

Tail bounds and Expectations for Random Arc Allocation and Applications

Peter Sanders and Berthold Vöcking*

February 26, 2002

Abstract

The paper considers a generalization of the well known random placement of balls into bins. Given n circular arcs of lengths α_i , $0 \leq i < n$ we study the maximum number of overlapping arcs on a circle if the starting points of the arcs are chosen randomly. We give almost exact tail bounds on the maximum overlap of the arcs. These tail bounds yield a complete characterization of the expected maximum overlap that is tight up to constant factors in the lower order terms.

We illustrate the strength of our results by presenting new performance guarantees for several application: Minimizing rotational delays of disks, scheduling accesses to parallel disks and allocating memory to limit cache interference misses.

*Max Planck Institut für Informatik Saarbrücken, Germany, [voecking, sanders]@mpi-sb.mpg.de, partially supported by the Future and Emerging Technologies programme of the EU under contract number IST-1999-14186 (ALCOM-FT).

1 Introduction

Randomly assigning tasks or data to computational resources has proved an important load balancing technique for many algorithmic problems. The model we study here was motivated by three different applications one of the authors became involved with where the lack of an accurate analysis of such a simple became an obstacle. Section 4 gives more details on these applications: minimizing rotational delays in disk scheduling, scheduling parallel disks, and memory allocation to limit cache conflicts. The common theme there is that jobs (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). In all three applications randomization is used to avoid make worst case situations unlikely.




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

The following model describes all three applications. It is so simple that we expect further applications. An *arc allocation* describes the arrangement of n circular arcs that are pieces of a *unit circle*, i.e., a circle with circumference 1. We represent points on this circle by numbers from the half open interval $[0, 1)$. Let $0 \leq a_i < 1$ denote the starting point of *left endpoint* of arc i and α_i the *arc length* of arc i . Arc i spans the half open interval $[a_i, a_i + \alpha_i \bmod 1)$ where $x \bmod 1$ denotes the fractional part of x and where $[a, b)$ for $a > b$ denotes the set $[a, 1) \cup [0, b)$ in this paper. Figure 1 gives an example. Let $\alpha = \sum_i \alpha_i / n$ denote the *average arc length*. If all arc length are identical, $\alpha_i = \alpha$ for all $0 \leq i < n$, we have a *uniform arc allocation*. 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 and derive tail bounds for L .

Let us go back to the interpretation of arcs as jobs to be executed/allocated. The maximum overlap L is important because all jobs can be executed using between L and $L + 1$ trips around the circle if all jobs are known in advance (see Section 4.1). Furthermore, there is a natural online allocation algorithm that needs at most $2L$ trips around the circle (see Section 4.3).

1.1 New Results for Random Arc Allocation

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 are able to describe

$\mathbf{E}[\Delta]$ almost exactly in terms of Lambert's W function [7]. This function is discussed in more detail below. We obtain the following estimates for $\mathbf{E}[\Delta]$.

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

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

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

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

Our estimates on $\mathbf{E}[\Delta]$ in all four cases are essentially tight in the sense that they describe the case of uniform arc length exactly up to constant factors. Observe that the cases D) and C) are symmetric to the cases A) and B) in α and $1 - \alpha$, resp. In fact, we will derive these bounds by a simple symmetry arguments treating holes (i.e., the uncovered pieces of the circle) like arcs. In case A) it holds $\mathbf{E}[L] = \mathbf{E}[\Delta]$. We will see that exact estimates for this case can be derived relatively easily using Chernoff bounds. More interesting is case D. To obtain tight estimates in this case, we need to combine Chernoff bounds with a random walk analysis.

Now let us come to a discussion of Lambert's W function. This function is defined to be the unique positive solution to the equation $W(x) \cdot \exp(W(x)) = x$ for $x > 0$. Of particular interest for us is the function $\exp(W(x))$. Asymptotically, this function can be estimated by $\lim_{x \rightarrow \infty} \exp(W(x)) = x / \ln(x)$. For example, consider the estimate of $\mathbf{E}[\Delta]$ for the subcase that arcs are very short, say $\alpha = O(1/n)$. In this case, the characterization above gives $\mathbf{E}[L] = \mathbf{E}[\Delta] = O((\ln n) / \ln \ln n)$. Furthermore, if $\alpha = \Theta(\frac{\log n}{n})$ then we obtain $\mathbf{E}[L] = \mathbf{E}[\Delta] = O((\ln n) / \ln \ln n)$.

