incoming sphere) when all spheres have unit radius. For this model $\alpha_1 = 1$, $\alpha_2 = \frac{3}{4}$ and $\alpha_{12} = 0$. We note that the decrease in α_k is very pronounced when $k = 9, 10, 11$. An interpretation is that the 'packing constraint' enforces very small probabilities for large k.

The results for the cow parking simulation appear in Table 2. The estimate of π_k is zero for $1 \leqq k \leqq 6$. We note the obstruction probabilities increase dramatically to near 1 at $k = 10$. Again this is in accordance with the packing interpretation. Estimates of ζ_k (the

TABLE 2
Estimates of mean blocking probability π_k from 1.1×10^7 trials. Spheres with fixed radius

k	π_k	S.D. (est.)	Bias (est.)	ζ_k (est.)
$k \leqq 6$	0.0	—	—	0.0
7	0.0008473	2.5×10^{-6}	-1.2×10^{-9}	0.0021
8	0.060197	1.4×10^{-5}	-2.2×10^{-7}	0.087
9	0.459811	2.0×10^{-5}	-6.7×10^{-6}	0.517
10	0.905194	1.2×10^{-5}	-9.4×10^{-5}	0.934
11	1.000000	2.7×10^{-6}	-3.6×10^{-3}	1.000
12	1 (exact)	—	—	1

probability that a cap will fail to be fitted within 10^4 trials) are also given in Table 2. We note appreciable differences between the estimates of ζ_k and π_k, indicating that the distribution of A_k has appreciable mass near zero, in other words there is an appreciable probability that a configuration will have a non-zero but very small probability of admitting a further cap.

When the sphere radii are allowed to follow a uniform distribution over the range 0.82 to 1.18 (corresponding to a coefficient of variation of approximately 10 percent), the results for small k are little changed, but for larger k there are striking differences (Tables 3 and 4). In Table 3, α_2 appears to decline very slightly, as expected, because if the first

TABLE 3
Estimates of mean accessible surface area α_k from 10^8 trials. Spheres with variable radius

k	α_k	S.D. (est.)	Bias (est.)
1	1 (exact)	0	0
2	0.749281	4.3×10^{-5}	0
3	0.532630	4.9×10^{-5}	-2×10^{-9}
4	0.357163	2.8×10^{-5}	-5×10^{-9}
5	0.226085	1.4×10^{-5}	-9×10^{-9}
6	0.136442	7.6×10^{-6}	-1×10^{-8}
7	0.077948	3.2×10^{-6}	-1×10^{-8}
8	0.041302	1.3×10^{-6}	-2×10^{-8}
9	0.0188730	4.1×10^{-7}	-2×10^{-8}
10	0.0069556	1.0×10^{-7}	-2×10^{-8}
11	0.00208331	1.9×10^{-8}	-2.6×10^{-8}

attached sphere is larger than unit radius, it presents a disproportionately low accessible surface area. Meanwhile, α_k for $k = 9, 10, 11$ have risen by factors of as much as 10^4. We attribute this to the possibility that many of the attached spheres could be smaller than unit radius, softening the packing effect.

In the constant radius case, the packing effect is so sharp that it could be roughly approximated by a constant constraint $k < 10$. For variable radius, the effect is smoother than this.

TABLE 4

Estimates of long-term blocking probability ζ_k from 1.1×10^7 trials. Spheres with variable radius

k	ζ_k (est.)	S.D. (est.)	Bias (est.)
$k \leqq 5$	0.0	—	—
6	2.32×10^{-6}	2.3×10^{-7}	-8×10^{-13}
7	0.0025094	5.3×10^{-6}	-2×10^{-9}
8	0.0075166	7.1×10^{-6}	-1×10^{-8}
9	0.027782	9×10^{-6}	-1.2×10^{-7}
10	0.13299	1.1×10^{-5}	-2×10^{-6}
11	0.41551	1.2×10^{-5}	-2.5×10^{-5}

The estimates of standard deviation in Tables 1–4 strongly dominate the corresponding estimates of bias, until we reach about $k = 11$, when S_k, S_{k-1}, N_k become relatively small and their covariance appreciable.

Table 5 compares the estimated values of α_k in Table 1 with the Wodak–Janin approximation and the coarse bounds (1). Neither approximation is adequate for large k.

