Combinatorics, Probability and Computing (2003) 12, 301–318. © 2003 Cambridge University Press
DOI: 10.1017/S0963548303005601 Printed in the United Kingdom

Tail Bounds and Expectations for Random Arc Allocation and Applications

PETER SANDERS[1†] and BERTHOLD VÖCKING[2†]

[1] Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany
(e-mail: sanders@mpi-sb.mpg.de)

[2] Department of Computer Science, Universität Dortmund, Informatik 2, Baroper Str. 301, 44221 Dortmund, Germany
(e-mail: voecking@ls2.cs.uni-dortmund.de)

Received 10 March 2002; revised 11 July 2002

Suppose n circular arcs of lengths $\alpha_i \in [0,1]$, $0 \leqslant i < n$, are placed uniformly at random on a unit length circle. We study the *maximum overlap*, *i.e.*, the number of arcs that overlap at the same position of the circle. In particular, we give almost exact tail bounds for this random variable. By applying these tail bounds we can characterize the expected maximum overlap exactly up to constant factors in lower order terms. We illustrate the strength of our results by presenting new performance guarantees for three algorithmic applications: minimizing rotational delays for disks, scheduling accesses to parallel disks, and allocating memory blocks to limit cache interference misses.

1. Introduction

Randomly assigning tasks or data to computational resources has proved to be a universal load balancing technique that can be applied to many algorithmic problems. The model that we study in this paper is motivated by several different applications from memory allocation and disk scheduling. The common theme of these applications is that objects (disk blocks, strings of disk blocks, memory segments) have to be allocated to a contiguous resource that wraps around (disk tracks, striped disks, memory locations mod cache size).

On a higher abstraction level, these problems can be described by an *arc allocation*, *i.e.*, a mapping of a set of n circular arcs to a circle with circumference 1. We represent points on this circle by numbers from the half-open interval $[0,1)$. An arc i with starting position $a_i \in [0,1)$ and length $\alpha_i \in [0,1]$ covers the points in the interval $[a_i, a_i + \alpha_i)$, if $a_i + \alpha_i < 1$, and $[a_i, 1) \cup [0, a_i + \alpha_i - 1)$, otherwise. Figure 1 gives an example.

† Partially supported by the Future and Emerging Technologies programme of the EU under contract number IST-1999-14186 (ALCOM-FT).

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

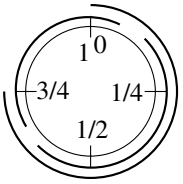


Figure 1. An arc allocation with $n = 3$, $L = 2$, $\alpha = 5/12$, $a_0 = 0$, $a_1 = 1/8$, $a_2 = 2/3$, $\alpha_0 = 3/4$, $\alpha_1 = 1/2$, $\alpha_2 = 5/12$

Let $\alpha = \sum_i \alpha_i / n$ denote the *average arc length*. In a *uniform* arc allocation all arc lengths are identical, $\alpha_i = \alpha$, for $0 \leqslant i < n$. In a *random* arc allocation, the starting points are chosen independently and uniformly at random. Let $L(x)$ denote the number of arcs containing x. Let $L = \sup_{x \in [0,1)} L(x)$ denote the *maximum overlap* of an arc allocation. In this paper, we estimate the expectation $\mathbf{E}[L]$ of the maximum overlap for random arc allocations and derive tail bounds for L.

Random arc allocations have been studied before. For example, the length and number of arcs needed to cover the circle has been investigated [17]. Other work is motivated by the computation of the maximum clique in so-called circular arc graphs [3, 11, 9]. An arc allocation defines a *circular arc graph* with n nodes corresponding to the arcs and edges between nodes if the corresponding arcs overlap. A set of overlapping arcs defines a clique of the circular arc graph. In this terminology, we are studying the size of the maximum overlap clique of a random circular arc graph.

In this paper we present almost exact tail bounds on the maximum overlap for random arc allocation. These tail bounds yield a complete characterization of the expected maximum overlap. Let $\Delta = L - \alpha n$ denote the difference between maximum and average load. We prove tail bounds of the following form. For every $\epsilon > 0$,

$$\mathbf{Pr}[\Delta \geqslant \epsilon \alpha n + k_0] \leqslant k_1 \min\{\alpha^{-1}, n\} \left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}} \right)^{k_2 \alpha n},$$

where $k_0, k_1, k_2 > 0$ denote suitable constants. Exact constants and more details concerning these tail bounds are described in Section 3.

Based on these tail bounds, we are able to describe $\mathbf{E}[\Delta]$ and hence $\mathbf{E}[L]$ almost exactly. Our bounds for these expectations are expressed in terms of Lambert's W function [6]. This function is defined to be the unique positive solution to the equation $\mathrm{W}(x) \cdot \exp(\mathrm{W}(x)) = x$ ($x \geqslant 0$). Of particular interest for us is the function $\exp(\mathrm{W}(x))$. Asymptotically, $\exp(\mathrm{W}(x))$ can be estimated by $x/\ln(x)$ as $\lim_{x \to \infty} \exp(\mathrm{W}(x)) = x/\ln(x)$. We obtain the following upper bound:

$$\mathbf{E}[\Delta] = O\left(\min\left\{ \alpha n \exp\left(\mathrm{W}\left(\frac{\ln(1/\alpha)}{\alpha n} \right) \right), \sqrt{\alpha n \ln\left(\frac{1}{\alpha} \right)} \right\} \right).$$

Recall that $\mathbf{E}[L] = \alpha n + \mathbf{E}[\Delta]$. Let us consider two subcases for which the above estimate

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

simplifies significantly. For example, if arcs are very short, say $\alpha = 1/n$, then $\mathbf{E}[L] = 1 + \mathbf{E}[\Delta] = O\left(\frac{\ln n}{\ln \ln n}\right)$. Furthermore, if $\alpha \geqslant \frac{\ln n}{n}$ then the second term of the minimum dominates and we obtain

$$\mathbf{E}[L] = \alpha n + \mathbf{E}[\Delta] = \alpha n + O\left(\sqrt{\alpha n \ln\left(\frac{1}{\alpha}\right)}\right).$$

These bounds are tight up to constant factors hidden by the Oh-notation. More details concerning the expected overlap are given in Section 4.

Outline. In the following section, we introduce the main technical tools, Chernoff bounds and deviation bounds for random walks, that we need for our analysis. The new tail bounds for random arc allocations are presented in Section 3. From these tail bounds we derive an almost exact characterization of the expected maximum overlap for different values of α and n. These bounds are presented in Section 4. Furthermore, we compare these estimates with simulation results. Finally, in Section 5, we show how our results can be applied to discrete applications from memory allocation and disk scheduling. Additional related work for discrete allocation models (chains into bins) is described in Section 5 along with the applications.

2. Tools

We will frequently apply the following tail bound, which is a slight variant of the well-known Chernoff bound. This tail bound differs from the standard bound, in that it describes the deviation in terms of an upper bound on the expectation rather than the expectation itself. Our analysis depends heavily on this technicality, so we give an explicit proof for this tail bound.