Finally, in some applications there might be arcs wrapping around the circle several times, i.e., $\alpha_i > 1$. Clearly, in this case our bounds for $\mathbf{E}[\Delta]$ transfer immediately when using $\alpha' = \frac{1}{n} \sum_i \alpha_i \bmod 1$ instead of α . In Section 2 we prove these bounds for the uniform case. The generalization to variable arc lengths is discussed in Section 3.

1.2 Results for Chains-into-bins

Many applications in computer science also require a discrete variant of arc allocation where the circle is subdivided into M equal bins and where arc end points are multiples of a bin size. We note that our proof techniques and hence our upper bounds 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 that are arranged in a circle. A chain is allocated by throwing its first ball into a bin chosen independently, uniformly at random and

by putting its remaining balls into adjacent bins in a round robin fashion. In this notation, the bounds in A) and B) become

$$\mathbf{E}[L^{(\text{cb})}] = \Theta\left(\frac{N}{M} W^*\left(\frac{\ln\left(\frac{Mn}{N}\right)}{N/M}\right)\right) \text{ if } N \geq M, \quad (1)$$

$$\mathbf{E}[\Delta^{(\text{cb})}] = \Theta\left(\sqrt{\frac{N}{M} \ln\left(\frac{Mn}{N}\right)}\right) \text{ if } N = \Omega(M \ln(Mn/N)) \text{ and } \alpha \leq M/2. \quad (2)$$

where $L^{(\text{cb})}$ is the number of balls in the fullest bin and $\Delta^{(\text{cb})} = L^{(\text{cb})} - N/M$. In the way one can translate the results in the cases C) and D).

Let us compare our results for chains-into-bins to the well known results for balls-into-bins processes. These processes are among of the most intensively studied stochastic processes in the context of algorithm analysis (e.g., [14, 23, 17, 12, 2, 24, 16, 20, 4]). The simplest balls-into-bins process assumes that N balls are placed at random into M bins [14, 23, 17]. Balls-into-bins are the special case of chains-into-bins where all chains consist of a single ball, i.e., $n = N$. We get

$$\mathbf{E}[L^{(\text{bb})}] = \Theta\left(\frac{N}{M} W^*\left(\frac{\ln M}{N/M}\right)\right) \text{ if } N \geq M, \quad (3)$$

$$\mathbf{E}[\Delta^{(\text{bb})}] = \Theta\left(\sqrt{\frac{N}{M} \ln M}\right) \text{ if } N = \Omega(M \ln M). \quad (4)$$

The Bounds (3) and (4) are well known although other papers [14, 23, 17] use a different, slightly more complicated notation that yields more information about constant factors.

Another instructive perspective is that arc allocations are related to balls-into-bins systems with $1/\alpha$ bins. Our analyses for L and Δ will give further insights into the relationship between the two different random processes.

1.3 Previous Results

Barve et al. [3] introduce the chains-into-bins problem and show why several tail bounds for the case $N = n$ also apply to the general case. Apparently, $\mathbf{E}[L^{(\text{cb})}]$ can only grow if chains are atomized into individual balls (although this is not proven yet). Our bounds improve these results by showing that $\Delta^{(\text{cb})}$ can be much smaller if $n \ll N$, i.e., if chains are long.

Chains-into-bins have been analyzed asking what is the expected number of bins with at least a balls [19]. This measure was needed to estimate the number of cache misses in executing a class of cache-efficient algorithms. Refer to Section 4.3 for more details.

Arc allocations have been studied in mathematics under the aspect of when the arcs cover the circle (e.g., [21]). This is related to the minimum overlap which seems to be more important for most computer science applications. We have adopted the convention from these papers to measure arc lengths between 0 and 1 rather than 0 and 2π in order to avoid notational overhead.

An arc allocation defines a *circular arc graph* [13, 10] with n nodes where there is an edge between nodes i and j 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. But note that the maximum overlap clique is not necessarily maximum clique of a circular arc graph [5].