TABLE 5

Comparison of coarse approximations to α_k in Table 1

k	α_k	Wodak–Janin	Absolute bounds
1	1	1	1
2	0.750	0.750	[0.750, 0.933]
3	0.531	0.563	[0.500, 0.866]
4	0.348	0.422	[0.250, 0.799]
5	0.206	0.317	[0, 0.732]
6	0.105	0.237	[0, 0.665]
7	0.043	0.178	[0, 0.599]
8	0.013	0.133	[0, 0.531]
9	0.002	0.100	[0, 0.464]
10	1.2×10^{-4}	0.075	[0, 0.397]
11	3.4×10^{-7}	0.056	[0, 0.330]

Table 6 shows the distribution of stopping time K for Model 1 calculated from Tables 1 and 3. Fixed radius and random radius models agree reasonably well for $k \leqq 7$, after which the random-radius version tends to have larger probabilities. Also included in Table 6 is an approximation to the distribution of K, based on the Wodak–Janin approximation (2):

$$\mathbb{P}(K = k) \approx \left(\frac{3}{4}\right)^{k(k-1)/2}\left[1 - \left(\frac{3}{4}\right)^k\right].$$

This is a reasonably good approximation for $k \leqq 5$ only.

TABLE 6

Distribution of stopping time K under Model 1. Calculated from estimates of α_k in Tables 1 and 3. Wodak–Janin approximation included for comparison

k	Fixed radius	Random radius	Wodak–Janin approx.
2	0.249999	0.250722	0.250
3	0.351685	0.350188	0.328
4	0.259432	0.256550	0.244
5	0.110242	0.110314	0.121
6	0.025618	0.027829	0.043
7	0.002892	0.004054	0.011
8	0.000129	0.000329	0.002
9	7.369×10^{-6}	1.389×10^{-5}	2.9×10^{-4}
10	3.143×10^{-9}	2.653×10^{-7}	2.9×10^{-5}
11	1.747×10^{-13}	1.854×10^{-9}	2.3×10^{-6}
12	0 (exact)	0 (exact)	1.3×10^{-7}

TABLE 7

Distribution of stopping time K under Model 2

k	Fixed radius	Random radius (estimated from ζ_k)
$k \leqq 5$	0	0
6	0.0008473	2.324×10^{-6}
7	0.0601460	0.002509
8	0.4317656	0.007498
9	0.4591516	0.027504
10	0.0480895	0.128001
11	0.0	0.346735
$k \geqq 12$	0 (exact)	0.487750

Table 7 shows the distribution of k for model 2 (random sequential packing) calculated from Tables 2 and 4. Here we note a dramatic difference between the fixed-radius and random-radius cases. For fixed radius spheres, $\mathbb{P}(K = k)$ is unimodal, and it is very likely (probability 0.881) that random sequential packing will stop at $k = 9$ or 10.

For random radii, the Model 2 predictions are rough guesses in that they were computed from ζ_k rather than π_k estimates, and $\mathbb{P}(K = k)$ was recorded only for $k \leqq 11$. These values are increasing in k and it is highly likely (probability 0.94) that packing will not stop before $k = 10$. Since the distribution of sphere radius was uniform on [0.82, 1.18] and thus bounded below, $\mathbb{P}(K = k)$ will drop to zero at some finite k.

It is interesting to compare the predicted values of packing density for this model with those obtained for random sequential packing in $\mathbb{R}^n$. For one-dimensional segments, Rényi [13] obtained an analytic expression for the packing density c giving $c \approx 0.7476$. In $\mathbb{R}^2$, simulations give $c \approx 0.562$ for aligned squares (e.g. [8]) and $c \approx 0.544$ for discs (e.g. [16], [5]). Using the present results for random sequential sphere packing (Model 2 with fixed radius) we have $c = b\mathbb{E}(K) \approx 0.6282$.

9. Discussion

The nested Monte Carlo design used here seems to have proved its worth, since even with 10^8 trials of a simpler scheme of independent repetitions, sensible estimates of α_k, π_k could not have been obtained for $k > 7$.

A remaining caveat concerns the possible effects of lattice regularity in random number generators, as described by Ripley [14]. We do not have sufficient information about the Cray XMP random number generator's performance to resolve this.