Lemma 2.1. *Let $X_1, \ldots, X_n$ denote independent* 0–1 *random variables. Define $X = \sum X_i$ and let $\mu \geqslant \mathbf{E}[X]$. For every $\epsilon > 0$,*

(a) $\mathbf{Pr}\left[X - \mathbf{E}[X] \geqslant \epsilon\mu\right] \leqslant \left(\dfrac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu}$,

(b) $\mathbf{Pr}\left[\mathbf{E}[X] - X \geqslant \epsilon\mu\right] \leqslant \left(\dfrac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu}$.

Proof. By applying a standard Chernoff bound (see, *e.g.*, [10], equation 5) we immediately obtain

$$\mathbf{Pr}\left[X - \mathbf{E}[X] \geqslant \epsilon\mathbf{E}[X]\right] \leqslant \left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mathbf{E}[X]}.$$

For $\mu = \mathbf{E}[X]$ this bound corresponds directly to the bound stated in (a). For $\mu > \mathbf{E}[X]$,

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

we need some further calculations. Define $\delta = \mu/\mathbf{E}[X] > 1$. Then

$$\begin{aligned}\mathbf{Pr}[X - \mathbf{E}[X] \geqslant \epsilon\mu] &= \mathbf{Pr}[X - \mathbf{E}[X] \geqslant \epsilon\delta\mathbf{E}[X]] \\ &\leqslant \left(\frac{\mathrm{e}^{\epsilon\delta}}{(1+\epsilon\delta)^{1+\epsilon\delta}}\right)^{\mathbf{E}[X]} \\ &\overset{(*)}{\leqslant} \left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu}.\end{aligned}$$

Inequality $(*)$ follows because $\mu = \delta\mathbf{E}[X]$ and $(1+\epsilon)^{(1+\epsilon)\delta} \leqslant (1+\epsilon\delta)^{1+\epsilon\delta}$, for every $\epsilon \geqslant 0$ and $\delta \geqslant 1$. The latter inequality can be shown by taking the logarithm on both sides, leading to the following equivalent inequality: $(1+\epsilon)\delta\ln(1+\epsilon) \leqslant (1+\epsilon\delta)\ln(1+\epsilon\delta)$. Observe that both of these terms are equal to zero for $\epsilon = 0$. Next we compute the derivative of the two terms with respect to ϵ. The expressions that we obtain are $\delta + \delta\ln(1+\epsilon)$ and $\delta + \delta\ln(1+\epsilon\delta)$ for the left- and right-hand side, respectively. Obviously, the derivative of the term on the left-hand side is bounded from above by the derivative of the term on the right-hand side. Thus the same relationship holds for the original terms, and hence bound (a) is shown. Bound (b) follows completely analogously, starting with the Chernoff bound given in [10] (equation 7 instead of equation 5). □

We will combine tail bounds and random walks. In particular, we will apply the following lemma bounding the maximum of a random walk.

Lemma 2.2. *Consider a random walk $S_0, S_1, S_2, \ldots$ on the integer numbers with $S_0 = 0$ and, for $i \geqslant 1$, $\mathbf{Pr}[S_i = S_{i-1} + 1] = \mathbf{Pr}[S_i = S_{i-1} - 1] = \frac{1}{2}$. For every $m \geqslant 0$, $\epsilon > 0$,*

$$\mathbf{Pr}[\exists i \in \{1,\ldots,m\} : S_i > \epsilon m] \leqslant 2\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{m/2}.$$

Proof. Theorem III.7.1 from the book by Feller [7] gives

$$\mathbf{Pr}[\exists i \in \{1,\ldots,m\} : S_i > r] = \sum_{k=r+1}^{\infty} \mathbf{Pr}[S_m = k] + \mathbf{Pr}[S_m = k+1] \leqslant 2\,\mathbf{Pr}[S_m > r],$$

for every $r \geqslant 0$. Thus, it is sufficient to investigate S_m rather than $S_1, \ldots, S_m$. Now observe that the random variable $X_m = S_m/2 + m/2$ follows the binomial distribution $B\left(m, \frac{1}{2}\right)$. In particular, $\mathbf{E}[X_m] = \frac{m}{2}$, so that the tail bound from Lemma 2.1(a) gives

$$\mathbf{Pr}[S_m > \epsilon m] \leqslant \mathbf{Pr}\left[X_m > (1+\epsilon)\frac{m}{2}\right] \leqslant \left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{m/2}.$$

Combining these two inequalities gives the lemma. □

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

3. Tail bounds

3.1. A tail bound for relatively short arcs

In this section, we prove a first tail bound that holds for general arc lengths but is best suited for arcs of small length, *i.e.*, $\alpha = O(\frac{\ln n}{n})$. Recall, for $x \in [0,1)$, $L(x)$ denotes the overlap at position x and $\Delta(x) = L(x) - \mathbf{E}[L(x)] = L(x) - \alpha n$. Furthermore, $\Delta = \sup_{x\in[0,1)} \Delta(x)$.

Theorem 3.1. *Suppose n arcs of average length α are placed at random onto the unit circle. Let $\mu \geqslant \alpha n$. For every $\epsilon > 0$,*

$$\mathbf{Pr}[\Delta \geqslant \epsilon\mu + 1] \leqslant n\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu}.$$

Proof. The proof idea is to estimate the maximal overlap at exactly n positions on the circle, namely at the left endpoints of the arcs. Fix an arbitrary arc i and let a_i denote the position of the left endpoint of this arc. Taking into account the contribution of arc i to the overlap at position a_i, we obtain

$$\mathbf{E}[L(a_i)] \leqslant \alpha n + 1 \leqslant \mu + 1.$$

Now, as all other arcs are placed independently from arc i, we can apply the tail bound from Lemma 2.1(a), which gives

$$\mathbf{Pr}[\Delta(a_i) \geqslant \epsilon\mu + 1] \leqslant \left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu}.$$

Clearly, the maximum overlap at the left endpoints is an upper bound on the maximum overlap over all positions on the cycle, because the load can increase only at those positions that are left endpoints of arcs when traversing the circle in clockwise fashion. Therefore, we can conclude

$$\mathbf{Pr}[\Delta \geqslant \epsilon\mu + 1] \leqslant \sum_{i=1}^{n} \mathbf{Pr}[\Delta(a_i) \geqslant \epsilon\mu + 1] \leqslant n\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu},$$

which proves the theorem. □

3.2. A tail bound for uniform, relatively long arcs

In the preceding section, we proved a tail bound on L by applying a Chernoff bound at n positions on the circle. This leads to the factor n in front of the term known from the standard Chernoff tail bound. Do we really need to lose this factor? The following theorem answers this question negatively. It gives a tail bound that loses only a factor $\frac{1}{\alpha}$ instead of n. This bound is stronger than the bound given in Theorem 3.1 if arcs are relatively long, in particular in the case $\alpha \gg \frac{1}{n}$. We will prove in Section 4.2 that this tail bound is tight in the sense that the factor $\frac{1}{\alpha}$ in front of the standard Chernoff bound term is unavoidable.

In this section, we assume that all arc lengths are identical. A more complicated calculation that loses some large constant factors, and hence is less practical, but holds for

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

non-uniform arcs, is presented in Section 3.4. (The analysis for non-uniform arcs is based on the analysis for uniform arcs.) Furthermore, we assume that the arc length is at most $\frac{1}{2}$. In Section 3.3, we will show that a bound for longer arcs can be obtained by a simple symmetry argument.