Additional related work for the applications is described in Section 4 along with the applications themselves.

2 Uniform Arcs

In this section we assume that all arcs have the same length. At first, we present a simple tail bound that is best suited in the case of short arcs. This tail bound gives the bound stated in case A) of our characterization of $\mathbf{E}[\Delta]$. At second, we present a tail bound best suited for relatively long arcs This tail bound implies the estimates given in case B) of our characterization. At third, we show that the estimates given for the cases C) and D) follow from a simple symmetry argument. Finally, we present a majorization argument as well as some simulations showing that our tail bounds are essentially tight.

2.1 Tail bound for 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. In particular, if the average arc length is $\alpha = O(\frac{\log n}{n})$ then this tail bound can be used to obtain the tight estimates that we presented in case A) of our characterization for $\mathbf{E}[\Delta]$.

In fact, the proof for this tail bound only needs one simple idea. The idea is that one only need to estimate the maximal overlap at exactly n positions on the circle, namely at the left endpoints of the arcs. The maximum overlap at these positions is estimated by using Chernoff bounds.

Theorem 1. *Suppose n arcs of length α are placed at random onto the unit circle Let $\mu \geq \alpha n$ denote an upper bound on the average overlap. Then, for every $\varepsilon > 0$,*

$$\Pr[\Delta \geq \varepsilon\mu + 1] \leq n \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^\mu .$$

Proof. 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

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

Now as all other arcs are placed independently from arc i , we can apply a Chernoff bound, which gives

$$\Pr[\Delta(a_i) \geq \varepsilon\mu + 1] \leq \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^\mu .$$

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

$$\Pr[\Delta \geq \varepsilon\mu + 1] \leq \sum_{i=1}^n \Pr[\Delta(a_i) \geq \varepsilon\mu + 1] \leq n \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^\mu ,$$

which proves the theorem. \square

2.2 Tail bound for 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 approach loses a factor n in comparison to 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 losing only a factor $\frac{1}{\alpha}$ in comparison to the Chernoff bound. This bound is stronger than the bound given in Theorem 1 if arcs are relatively long, in particular in the case $\alpha = \Omega\left(\frac{\log n}{n}\right)$. We will prove in Section 2.4 that this tail bound is tight in the sense that the factor $\frac{1}{\alpha}$ in front of the typical Chernoff bound term is unavoidable.

Theorem 2. *Suppose n arcs of length α are placed at random onto the unit circle. Assume $\mu \geq \alpha n$. Then, for every $\varepsilon > 0$,*

$$\Pr[\Delta \geq 5\varepsilon\mu] \leq \frac{6}{\alpha} \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^\mu.$$

Proof. For simplicity let us assume that $\alpha \leq \frac{1}{2}$. In Section 2.3 we will show that the proof can be generalized to the case $\alpha > \frac{1}{2}$ in such a way that we even obtain a much stronger tail bound for arc length close to 1.

Define $\kappa = \lceil 1/\alpha \rceil \geq 2$. Let $x_0, \dots, x_{\kappa-1}$ denote κ points on the circle that decompose the circle into κ intervals $X_i = [x_i, x_{i+1})$ of identical length $\lceil \frac{1}{\kappa} \rceil \approx \alpha$. (Here and in the following $i+1$ abbreviates $(i+1) \bmod \kappa$.) Observe that every arc has at most one endpoint in each interval. 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. *For every $\varepsilon > 0$, $\Pr[\Delta_i \geq \varepsilon\mu] \leq \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^\mu$.*

Claim 4. *For every $\varepsilon > 0$, $\Pr[\Delta'_i \geq \varepsilon(2+2\varepsilon)\mu \mid \Delta_i + \Delta_{i+1} \leq 2\varepsilon\mu] \leq 2 \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^{(1+\varepsilon)\mu}$.*