However, the uniformity of the same generator in combination with Marsaglia's method was examined by Dr Mark Durst of the Lawrence Livermore Laboratories, Livermore, California. A test of the uniform spherical distribution of points was performed by delineating 100 equal-area cells on the unit sphere, generating 555 sets of 1000 points on the sphere, and classifying the points over the 100 cells. A χ^2 statistic was computed for the cell counts of each 1000-point set. The observed quantiles of the χ^2_{99} distribution were then tested for uniformity with a χ^2 test on [0, 1] employing eight cells (with cell boundaries 0.01, 0.05, 0.1, 0.5, 0.9, 0.95, and 0.99). The test accepted the hypothesis of uniformity (observed significance level $p = 0.73$).

To test the randomness of pairs of points, 550 sets of 6250 pairs of points were generated, and each point was classified into one of 25 equal-area cells (for a total of 625 cell pairs); a similar analysis to the above was performed, and again, the hypothesis of uniformity was accepted ($p = 0.30$).

Dr Durst's results suggest at least that pairs of generated spheres should have the correct joint distribution. Further tests of the Cray random number generator would be required to put this issue to rest.

Separate tests of the entire simulation program were performed on a Sun-4/60 (Sparcstation-1) using the Wichmann–Hill random number generator [17], [18], [11]; they yielded results very similar to those above.

We are grateful to Professor M. Gordon for his continuing interest and involvement in this work. Many helpful comments were also made by Professor Richard Cowan. We thank the referee for some very helpful criticisms.

References

[1] AKEDA, Y. AND HORI, M. (1976) On random sequential packing in two and three dimensions. *Biometrika* **63**, 361–366.

[2] BLAISDELL, B. E. AND SOLOMON, H. (1970) On random sequential packing in the plane and a conjecture of Palasti. *J. Appl. Prob.* **7**, 667–698.

[3] BLAISDELL, B. E. AND SOLOMON, H. (1982) Random sequential packing in Euclidean spaces of dimensions three and four and a conjecture of Palasti. *J. Appl. Prob.* **19**, 382–390.

[4] COWAN, R. (1984) A model for random packing of disks in the neighbourhood of one disk. *SIAM J. Appl. Math.* **44**, 839–853.

[5] FEDER, J. (1980) Random sequential adsorption. *J. Theor. Biol.* **87**, 237–254.

[6] FINEGOLD, L. AND DONNELL, J. T. (1979) Maximum density of random placing of membrane particles. *Nature* **278**, 443–445.

[7] GOTOH, K. AND FINNEY, J. L. (1974) Statistical geometrical approach to random packing density of equal spheres. *Nature* **252**, 202–205.

[8] JODREY, W. S. AND TORY, E. M. (1980) Random sequential packing in R^n. *J. Statist. Comp. Simul.* **10**, 87–93.

[9] KNUTH, D. (1969) *The Art of Computer Programming. Vol 2. Seminumerical Algorithms.* Addison-Wesley, Reading, Mass.

[10] MARSAGLIA, G. (1972) Choosing a point from the surface of a sphere. *Ann. Math. Statist.* **43**, 645–646.

[11] McLEOD, A. I. (1985) A remark on Algorithm AS 183. *Appl. Statist.* **34**, 198–200.

[12] MILES, R. E. (1971) Random points, sets, and tessellations on the surface of a sphere. *Sankhyā* A **33**, 145–174.

[13] RÉNYI, A. (1958) On a one-dimensional problem concerning place filling (in Hungarian). *Magyar Tud. Akad. Mat. Intézet Közl. (Publ. Math. Inst. Hungar. Acad. Sci.)* **3**, 109–127.

[14] RIPLEY, B. D. (1987) *Statistical Simulation.* Wiley, New York.

[15] STEPHENS, M. A. (1964) The testing of unit vectors for randomness. *J. Amer. Statist. Assoc.* **59**, 160–167.

[16] TORY, E.M., JODREY, W. S. AND PICKARD, D. K. (1983) Simulation of random sequential adsorption: efficient methods and resolution of conflicting results. *J. Theor. Biol.* **102**, 439–445.

[17] WICHMANN, B. A. AND HILL, I.D. (1982) Algorithm AS 183: An efficient an portable pseudo-random number generator. *Appl. Statist.* **31**, 188–190.

[18] WICHMANN, B. A. AND HILL, I. D. (1984) Correction to Algorithm AS 183. *Appl. Statist.* **33**, 123.

[19] WODAK, S. J. AND JANIN, J. (1980) Analytical approximations to the accessible surface area of proteins. *Proc. Nat. Acad. Sci. USA* **77**, 1736–1740.