Theorem 3.2. *Suppose n arcs of length $\alpha \leqslant \frac{1}{2}$ are placed at random onto the unit circle. Assume $\mu \geqslant \alpha n$. For every $\epsilon > 0$,*

$$\mathbf{Pr}[\Delta \geqslant 5\epsilon\mu] \leqslant \frac{6}{\alpha}\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu}.$$

Proof. Define $\kappa = \lceil 1/\alpha \rceil \geqslant 2$. Let $x_0, \ldots, x_{\kappa-1}$ denote κ points on the circle that decompose the circle into κ intervals $X_i = [x_i, x_{i+1})$ of identical length $\frac{1}{\kappa} \approx \alpha$. (Here and in the following $i+1$ abbreviates $(i+1) \bmod \kappa$.) Observe that every arc has at most one endpoint in each of these intervals. Define $\Delta_i = L(x_i) - \alpha n$ and $\Delta_i' = \sup_{x \in X_i}(L(x) - L(x_i))$. In this way, the maximum overlap in interval X_i is exactly $\alpha n + \Delta_i + \Delta_i'$.

Claim 3.3. *For every $\epsilon > 0$,*

$$\mathbf{Pr}[\Delta_i \geqslant \epsilon\mu] \leqslant \left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu}.$$

Claim 3.4. *For every $\epsilon > 0$,*

$$\mathbf{Pr}[\Delta_i' \geqslant \epsilon(2+2\epsilon)\mu \mid \Delta_i + \Delta_{i+1} < 2\epsilon\mu] \leqslant 2\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{(1+\epsilon)\mu}.$$

Before we come to the proofs of these claims, let us show that they imply the theorem. First, suppose $\epsilon \geqslant 1$. The maximum overlap in interval X_i is bounded above by $L(x_i) + L(x_{i+1}) = 2\alpha n + \Delta_i + \Delta_{i+1}$, because every arc overlapping with interval X_i covers x_i or x_{i+1}. This implies the inequality $\Delta \leqslant \max_i\{\alpha n + \Delta_i + \Delta_{i+1}\}$ (a1). Furthermore, we have $\alpha n \leqslant \epsilon\mu$ (a2) and $\kappa = \lceil \frac{1}{\alpha} \rceil$ (a3). Thus we obtain

$$\begin{aligned}
\mathbf{Pr}[\Delta \geqslant 3\epsilon\mu] &\overset{\text{(a1)}}{\leqslant} \mathbf{Pr}[\exists_{i=0}^{\kappa-1} : \alpha n + \Delta_i + \Delta_{i+1} \geqslant 3\epsilon\mu] \\
&\overset{\text{(a2)}}{\leqslant} \mathbf{Pr}[\exists_{i=0}^{\kappa-1} : \Delta_i \geqslant \epsilon\mu] \\
&\overset{\text{(Claim 3.3)}}{\leqslant} \kappa\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu} \\
&\overset{\text{(a3)}}{\leqslant} \frac{2}{\alpha}\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu}.
\end{aligned}$$

Now assume that $\epsilon < 1$. Observe that this implies $\epsilon \geqslant \epsilon(3+2\epsilon)/5$ (b1). Recall that

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

$\Delta = \max_i\{\Delta_i + \Delta'_i\}$ (b2). Applying these facts, we obtain

$$\begin{aligned}
\mathbf{Pr}[\Delta \geqslant 5\epsilon\mu] &\overset{\text{(b1)}}{\leqslant} \mathbf{Pr}[\Delta \geqslant \epsilon(3+2\epsilon)\mu] \\
&\overset{\text{(b2)}}{\leqslant} \mathbf{Pr}[\exists_{i=0}^{\kappa-1} : \Delta_i \geqslant \epsilon\mu \ \vee \ \exists_{i=0}^{\kappa-1} : \Delta'_i \geqslant \epsilon(2+2\epsilon)\mu] \\
&\overset{\text{(b3)}}{\leqslant} \mathbf{Pr}[\exists_{i=0}^{\kappa-1} : \Delta_i \geqslant \epsilon\mu \ \vee \ \exists_{i=1}^{\kappa-1} : \Delta'_i \geqslant \epsilon(2+2\epsilon)\mu] \\
&\qquad + \mathbf{Pr}[\Delta'_0 \geqslant \epsilon(2+2\epsilon)\mu \mid \Delta_0 + \Delta_1 < 2\epsilon\mu] \\
&\ \ \vdots \\
&\overset{\text{(b3)}}{\leqslant} \mathbf{Pr}[\exists_{i=0}^{\kappa-1} : \Delta_i \geqslant \epsilon\mu] \\
&\qquad + \sum_{i=0}^{\kappa-1} \mathbf{Pr}[\Delta'_i \geqslant \epsilon(2+2\epsilon)\mu \mid \Delta_i + \Delta_{i+1} < 2\epsilon\mu] \\
&\overset{\text{(Claims 3.3, 3.4)}}{\leqslant} \kappa\left(\frac{\mathrm{e}^\epsilon}{(1+\epsilon)^{1+\epsilon}}\right)^\mu + 2\kappa\left(\frac{\mathrm{e}^\epsilon}{(1+\epsilon)^{1+\epsilon}}\right)^{(1+\epsilon)\mu} \\
&\leqslant \frac{6}{\alpha}\left(\frac{\mathrm{e}^\epsilon}{(1+\epsilon)^{1+\epsilon}}\right)^\mu.
\end{aligned}$$

For (b3) we iteratively apply the following basic fact from probability theory: For events X, X', and Y with $X' \subseteq X$ we have $\mathbf{Pr}[X \vee Y] \leqslant \mathbf{Pr}[X] + \mathbf{Pr}[Y \setminus X'] \leqslant \mathbf{Pr}[X] + \mathbf{Pr}[Y \mid \neg X']$. In the jth iteration ($1 \leqslant j \leqslant \kappa$) of (b3), we set

$$\begin{aligned}
X &= \left(\exists_{i=0}^{\kappa-1} : \Delta_i \geqslant \epsilon\mu \ \vee \ \exists_{i=j}^{\kappa-1} : \Delta'_i \geqslant \epsilon(2+2\epsilon)\mu\right), \\
X' &= (\Delta_{j-1} + \Delta_j \geqslant 2\epsilon\mu) \subseteq X, \\
Y &= (\Delta'_{j-1} \geqslant \epsilon(2+2\epsilon)\mu).
\end{aligned}$$

Thus the two claims imply the theorem, and it only remains to prove the claims. Observe that $\mathbf{E}[L(x_i)] \leqslant \mu$, so that the bound in Claim 3.3 follows directly from the tail bound in Lemma 2.1(a). In the following we present the proof for Claim 3.4.