Before we turn to proving these claims let us show that, in fact, these claims imply the theorem. First suppose $\varepsilon \geq 1$. 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 a1: $\Delta \leq \max_i \{\alpha n + \Delta_i + \Delta_{i+1}\}$ so that we obtain

$$\begin{aligned} \Pr[\Delta \geq 3\varepsilon\mu] &\stackrel{(a1)}{\leq} \Pr[\exists_{i=0}^{\kappa-1} : \alpha n + \Delta_i + \Delta_{i+1} \geq 3\varepsilon\mu] \\ &\stackrel{(a2)}{\leq} \Pr\left[\exists_{i=0}^{\kappa-1} : \Delta_i \geq \varepsilon\mu\right] \\ &\stackrel{(a3)}{\leq} \frac{1}{\kappa} \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^\mu \\ &\stackrel{(a4)}{\leq} \frac{2}{\alpha} \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^\mu, \end{aligned}$$

where inequality (a2) follows from $\alpha n \leq \varepsilon \mu$, (a3) follows from Claim 3, and (a4) follows from $\kappa = \lceil \frac{1}{\alpha} \rceil$.

Now assume $\varepsilon < 1$. Observe that this implies b1: $\varepsilon \geq \varepsilon(3 + 2\varepsilon)/5$. Furthermore, we apply b2: $\Delta = \max_i \{\Delta_i + \Delta'_i\}$ and obtain

$$\begin{aligned}
\Pr[\Delta \geq 5\varepsilon\mu] &\stackrel{(b1)}{\leq} \Pr[\Delta \geq \varepsilon(3 + 2\varepsilon)\mu] \\
&\stackrel{(b2)}{\leq} \Pr[\exists_{i=0}^{\kappa-1} : \Delta_i \geq \varepsilon\mu \vee \Delta'_i \geq \varepsilon(2 + 2\varepsilon)\mu] \\
&\stackrel{(b3)}{\leq} \sum_{i=0}^{\kappa-1} \Pr[\Delta_i \geq \varepsilon\mu] + \Pr[\Delta'_i \geq \varepsilon(2 + 2\varepsilon)\mu \mid \Delta_i + \Delta_{i+1} \leq 2\varepsilon\mu] \\
&\stackrel{(b4)}{\leq} \kappa \left(\frac{e^\varepsilon}{(1 + \varepsilon)^{1+\varepsilon}} \right)^\mu + 2\kappa \left(\frac{e^\varepsilon}{(1 + \varepsilon)^{1+\varepsilon}} \right)^{(1+\varepsilon)\mu} \\
&\leq \frac{6}{\alpha} \left(\frac{e^\varepsilon}{(1 + \varepsilon)^{1+\varepsilon}} \right)^\mu .
\end{aligned}$$

The following basic fact from probability theory implies inequality (b3). For events X, X' , and Y with $X' \subseteq X$ it holds $\Pr[X \vee Y] \leq \Pr[X] + \Pr[Y \setminus X'] \leq \Pr[X] + \Pr[Y|X']$. Furthermore, inequality (b4) follows from the Claims 3 and 4.

It remains to prove the two claims. Observe that $\mathbf{E}[L(x_i)] \leq \mu$ so that the bound in Claim 3 follows directly from a Chernoff bound. Hence it only remains to show the Claim 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 $\Delta_i + \Delta_{i+1} \leq 2\varepsilon\mu$, we can assume

$$m \leq L(x_i) + L(x_{i+1}) = 2\mu + \Delta_i + \Delta_{i+1} \leq (2 + 2\varepsilon)\mu .$$

For the time being, assume that m is fixed. Let y_1, \dots, 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, \dots, 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 $\Delta_i + \Delta_{i+1} \leq 2\varepsilon\mu$. As this assumption does not affect the arc's orientation, the variables s_1, \dots, s_m are independent and $\Pr[s_j = 1] = \Pr[s_j = -1] = \frac{1}{2}$, for $1 \leq j \leq m$.

Now let us define $S_j = \sum_{k=1}^j s_k$, for $0 \leq j \leq m$. Notice that the sequence S_0, S_1, \dots, 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, $\Delta'_i = \max_{0 \leq i \leq m} (S_i)$. Hence, we can estimate Δ'_i by analyzing this random walk. Applying Theorem III.7.1 from [8] we can derive the following probability bound. For every $r \geq 0$,

$$\Pr[\exists j \{1, \dots, m\} : S_j \geq r] = \sum_{k=r+1}^{\infty} \Pr[S_m = k] + \Pr[S_m = k+1] \leq 2\Pr[S_m > r] .$$