We will estimate Δ'_i by investigating the following random walk. Recall that each arc has at most one endpoint in interval X_i. Let m denote the number of those arcs that have an endpoint in X_i. As we condition on $\Delta_i + \Delta_{i+1} < 2\epsilon\mu$, we can assume

$$m \leqslant L(x_i) + L(x_{i+1}) \leqslant 2\mu + \Delta_i + \Delta_{i+1} \leqslant (2+2\epsilon)\mu.$$

For the time being, assume that m is fixed. Let $y_1, \ldots, y_m$ denote the endpoints in X_i, sorted from left to right. If we are given these endpoints without further information, then the *orientation* of the corresponding arcs (*i.e.*, whether y_j is a left or right endpoint) defines a random experiment that can be described in terms of binary random variables, as follows. Let $s_1, \ldots, s_m$ denote random variables with

$$s_j = \begin{cases} +1 \text{ if } y_j \text{ is a left endpoint, and} \\ -1 \text{ if } y_j \text{ is a right endpoint.} \end{cases}$$

The only assumption that we made about the allocation of the arcs is $\Delta_i + \Delta_{i+1} \leqslant 2\epsilon\mu$.

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

As this assumption does not affect the arc's orientation, the variables $s_1, \ldots, s_m$ are independent and $\mathbf{Pr}[s_j = 1] = \mathbf{Pr}[s_j = -1] = \frac{1}{2}$, for $1 \leqslant j \leqslant m$.

Now let us define $S_j = \sum_{k=1}^{j} s_j$, for $0 \leqslant j \leqslant m$. Notice that the sequence $S_0, S_1, \ldots, S_m$ corresponds to a random walk in which a particle starts at position 0 and goes up or down by one position in each step with probability $\frac{1}{2}$ each, that is, $\Delta_i' = \max_{0 \leqslant i \leqslant m}(S_i)$. Hence, we can estimate Δ_i' by analysing this random walk. Clearly, Δ_i' is monotonically increasing in m. Therefore, the worst-case choice for m is $m = (2 + 2\epsilon)\mu$. (Recall that we condition on $\Delta_i + \Delta_{i+1} \leqslant 2\epsilon\mu$ so that the maximum number of endpoints in interval X_i is $(2 + 2\epsilon)\mu$.) Now we apply Lemma 2.2 and we obtain

$$\mathbf{Pr}[\Delta_i' > \epsilon(2 + 2\epsilon)\mu \mid \Delta_i + \Delta_{i+1} < 2\epsilon\mu] \leqslant 2\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{(1+\epsilon)\mu},$$

which proves Claim 3.4. This completes the proof of Theorem 3.2. □

3.3. A symmetry argument

In this section, we describe a general method for dealing with arcs that are longer than $\frac{1}{2}$. Instead of considering these arcs directly, we investigate the corresponding holes. These holes are defined as follows. Suppose an arc covers the interval $[a, b)$. Then the corresponding *hole* covers the interval $[b, a)$. Thus if an arc has length α then the corresponding hole has length $1 - \alpha$.

Observe that the holes are placed randomly on the circle in the same way as the arcs. Here we are interested, however, in the minimum overlap instead of the maximum overlap. Let $\bar{\alpha} = 1 - \alpha$ denote the average length of the holes and $\bar{L}(x)$ the number of holes that cover position x. Then we come to the following identity that holds for every point x, $L(x) - \alpha n = \bar{\alpha} n - \bar{L}(x)$. Consequently,

$$\Delta = \max_x(L(x)) - \alpha n = \bar{\alpha} n - \min_x(\bar{L}(x)).$$

Moreover, $\bar{\alpha} n - \min_x(\bar{L}(x))$ can be estimated using a combination of tail bounds and random walks analogous to the analyses we presented in the previous sections: that is, instead of the tail bound of Lemma 2.1(a) we apply the symmetric tail bound of Lemma 2.1(b) and we consider the minimum rather than the maximum of the random walks. The rest of the analysis is completely analogous to the proofs of Theorems 3.1 and 3.2. In this way, we obtain the following tail bounds.

Theorem 3.5. *Suppose n arcs of average length α are placed at random onto the unit circle. Let $\bar{\mu} \geqslant (1 - \alpha)n$. For every $\epsilon > 0$,*

$$\mathbf{Pr}[\Delta \geqslant \epsilon\bar{\mu} + 1] \leqslant n\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\bar{\mu}}.$$

Theorem 3.6. *Suppose n arcs of length $\alpha \leqslant \frac{1}{2}$ are placed at random onto the unit circle. Assume $\bar{\mu} \geqslant (1 - \alpha)n$. For every $\epsilon > 0$,*

$$\mathbf{Pr}[\Delta \geqslant 5\epsilon\bar{\mu}] \leqslant \frac{6}{\alpha}\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\bar{\mu}}.$$

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

Observe that these tail bounds are completely symmetrical to the tail bounds given in Theorems 3.1 and 3.2. If $\alpha > \frac{1}{2}$ then the bounds w.r.t. $\bar{\mu}$ are stronger; otherwise the corresponding bounds w.r.t. μ give the better results. The fact that these tail bounds are completely symmetric does not imply, however, that the actual behaviour of the random variables is completely symmetrical, too. In fact, our simulations in Section 4 show small asymmetries that are not covered by our analysis.

3.4. Tail bound for non-uniform arc lengths

The tail bounds in Theorems 3.2 and 3.6 hold only for uniform arcs. Observe that the proof of Theorem 3.2 fails in the case of non-uniform arcs, as it can no longer be guaranteed that arcs have at most one endpoint in each interval. Nevertheless, we will be able to use the results obtained for uniform arc lengths as tools in the following analysis for non-uniform arc lengths. The following theorem generalizes the tail bounds in Theorems 3.2 and 3.6 towards non-uniform arc lengths.

Theorem 3.7. *Suppose n arcs of average length α are placed at random onto the unit circle. Assume $\mu \geqslant \alpha n$ is an upper bound on the average overlap of the arcs. Then, for every $\epsilon > 0$,*

$$\mathbf{Pr}[\Delta \geqslant \epsilon\mu + 2] \leqslant \frac{k_1}{\alpha}\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{k_2\mu},$$

for suitable constants $k_1, k_2 > 0$.

Remark. Following the arguments from Section 3.3, the same tail bound can be shown w.r.t. the holes, *i.e.*, μ can be replaced by $\bar{\mu} \geqslant (1-\alpha)n$.

Proof. We distinguish two cases depending on the value of ϵ.

Case 1: Suppose $\epsilon \geqslant \frac{1}{8}$.
In this case, Theorem 3.1 implies the following bound:

$$\mathbf{Pr}[\Delta \geqslant \epsilon\mu + 1] \leqslant n\left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu}.$$

Essentially, we have to replace the factor n by $\frac{1}{\alpha}$. Let us define $f(\epsilon,\mu) = \left(\frac{\mathrm{e}^{\epsilon}}{(1+\epsilon)^{1+\epsilon}}\right)^{\mu/3}$. Clearly, if $n \leqslant f(\epsilon,\mu)^{-2}$, then the bound above can be estimated by $\mathbf{Pr}[\Delta \geqslant \epsilon\mu + 1] \leqslant n f(\epsilon,\mu)^3 \leqslant f(\epsilon,\mu) \leqslant \frac{1}{\alpha} f(\epsilon,\mu)$, which corresponds to the tail bound that we have to prove. Now assume $n > f(\epsilon,\mu)^{-2}$. Solving this inequality for μ gives $\mu < 200 \ln n$. Now applying $\mu \geqslant \alpha n$ gives $\frac{200}{\alpha} > \frac{n}{\ln n}$. Consequently,

$$\frac{200}{\alpha} f(\epsilon,\mu) > \frac{n}{\ln n} f(\epsilon,\mu) \geqslant \sqrt{n} f(\epsilon,\mu) p \geqslant 1,$$

which means that the tail bound stated in the theorem is trivially true for $k_1 = 200$ and $k_2 = \frac{1}{3}$.

Case 2: Now suppose $\epsilon < \frac{1}{8}$.
In this case, we partition the set of arcs into subsets according to their length. For a subset

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

A of the arcs, let L_A denote the maximum overlap among the arcs in A. Furthermore, let α_A denote the average length of the arcs in A, and define $\Delta_A = L_A - \alpha_A |A|$.

First we bound the deviation due to those arcs that have length in the interval $(\alpha, 1-\alpha)$. Let T denote the set of these arcs. The arcs in T can be treated completely analogously to the uniform arcs in the proof of Theorem 3.2, as the length restriction ensures the crucial property that each interval X_i contains at most one endpoint of each arc in T. Clearly, $\mu \geqslant \alpha_T |T|$. Hence, applying the bound stated in Theorem 3.2 gives

$$\mathbf{Pr}[\Delta_T \geqslant 5\epsilon\mu] \leqslant \frac{6}{\alpha_T}\left(\frac{e^\epsilon}{(1+\epsilon)^{1+\epsilon}}\right)^\mu \leqslant \frac{6}{\alpha}\left(\frac{e^\epsilon}{(1+\epsilon)^{1+\epsilon}}\right)^\mu$$

because $\alpha_T \geqslant \alpha$, as all arcs in T have length at least α. Thus it remains only to estimate the deviation of the arcs not in T, namely the arcs of length at most α and the arcs of length at least $1-\alpha$. Let S denote the set of arcs of length at most α. The arcs of length at least $1 - 1/\alpha$ can be treated symmetrically by considering holes instead of arcs as explained in Section 3.3. Therefore, we need only consider the arcs in S.

Claim 3.8. *For a set S with at most n arcs of length at most α,*

$$\mathbf{Pr}[\Delta_S \geqslant \epsilon\alpha n + 1] \leqslant \frac{k_1}{\alpha}\exp(-k_2\epsilon^2\alpha n),$$

for every $\epsilon \in (0, \frac{1}{8}]$ and suitable, absolute constants $k_1, k_2 > 0$.

The following approximation is typically used in the context of Chernoff bounds (*cf. e.g.*, [10, equation 3]). For every $\epsilon \leqslant 1$,

$$\exp(-\epsilon^2/2) \leqslant \frac{e^\epsilon}{(1+\epsilon)^{1+\epsilon}} \leqslant \exp(-\epsilon^2/3).$$

From this approximation we can conclude that Claim 3.8 implies Theorem 3.7, so it remains only to prove this claim.

We prove Claim 3.8 using a recursive approach. The recursion is on the maximum arc length. We will see that decreasing the maximum arc length implicitly allows us to increase the value of ϵ: that is, in each step of our recursion, the value of ϵ is increased by the factor 4. In this way, the recursion is guaranteed to terminate in Case 1. Observe that our analysis for this case yields Claim 3.8 for $\epsilon \geqslant \frac{1}{8}$ by applying $\frac{e^\epsilon}{(1+\epsilon)^{1+\epsilon}} \leqslant \exp(-\epsilon^2/3)$.

The recursion works as follows. Let R denote the subset of arcs of length at most $\beta = \alpha/8$. Furthermore, let $\bar{R} = S \backslash R$, that is, $\bar{R}$ contains the arcs with length from the interval $(\beta, \alpha]$. As recursion hypothesis, we can assume that Claim 3.8 holds for arcs of length at most β, that is, the claim holds for the arcs in R. It remains to check that the claim still holds when additionally taking into account the arcs in $\bar{R}$. More precisely, by applying our hypothesis, we obtain

$$\mathbf{Pr}[\Delta_R \geqslant \gamma\beta n + 1] \leqslant \frac{k_1}{\beta}\exp(-k_2\gamma^2\beta n),$$

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

for every $\gamma > 0$. We choose $\gamma = \frac{\epsilon\alpha}{2\beta} = 4\epsilon$. Hence,

$$\begin{aligned}\mathbf{Pr}\left[\Delta_R \geqslant \frac{\epsilon\alpha n}{2} + 1\right] &\leqslant \mathbf{Pr}\left[\Delta_R \geqslant \left(\frac{\epsilon\alpha}{2\beta}\right)\beta n + 1\right] \\ &\leqslant \frac{k_1}{\beta}\exp\left(-k_2\left(\frac{\epsilon\alpha}{2\beta}\right)^2\beta n\right) \\ &\leqslant \frac{8k_1}{\alpha}\exp(-2k_2\epsilon^2\alpha n) \\ &\leqslant \frac{k_1}{2\alpha}\exp(-k_2\epsilon^2\alpha n).\end{aligned}$$

The last inequality follows from $\exp(-2x) \leqslant \frac{1}{16}\exp(-x)$, for every $x \geqslant 3$. Here we additionally assume $k_1 \geqslant \mathrm{e}^3$ and $k_2\epsilon^2\alpha n \geqslant \ln(k_1) \geqslant 3$. Observe that we can make the last assumption w.l.o.g., as $k_2\epsilon^2\alpha n < \ln(k_1)$ would imply that the right-hand side of the bound in Claim 3.8 exceeds the value of 1, so that the stated probability bound would hold trivially.

For the arcs in $\bar{R}$ we have upper and lower bounds on the arc length deviating only by a constant factor. Hence we can use the same techniques as in the proof of Theorem 3.2 in order to estimate $\Delta_{\bar{R}}$. In this way, we can show

$$\mathbf{Pr}\left[\Delta_{\bar{R}} \geqslant \frac{\epsilon\alpha n}{2}\right] \leqslant \frac{k_3}{\alpha}\exp\left(-k_4\left(\frac{\epsilon}{2}\right)^2\alpha n\right),$$

for suitable constants $k_3, k_4 > 0$. Finally, we choose k_1 and k_2 such that $k_1 \geqslant 2k_3$ and $k_2 \leqslant k_4/4$. In this way, we obtain

$$\begin{aligned}\mathbf{Pr}\left[\Delta_S \geqslant \epsilon\alpha n + 1\right] &\leqslant \mathbf{Pr}\left[\Delta_R + \Delta_{\bar{R}} \geqslant \epsilon\alpha n + 1\right] \\ &\leqslant \frac{k_1}{2\alpha}\exp(-k_2\epsilon^2\alpha n) + \frac{k_3}{\alpha}\exp\left(-k_4\left(\frac{\epsilon}{2}\right)^2\alpha n\right) \\ &\leqslant \frac{k_1}{\alpha}\exp(-k_2\epsilon^2\alpha n),\end{aligned}$$

which completes the proof of Theorem 3.7. □

4. Expectations

4.1. A characterization of the expected maximum overlap

Applying our tail bounds, we can derive the following bounds on the expected maximum overlap.

(A) If $\alpha \leqslant \frac{\ln n}{n}$ then $\mathbf{E}[L] = O\left(\alpha n \exp\left(\mathrm{W}\left(\frac{\ln(1/\alpha)}{\alpha n}\right)\right)\right)$.