Next we observe that the random variable $X_m = S_m/2 + m/2$ follows the binomial distribution and, hence, can be estimated using a Chernoff bound. In this way, we obtain

$$\Pr[S_m > \varepsilon m] = \Pr[X_m > (1 + \varepsilon)m/2] \leq \left(\frac{e^\varepsilon}{(1 + \varepsilon)^{1+\varepsilon}} \right)^{m/2},$$

for every $\varepsilon > 0$. As a consequence,

$$\Pr[\Delta'_i > \varepsilon m] \leq 2\Pr[S_m > \varepsilon m] \leq 2 \left(\frac{e^\varepsilon}{(1 + \varepsilon)^{1+\varepsilon}} \right)^{m/2}.$$

Clearly, Δ'_i is monotonically increasing in m . Therefore, the worst-case choice for m is $m = (2 + 2\varepsilon)\mu$, and we obtain,

$$\Pr[\Delta'_i > \varepsilon(2 + 2\varepsilon)\mu] \leq 2 \left(\frac{e^\varepsilon}{(1 + \varepsilon)^{1+\varepsilon}} \right)^{(1+\varepsilon)\mu},$$

which proves Claim 4. Thus Theorem 2 is shown. \square

2.3 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 , $L(x) - \alpha n = \bar{L}(x) - \bar{\alpha}n$. 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 Chernoff bounds and random walks analogous to the analyses we presented in the previous sections, apart from the fact that we are now bounding the deviation below instead of above the mean. In this way we can derive the following tail bounds.

Theorem 5. *Suppose n arcs of length α are placed at random onto the unit circle. Assume $\bar{\mu} \geq (1 - \alpha)n$. Then, for every $\varepsilon > 0$,*

$$\Pr[\Delta \geq \varepsilon\bar{\mu} + 1] \leq \frac{1}{n} \left(\frac{e^\varepsilon}{(1 + \varepsilon)^{1+\varepsilon}} \right)^{\bar{\mu}},$$

and $\Pr[\Delta \geq 6\varepsilon\bar{\mu}] \leq \frac{6}{1 - \alpha} \left(\frac{e^\varepsilon}{(1 + \varepsilon)^{1+\varepsilon}} \right)^{\bar{\mu}}.$

Observe that these tail bounds are completely symmetrical to the tail bounds given in Theorem 1 and 2. If $\alpha > \frac{1}{2}$ then these bounds are stronger, if $\alpha \leq \frac{1}{2}$ then the other bounds are stronger.

2.4 A lower bound in terms of balls-into-bins

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 with $N = n$ balls and $M = 1/\alpha$ 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 can not exceed the maximum overlap of the random arc allocation. Let us summarize this simple observation formally.

Observation 6. *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 bounds given in the characterization for $\mathbf{E}[\Delta]$ match the bounds for balls-into-bins processes with $N = n$ balls and $M = 1/\alpha$ bins. As a consequence, all these bounds are essentially tight in the sense that they describe the case of uniform arc length exactly up to constant factors.

2.5 Simulations

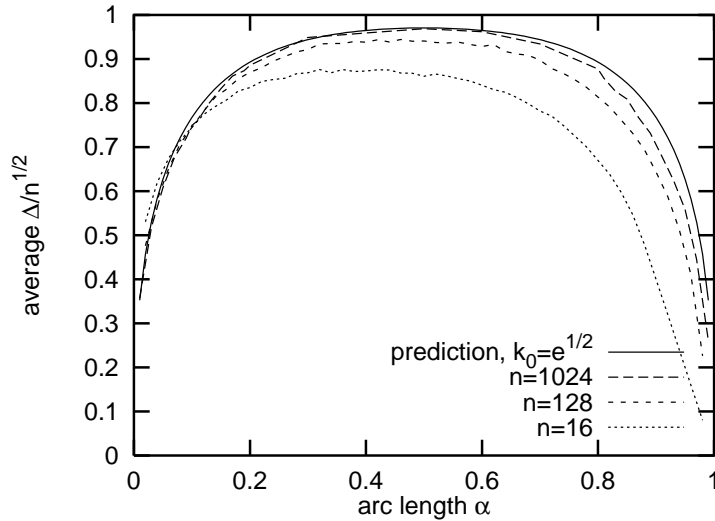


Figure 2: Comparison of the theoretical prediction $\Delta = \sqrt{en \cdot \alpha(1 - \alpha) \ln(1/(\alpha(1 - \alpha)))}$ with simulations for different n . The measurements are averages of 10000 repetitions for $n \leq 128$ and 1000 repetitions for $n = 1024$.