(B) For $\alpha \in \left[\frac{\ln n}{n}, \frac{1}{2}\right]$, $\mathbf{E}[L] = \alpha n + O\left(\sqrt{\alpha n \ln\left(\frac{1}{\alpha}\right)}\right)$.

(C) For $\alpha \in \left[\frac{1}{2}, 1 - \frac{\ln n}{n}\right]$, $\mathbf{E}[L] = \alpha n + O\left(\sqrt{(1-\alpha)n \ln\left(\frac{1}{1-\alpha}\right)}\right)$.

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

(D) If $\alpha \geqslant 1 - \frac{\ln n}{n}$ then $\mathbf{E}[L] = O\left((1-\alpha)n \exp\left(\mathrm{W}\left(\frac{\ln(1/(1-\alpha))}{(1-\alpha)n}\right)\right)\right)$.

In fact, each of these bounds holds for general $\alpha \in [0,1]$. The intervals specify only for which values of α these bounds are best suited. Observe that the cases (D) and (C) are symmetric to the cases (A) and (B). The bound in (A) follows from the tail bound in Theorem 3.1, the bound in (B) follows from the tail bound in Theorem 3.7, and (D) and (C) follow from the symmetric counterparts of these tail bounds w.r.t. $\bar{\mu}$.

As an example of how to derive these estimates from the corresponding tail bounds, let us formally prove the bound on $\mathbf{E}[L]$ given in (A). From the tail bound in Theorem 3.1 we conclude

$$\mathbf{Pr}[L \geqslant q\alpha n + 1] \leqslant n\left(\frac{\mathrm{e}}{q}\right)^{q\alpha n},$$

for every $q > 1$. Next we substitute the right-hand side of this tail bound by $\exp(-t)$ ($t \geqslant 0$), that is, we set

$$n\left(\frac{\mathrm{e}}{q}\right)^{q\alpha n} = \exp(-t)$$

and solve for q. In this way, Lambert's W function comes into play, as $\frac{q}{\mathrm{e}} \ln\left(\frac{q}{e}\right) = \frac{t+\ln n}{\mathrm{e}\alpha n}$ gives $\ln\left(\frac{q}{e}\right) = \mathrm{W}\left(\frac{t+\ln n}{\mathrm{e}\alpha n}\right)$ so that

$$q = \mathrm{e}\exp\left(\mathrm{W}\left(\frac{t+\ln n}{\mathrm{e}\alpha n}\right)\right).$$

The function $\exp(\mathrm{W}(\cdot))$ is concave and $\exp(\mathrm{W}(0)) = 1$ so that $\exp(\mathrm{W}(x+y)) \leqslant \exp(\mathrm{W}(x)) + \exp(\mathrm{W}(y)) + 1$, for every $x, y \geqslant 0$. Furthermore, $\exp(\mathrm{W}(x)) \leqslant x + 1$. Applying these inequalities gives

$$\begin{aligned}
q\alpha n + 1 &= \mathrm{e}\exp\left(\mathrm{W}\left(\frac{t+\ln n}{\mathrm{e}\alpha n}\right)\right)\alpha n + 1 \\
&\leqslant \mathrm{e}\left(\exp\left(\mathrm{W}\left(\frac{\ln n}{\mathrm{e}\alpha n}\right)\right) + \exp\left(\mathrm{W}\left(\frac{t}{\mathrm{e}\alpha n}\right)\right) + 1\right)\alpha n + 1 \\
&\leqslant \mathrm{e}\left(\exp\left(\mathrm{W}\left(\frac{\ln n}{\mathrm{e}\alpha n}\right)\right) + \frac{t}{\mathrm{e}\alpha n} + 2\right)\alpha n + 1 \\
&\leqslant \mathrm{e}\left(\exp\left(\mathrm{W}\left(\frac{\ln n}{\mathrm{e}\alpha n}\right)\right) + 2\right)\alpha n + t + 1.
\end{aligned}$$

Now, substituting these terms back into our tail bound gives

$$\mathbf{Pr}\left[L \geqslant \mathrm{e}\left(\exp\left(\mathrm{W}\left(\frac{\ln n}{\mathrm{e}\alpha n}\right)\right) + 2\right)\alpha n + t + 1\right] \leqslant \exp(-t).$$

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

As a consequence,

$$\begin{aligned}\mathbf{E}[L] &\leqslant \mathrm{e}\left(\exp\left(\mathrm{W}\left(\frac{\ln n}{\mathrm{e}\alpha n}\right)\right)+2\right)\alpha n+\left(\sum_{t=0}^{\infty}\exp(-t)\right)+1\\ &\leqslant \mathrm{e}\left(\exp\left(\mathrm{W}\left(\frac{\ln n}{\mathrm{e}\alpha n}\right)\right)+2\right)\alpha n+3\\ &\overset{(*)}{\leqslant} \mathrm{e}\left(2\exp\left(\mathrm{W}\left(\frac{2\ln(1/\alpha)}{\mathrm{e}\alpha n}\right)\right)+2\right)\alpha n+3\\ &\overset{(**)}{\leqslant} 8\mathrm{e}\exp\left(\mathrm{W}\left(\frac{\ln n}{\mathrm{e}\alpha n}\right)\right)\alpha n+3.\end{aligned}$$

Inequality (*) requires some further explanation. Suppose $1/\alpha \geqslant \sqrt{n}$. Then the inequality holds because $\ln n = 2\ln\sqrt{n} \leqslant 2\ln(1/\alpha)$. Now assume $1/\alpha < \sqrt{n}$. In this case, the term $\ln n/\mathrm{e}\alpha n$ is smaller than 1, such that $\exp\big(\mathrm{W}\big(\ln n/\mathrm{e}\alpha n\big)\big) \leqslant \exp(\mathrm{W}(1)) \leqslant 2 \leqslant 2\exp(\mathrm{W}(x))$, for every $x \geqslant 0$. Furthermore, inequality (**) follows by applying first $\exp(\mathrm{W}(2x)) \leqslant 2\exp(\mathrm{W}(x))+1$ and then $\exp(\mathrm{W}(x)) \geqslant 1$.

This completes the proof for the bound in (A). The bounds in (B), (C) and (D) can be proved in basically the same way.

4.2. A lower bound in terms of balls-into-bins

In this section, we show that our estimates for $\mathbf{E}[L]$ are essentially tight by relating them to well-known estimates for balls-and-bins processes. For simplicity in notation, let us assume that $1/\alpha$ is integral. We relate the arc allocation process directly to a balls-into-bins process in which $N = n$ balls are thrown uniformly at random into $M = 1/\alpha$ bins.

We describe a simple coupling between these two processes. Identify $M = 1/\alpha$ points on the circle at equal distance. Observe that each allocated arc can overlap with at most one of these points. Now each of these points corresponds to a bin, and each arc corresponds to a ball. Then obviously the maximum occupancy in the balls-and-bins process cannot exceed the maximum overlap of the random arc allocation. Let us summarize this simple observation formally.

Observation 4.1. *The maximum overlap for random arc allocation with n arcs of length α stochastically dominates the maximum occupancy for a balls-into-bins process with $N = n$ balls and $M = \lfloor 1/\alpha \rfloor$ bins.*

Observe that all estimates given in the characterization for $\mathbf{E}[L]$ match the standard estimates for balls-into-bins processes with $N = n$ balls and $M = 1/\alpha$ bins. In fact, even our estimates for the lower-order term $\mathbf{E}[\Delta]$ are tight up to constant factors hidden in the Oh-notation.

4.3. Simulations

Figure 2 shows simulations for three different values of n and compares them with our analytic bound from the case (B) and (C) of our characterization for $\mathbf{E}[\Delta]$. Merging the

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

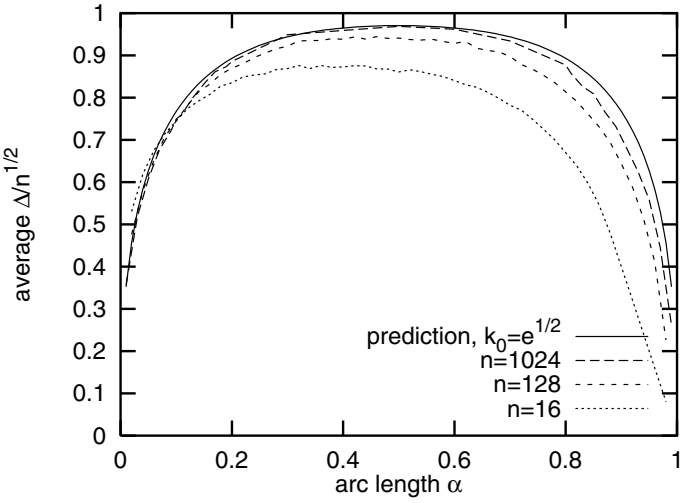


Figure 2. Comparison of the theoretical prediction $\Delta = \sqrt{e\,n \cdot \alpha(1-\alpha)\ln(1/(\alpha(1-\alpha)))}$ with simulations for different n. The measurements are averages of 10000 repetitions for $n \leqslant 128$ and 1000 repetitions for $n = 1024$

bounds in these two cases, we obtain the following elegant estimate:

$$\mathbf{E}[\Delta] \approx \sqrt{\alpha(1-\alpha)n \ln\left(\frac{1}{\alpha(1-\alpha)}\right)}.$$

(The term $\alpha(1-\alpha)$ approximates both the terms α and $1-\alpha$ within a factor of two.) In fact, the measured and predicted curves are quite close when we choose an appropriate constant factor (namely $\sqrt{e}$) in front of this estimate. Even for $n = 16$, where a significant influence of lower order terms can be expected, Δ is fairly well approximated. An interesting phenomenon is that the measured graphs for $\Delta(\alpha)$ are not completely symmetric around $\alpha = 1/2$. This asymmetry is not reflected by our theoretical analysis. The applied Chernoff bounds from Lemma 2.1(a) and (b) are symmetric, although the actual behaviour of 0–1 random variables above and below the mean is slightly asymmetric. In fact, the deviation of holes below the mean can be expected to be slightly less than the deviation of arcs above the mean, which explains the asymmetry that can be observed in the experiments.

5. Applications

Typical applications in computer science require a discrete variant of arc allocation, where the circle is subdivided into M equal bins and where arc endpoints are integral. We want to point out that our proof techniques, and hence our results, directly transfer to this discrete model. Observe that discrete arc allocation is equivalent to the following *chains-into-bins* problem. $N = \sum_i \alpha_i$ balls connected into n chains are allocated to M bins arranged in a circle. A chain is allocated by throwing its first ball into a bin chosen independently and uniformly at random, and by putting its remaining balls into adjacent bins in a round

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

robin fashion. Translating our results into this notation we obtain the following estimates for the *maximum occupancy* $L^{(\mathrm{cb})}$, *i.e.*, the number of balls in the fullest bin:

$$\mathbf{E}\left[L^{(\mathrm{cb})}\right] = \frac{N}{M} + O\left(\min\left\{\frac{N}{M}\exp\left(\mathrm{W}\left(\frac{\ln\left(M\frac{n}{N}\right)}{N/M}\right)\right), \sqrt{\frac{N}{M}\ln\left(M\frac{n}{N}\right)}\right\}\right).$$

Observe that, for $N = n$, these bounds correspond to well-known estimates for maximum occupancy when throwing N balls into M bins: see, *e.g.*, [12, 19, 14]. For $N > n$, our results improve significantly on previous results by Barve, Grove and Vitter [2], who only show that several tail bounds for the case $N = n$ can also be applied for $N > n$, whereas our tail bounds are much stronger for $N \gg n$.

We now present three examples of applications where randomization is used to avoid bad worst case behaviour. The performance of each of these applications corresponds directly to the maximum overlap of a randomized arc allocation. Therefore, our bounds can be used to predict the expected performance of such an application almost exactly.

5.1. Disks and drums

One of our main incentives for studying arc allocations was the desire to find disk scheduling algorithms that take rotational delays into account. For example, consider a database query selecting the set $S = \{x \in r : x.a = y\}$ from a relation r. Assume we have an index of r, with respect to an attribute a, which gives us a set of small disk blocks where we can find S. In this situation, the access time for retrieving S is dominated by rotational delays and seek times rather than the access or transmission time of the data [15]. Unfortunately, simultaneously minimizing seek times and rotational delays is NP-hard [1]. On the other hand, the explosive growth of disk storage densities in the last years suggests considering the case where the accessed file fits into a narrow zone on the disk. In this case, seek times can be bounded by a constant that is only slightly larger than the overhead for request initiation, data transmission, and settling the head into a stable state.[†] Such constant overheads can be absorbed into the size of the blocks, and we end up with a problem where only block lengths and rotational delays matter. Interestingly, this reasoning leads to a model logically identical to *drums* – rotating storage devices from the 1950s [5] that have been technologically outdated since the 1970s. Drums have a separate read/write head for every data track and hence suffer no seek delays.

To read a block on a drum one just has to wait until its start rotates to the position of the read head, and then read until the end of the block. Suppose we want to read a batch of n blocks efficiently. Each block can be modelled as an arc in an arc allocation problem. Obviously, L is a lower bound for the number of drum rotations needed to retrieve all n blocks. Fuller [8, 18] found optimal and near-optimal drum scheduling algorithms that can retrieve the blocks in at most $L + 1$ drum rotations. One such algorithm, *shortest latency time first* (SLTF), is very simple: when the drum is at point x, pick the unread block i whose starting point a_i is closest, read it, set $x = a_i + \alpha_i \bmod 1$ and iterate until all blocks are read.

[†] A typical current disk has rotational delays up to 8.33 ms, a constant seek overhead of about 1ms, and nonconstant seek overhead for small seeks of a few microseconds per track.

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

The question arises, how good is an optimal schedule? In the worst case, n rotations may be needed. For example, if all blocks have the same starting point. In this case, even a good scheduling algorithm is of little help. Our results provide us with very attractive performance guarantees if the starting points are randomized. We need time $n\alpha + O(\sqrt{n})$ rotations and hence, for large n, almost all of the schedule time is spent productively reading data.