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 bounds in

these two cases we obtain the following elegant estimate

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

In fact, the measured and predicted curves are quite close together 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 since we use the same Chernoff bounds to estimate the overlap of arcs in case of $\alpha \leq \frac{1}{2}$ and the overlap of holes in case of $\alpha > \frac{1}{2}$. In fact, however, 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.

3 Non-Uniform Arc Lengths

In case of non-uniform arcs, we can obtain almost the same tail bounds as in the uniform case. In fact, the tail bound in Theorem 1 and its symmetric counterpart given in Theorem 5 hold with out any change if we assume that α denotes the average arc length. However, in order to prove results as in Theorem 2 and its symmetric counterpart we need some further, more complicated arguments. Observe that the proof of this theorem fails in case of non-uniform arcs as it cannot longer be guaranteed that arcs have at most one endpoint in each interval. We obtain the following tail bound that is only slightly weaker than the bound in Theorem 2.

Theorem 7. *Suppose n arcs of average length α are placed at random onto the unit circle. Assume $\mu \geq \alpha n$. Then, for every $\varepsilon > 0$,*

$$\Pr[\Delta \geq \varepsilon\mu + 2] \leq \frac{k_1}{\alpha} \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^{k_2\mu},$$

for suitable constants $k_1, k_2 > 0$.

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

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

$$\Pr[\Delta \geq \varepsilon\mu + 1] \leq n \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^\mu.$$

Essentially, we have to replace the factor n by $\frac{1}{\alpha}$. Let us define $f(\varepsilon, \mu) = \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^{\mu/3}$. Clearly, if $n \leq f(\varepsilon, \mu)^{-2}$ then the bound above implies $\Pr[\Delta \geq \varepsilon\mu + 1] \leq f(\varepsilon, \mu)^{-1} \leq \frac{1}{\alpha} f(\varepsilon, \mu)^{-1}$, which corresponds to the tail bound that we have to prove. Otherwise, $n > f(\varepsilon, \mu)^{-2}$. Solving this

inequality for μ gives $\mu < c \ln n$, for some positive constant c . Now applying $\mu \geq \alpha n$ gives $\frac{c}{\alpha} > \frac{n}{\ln n}$. Consequently,

$$\frac{c}{\alpha} f(\varepsilon, \mu) > \frac{n}{\ln n} f(\varepsilon, \mu) \geq \sqrt{n} f(\varepsilon, \mu) \geq 1 ,$$

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

Case 2: Now suppose $\varepsilon < \frac{1}{8}$. In this case, we partition the set of arcs into subsets according to their length. For a subset 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 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, $\mu \geq \alpha_T |A|$. Hence, applying the bound stated in Theorem 2 gives

$$\Pr[\Delta_T \geq 5\varepsilon\mu] \leq \frac{6}{\alpha_T} \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^\mu \leq \frac{6}{\alpha} \left(\frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \right)^\mu$$

because $\alpha_T \geq \alpha$ 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 2.3.

Claim 8. *For a set S with at most n arcs of length at most α ,*

$$\Pr[\Delta_S \geq \varepsilon\alpha n + 1] \leq \frac{k_1}{\alpha} \exp(-k_2 \varepsilon^2 \alpha n) ,$$

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

The following approximation is typically used in the context of Chernoff bounds (cf., e.g., [11]). For every $\varepsilon \leq 1$,

$$\exp(-\varepsilon^2) \leq \frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}} \leq \exp(-\varepsilon^2/2) .$$

From this approximation we can conclude that Claim 8 implies Theorem 7 so that it remains only to prove this claim.