Performance guarantees for random starting points need not merely be predictions for the average case if we randomize the mapping of logical blocks to physical positions. Here is one scheme with particular practical appeal. Start with a straightforward non-randomized mapping of logical blocks to physical positions by filling one track after the other. Now rotate the mapping on each track by a random amount. This way, accesses to consecutive blocks mapped to the same track remain consecutive. A technical problem is that starting points are not completely independent, so our analysis does not strictly apply. However, starting points of two blocks in a batch of blocks to be read are either independent or the two blocks do not overlap (we merge consecutive blocks on the same track). Therefore it seems likely that our upper bounds still apply.

5.2. Parallel disk striping

Assume a file is to be allocated to M parallel disks numbered 0 to $M-1$, so that accesses to any consecutive range of data in the file can be done efficiently. A common allocation approach is *striping*: block i of the data stream is allocated to disk $i \bmod M$. The situation gets more complicated if n files are accessed concurrently. Barve, Grove and Vitter [2] propose *simple randomized striping* (SR), where block i of a file f is mapped to disk $r_f + i \bmod M$, where r_f is a random offset between 0 and $M-1$. The number of I/O steps needed to access N blocks from the n data streams and M disks is the variable $L^{(\mathrm{cb})}$ in the corresponding chains-into-bins problem. Our results improve the performance bounds shown in [2] for the case when $n \ll N$.

5.3. Adaptive mapping of data to heterogeneous Disks

An effective approach to using many disks in storage networks is to randomly map data blocks to disks. Brinkmann, Salzwedel and Scheideler [4] consider the problem of hashing data blocks to disks of different capacity such that all disks get data proportional to their capacity. In addition, when disks are added or removed, the amount of data to be migrated should be small. They propose the *SHARE* strategy, where each disk is assigned to a number of randomly allocated arcs such that the total length of these arcs is proportional to the disk capacity. A block is mapped in two stages. First it is mapped to a random position on the unit circle, and then to one of the arcs overlapping this position. The quality of load balancing in this scheme is governed by the maximum overlap of arcs. Our bounds therefore allow a refined analysis of the SHARE strategy.

5.4. Cache efficient allocation of DRAM memory

Many algorithms that are efficient on memory hierarchies are based on accessing n arrays in such a way that accesses in each array are sequential, but where it is input dependent how accesses to different arrays are interleaved. For example, multiway mergesort merges

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

several sorted arrays into a single sorted array. In [16] it is shown that these algorithms even work well on hardware caches where we have no direct control over the cache content, provided that the starting addresses of the arrays modulo the cache size M are chosen at random. (Essentially, the SR disk striping from [2] is applied to words in the cache.) Now the question arises, how to allocate the arrays. A naive approach allocates one array after the other in arbitrary order, leaving random gaps of size up to M between the arrays. On average this strategy wastes $nM/2$ of storage. A better way is to apply the simple SLTF algorithm we have seen for drum scheduling. This way we need space ML on average, wasting only $M\Delta$ of memory.

This bound for the offline algorithm also translates into a performance guarantee for an online allocation algorithm where requests for memory segments arrive one at a time. The following greedy algorithm can be easily proved to be 2-competitive. Keep a list of free memory intervals sorted by starting address. Allocate a new segment in the first free interval that has room at the right address offsets. Hence, space $2M(L + \Delta)$ memory suffices to fulfil all requests. In the final version of [13], citing our result replaces a lengthy derivation that shows a bound of $4eML$ for the online algorithm. No result on the offline case was available there.

Acknowledgements

We would like to thank Ashish Gupta and Ashish Rastogi for fruitful cooperation on disk scheduling algorithms.

References

[1] Andrews, M., Bender, M. A. and Zhang, L. (1996) New algorithms for the disk scheduling problem. In *37th Annual Symposium on Foundations of Computer Science*, IEEE Computer Society Press, pp. 550–559.
[2] Barve, R. D., Grove, E. F. and Vitter, J. S. (1997) Simple randomized mergesort on parallel disks. *Parallel Computing* **23** 601–631.
[3] Bhattacharya, B., Hell, P. and Huang, J. (1996) A linear algorithm for maximum weight cliques in proper circular arc graphs. *SIAM J. Discrete Math.* **9** 274–289.
[4] Brinkmann, A., Salzwedel, K. and Scheideler, C. (2002) Compact, adaptive placement schemes for non-uniform requirements. In *14th ACM Symposium on Parallel Algorithms and Architectures*, pp. 53–62.
[5] Cohen, A. A. (1950) Technical developments: Magnetic drum storage for digital information processing systems (in Automatic Computing Machinery). *Mathematical Tables and Other Aids to Computation* **4** 31–39.
[6] Corless, R. M., Gonnet, G. H., Hare, D. E. G., Jeffrey, D. J. and Knuth, D. E. (1996) On the Lambert W function. *Adv. Comput. Math.* **5** 329–359.
[7] Feller, W. (1968) *An Introduction to Probability Theory and its Applications*, Wiley, 3rd edn.
[8] Fuller, S. H. (1972) An optimal drum scheduling algorithm. *IEEE Trans. Comput.* **21** 1153–1165.
[9] Garey, M. R., Johnson, D. S., Miller, G. L. and Papadimitriou, C. H. (1980) The complexity of coloring circular arcs and chords. *SIAM J. Alg. Discrete Methods* **1** 216–227.
[10] Hagerup, T. and Rüb, C. (1990) A guided tour of Chernoff bounds. *Inform. Process. Lett.* **33** 305–308.
[11] Klee, V. (1969) What are the intersections graphs of arcs in a circle? *Amer. Math. Monthly* **76** 810–813.

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press

[12] Kolchin, V. F., Sevatyanov, B. A. and Chistiakov, V. P. (1978) *Random Allocations*, V. H. Winston.
[13] Mehlhorn, K. and Sanders, P. (2002) Scanning multiple sequences via cache memory. *Algorithmica* **35** 75–93.
[14] Raab, M. and Steger, A. (1998) 'Balls into bins': A simple and tight analysis. In *RANDOM: International Workshop on Randomization and Approximation Techniques in Computer Science*, Vol. 1518 of *Lecture Notes in Computer Science*, Springer, pp. 159–170.
[15] Ruemmler, C. and Wilkes, J. (1994) An introduction to disk drive modeling. *IEEE Computer* **27** 17–28.
[16] Sanders, P. (1999) Accessing multiple sequences through set associative caches. In *ICALP*, Vol. 1644 of *Lecture Notes in Computer Science*, Springer, pp. 655–664.
[17] Siegel, A. F. and Holst, L. (1982) Covering the circle with random arcs of random sizes. *J. Appl. Probab.* **19** 373–381.
[18] Stone, H. S. and Fuller, S. F. (1973) On the near-optimality of the shortest-access-time-first drum scheduling discipline. *Comm. Assoc. Comput. Mach.* **16** 352–353. Also published as: Technical Note No.12, DSL.
[19] Vitter, J. S. and Flajolet, P. (1990) Average case analysis of algorithms and data structures. In *Handbook of Theoretical Computer Science*, Vol. A: *Algorithms and Complexity*, Chapter 9, Elsevier, pp. 431–524.

https://doi.org/10.1017/S0963548303005601 Published online by Cambridge University Press