We prove Claim 8 using a recursive approach. In each step of the recursion, the value of ε is increased by a factor of 4 so that the recursion is guaranteed to terminate in case 1. Let R denote the subset of arcs of length at most $\beta = \alpha/8$. Furthermore, let $\bar{R} = S \setminus R$, that is, \bar{R} contains the arcs with length from the interval $(\alpha/8, \alpha]$. By recursion, we can assume

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

for every $\gamma > 0$. Hence,

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

The last inequality follows from $\exp(-2x) \leq \frac{1}{16} \exp(-x)$, for every $x \geq 3$. Here we additionally assume $k_1 \geq e^3$ and $k_2 \varepsilon^2 \alpha n \geq \ln(k_1) \geq 3$. Observe that we can make the last assumption w.l.o.g. as $k_2 \varepsilon^2 \alpha n < \ln(k_1)$ would imply that the right hand side of the bound in Claim 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 2 in order to estimate $\Delta_{\bar{R}}$. In this way, we can show

$$\Pr \left[\Delta_{\bar{R}} \geq \frac{\varepsilon \alpha n}{2} \right] \leq \frac{k_3}{\alpha} \exp \left(-k_4 \left(\frac{\varepsilon}{2} \right)^2 \alpha n \right) ,$$

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

$$\begin{aligned}
\Pr[\Delta \geq \varepsilon \alpha n + 1] &\leq \Pr[\Delta_R + \Delta_{\bar{R}} \geq \varepsilon \alpha n + 1] \\
&\leq \frac{k_1}{2\alpha} \exp(-k_2 \varepsilon^2 \alpha n) + \frac{k_3}{\alpha} \exp \left(-k_4 \left(\frac{\varepsilon}{2} \right)^2 \alpha n \right) \\
&\leq \frac{k_1}{\alpha} \exp(-k_2 \varepsilon^2 \alpha n) ,
\end{aligned}$$

which completes the proof of Theorem 7. □

4 Applications

We now present three examples, where our bounds on arc allocation yield performance guarantees. In Section 4.1 Bound (2) guarantees that rotational delays in accessing a single disks eventually become small compared to data access times. Section 4.2 gives a different disk scheduling application where load balancing data accesses over several disks is the objective. Whereas the first two example concern execution time, the last example in Section 4.3 bounds the memory consumption of a technique for reducing cache interference misses.

In all three applications very bad worst case behavior is avoided using randomization. The price paid is that the best case behavior gets worse. Since best case behavior may sometimes not be far from real inputs, it is crucial for our performance bounds to demonstrate that this possible penalty is very small. This is a quite strong incentive to study lower order terms in the performance bounds rather bounds that leave the constant factors open.

4.1 Disks and drums

One of our main motivations for studying arc allocations was the desire to find disk scheduling algorithms that take rotational delays into account. For example, consider a data base 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 attribute a that tells 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 [18]. 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 to consider 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 50s [6] that are technologically outdated since the 70s. 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 modeled as an arc in an arc allocation problem in the obvious way. Obviously, L is a lower bound for the number of drum rotations needed to retrieve all n blocks. Fuller [9, 22] 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.

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 that 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.

4.2 Parallel disk striping

Assume a file is to be allocated to M parallel disks numbered 0 through $M - 1$ 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 et al. [3] 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 $L^{(\text{cb})}$ in the corresponding chains-into-bins problem. Our results improve the performance bounds shown in [3] for the case that $n \ll N$.

4.3 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 unpredictable how accesses to different arrays are interleaved. In [19] 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 points of the arrays modulo the cache size M are chosen at random. (Essentially the SR disk striping from [3] is applied to words in the cache.) Now the question arises how to allocate the arrays. Assume we have a large contiguous amount of memory available starting at address $x \bmod M$. A naive approach allocates one array after the other in arbitrary order by picking a random location $y \bmod M$ and allocating the array so that it starts at point $x + (y - x \bmod M)$. On the average this strategy wastes $nM/2$ of storage. A better way is by applying the simple SLTF algorithm we have seen for drum scheduling. This way we need space ML on the 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 proven 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 fulfill all requests. In the final version of [15] 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] M. Andrews, M. A. Bender, and L. Zhang. New algorithms for the disk scheduling problem. In IEEE, editor, *37th Annual Symposium on Foundations of Computer Science*, pages 550–559. IEEE Computer Society Press, 1996.
- [2] Y. Azar, A. Z. Broder, A. R. Karlin, and Eli Upfal. Balanced allocations. *SIAM Journal on Computing*, 29(1):180–200, February 2000.

- [3] R. D. Barve, E. F. Grove, and J. S. Vitter. Simple randomized mergesort on parallel disks. *Parallel Computing*, 23(4):601–631, 1997.
- [4] P. Berenbrink, A. Czumaj, A. Steger, and B. Vöcking. Balanced allocations: The heavily loaded case. In *32th Annual ACM Symposium on Theory of Computing*, 2000.
- [5] Binay Bhattacharya, Pavol Hell, and Jing Huang. A linear algorithm for maximum weight cliques in proper circular arc graphs. *SIAM Journal on Discrete Mathematics*, 9(2):274–289, May 1996.
- [6] Arnold A. Cohen. Technical developments: Magnetic drum storage for digital information processing systems (in Automatic Computing Machinery). *Mathematical Tables and Other Aids to Computation*, 4(29):31–39, January 1950.
- [7] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the lambert w function. *Advances in Computational Mathematics*, 5:329–359, 1996.
- [8] W. Feller. *An Introduction to Probability Theory and its Applications*. Wiley, 3rd edition, 1968.
- [9] S. H. Fuller. An optimal drum scheduling algorithm. *IEEE Trans. on Computers*, 21(11):1153, November 1972.
- [10] M. R. Garey, D. S. Johnson, G. L. Miller, and C. H. Papadimitriou. The complexity of coloring circular arcs and chords. *SIAM Journal on Algebraic and Discrete Methods*, 1(2):216–227, June 1980.
- [11] T. Hagerup and C. Rüb. A guided tour of chernoff bounds. *Information Processing Letters*, 33:305–308, 1990.
- [12] R. M. Karp, M. Luby, and F. Meyer auf der Heide. Efficient PRAM simulation on a distributed memory machine. In *24th ACM Symp. on Theory of Computing*, pages 318–326, May 1992.
- [13] V. Klee. What are the intersections graphs of arcs in a circle? *Amer. Math. Monthly*, 76:810–813, 1969.
- [14] V. F. Kolchin, B. A. Sevastyanov, and V. P. Chistiakov. *Random Allocations*. V. H. Winston, 1978.
- [15] K. Mehlhorn and P. Sanders. Scanning multiple sequences via cache memory. *Algorithmica*, 2002. accepted subject to minor revisions, old version under <http://www.mpi-sb.mpg.de/~sanders/papers/icalpfull.ps.gz>.
- [16] M. Mitzenmacher, A. Richa, and R. Sitaraman. The power of two random choices: A survey of the techniques and results. In P. Pardalos, S. Rajasekaran, and J. Rolim, editors, *Handbook of Randomized Computing*. Kluwer, 2000.

- [17] M. Raab and A. Steger. “balls into bins” – A simple and tight analysis. In *RANDOM: International Workshop on Randomization and Approximation Techniques in Computer Science*. LNCS, 1998.
- [18] C. Ruemmler and J. Wilkes. An introduction to disk drive modeling. *Computer*, 27(3):17–28, March 1994.
- [19] P. Sanders. Accessing multiple sequences through set associative caches. In *ICALP*, number 1644 in LNCS, pages 655–664, 1999.
- [20] P. Sanders, S. Egner, and J. Korst. Fast concurrent access to parallel disks. In *11th ACM-SIAM Symposium on Discrete Algorithms*, pages 849–858, 2000.
- [21] A. F. Siegel and L. Holst. Covering the circle with random arcs of random sizes. *J. Appl. Probab.*, 19:373–381, 1982.
- [22] H. S. Stone and S. F. Fuller. On the near-optimality of the shortest-access-time-first drum scheduling discipline. *Communications of the ACM*, 16(6), June 1973. Also published in/as: Technical Note No.12, DSL.
- [23] J. S. Vitter and P. Flajolet. Average case analysis of algorithms and data structures. In *Handbook of Theoretical Computer Science*, volume A: Algorithms and Complexity, chapter 9, pages 431–524. Elsevier, 1990.
- [24] B. Vöcking. How asymmetry helps load balancing. In *40th Symposium on Foundations of Computer Science*, pages 131–140. IEEE, 1999.