

HDD average seek time

2021-05-25

The average seek distance of a random hard drive access is 1/3rd of the maximum seek distance, or is it? We re-examine this rule of thumb and find that it is not exactly true. We also provide the full statistics of the seek distance and discuss how to compute the seek time.

blog: https://tetraquark.vercel.app/posts/seek_time/?src=pdf

email: quarktetra@gmail.com

A Hard Disk Drive (HDD) has a spinning disk and a rotating arm to position the read-write head on target tracks, which are concentric circles. It takes a finite amount of time to move the actuator from one point to another, which is referred to as the seek time [1]. Given the task of locating many random sectors on a disk surface, the average time to find the tracks the sectors are located on is the average seek time, and in this post we will compute the statistics of it.

The geometry of the problem is simple: it is just a disk with a hollow center, as shown in Figure 1.




Figure 1: A disk of inner radius r_i and outer radius r_o with two sample tracks located at radii r_1 and r_2 .

Averaging

Consider two tracks at radii r_1 and r_2 where $r_{1,2} \in [r_i, r_o]$. The radial distance between them is $|r_2 - r_1|$. If we average over every possible position of r_1 and r_2 we get the average radial distance. Assuming that each r is equally likely to happen, the density will be $\frac{1}{r_o - r_i}$, and the average radial distance can be calculated as:

$$\begin{aligned}
\overline{\text{Radial Distance}} &= \frac{1}{(r_o - r_i)^2} \int_{r_i}^{r_o} \int_{r_i}^{r_o} dr_2 dr_1 |r_2 - r_1| = \frac{2}{(r_o - r_i)^2} \int_{r_i}^{r_o} dr_2 \int_{r_i}^{r_2} dr_1 (r_2 - r_1) \\
&= \frac{2}{(r_o - r_i)^2} \int_{r_i}^{r_o} dr_2 \left(r_2(r_2 - r_i) - \frac{r_2^2 - r_i^2}{2} \right) = 2 \int_{r_i}^{r_o} dr_2 \left(\frac{r_2^2}{2} - r_2 r_i + \frac{r_i^2}{2} \right) \\
&= \frac{2}{(r_o - r_i)^2} \left(\frac{r_o^3 - r_i^3}{6} - \frac{r_o^2 r_i - r_i^3}{2} + \frac{r_o r_i^2 - r_i^3}{2} \right) = \frac{1}{3(r_o - r_i)^2} (r_o^3 - 3r_o^2 r_i + 3r_o r_i^2 - r_i^3) \\
&= \frac{r_o - r_i}{3},
\end{aligned} \tag{1}$$

where we used the symmetries of the integral to split the range and convert it to a factor of 2.

The integrals were a bit involved since we had nonzero values for both the upper and lower limits. It would be wiser if we define $r_{1,2} = r_i + \tilde{r}_{1,2}$ and $\Delta r = r_o - r_i$. This would give a simpler integral to deal with:

$$\begin{aligned}
\overline{\text{Radial Distance}} &= \frac{1}{(\Delta r)^2} \int_0^{\Delta r} \int_0^{\Delta r} d\tilde{r}_2 d\tilde{r}_1 |\tilde{r}_2 - \tilde{r}_1| = \frac{2}{(\Delta r)^2} \int_0^{\Delta r} d\tilde{r}_2 \int_0^{\tilde{r}_2} d\tilde{r}_1 (\tilde{r}_2 - \tilde{r}_1) \\
&= \frac{2}{(\Delta r)^2} \int_0^{\Delta r} d\tilde{r}_2 \frac{\tilde{r}_2^2}{2} = \frac{\Delta r}{3} = \frac{r_o - r_i}{3}.
\end{aligned} \tag{2}$$

We are treating the radius as a continuous variable although in an HDD tracks are placed in discrete steps. Would it make any difference if we used discrete math? In that case we just need to switch to indices and summations. Let us define $r_1 = r_i + \Delta j$, where Δ is the radial separation of two adjacent tracks, and $j \in [0, N-1]$ is an integer that labels the N tracks on the disk. The probability of landing on a particular track is $1/N$. The averaging becomes:

$$\begin{aligned}
\overline{\text{Radial Distance}} &= \frac{\Delta}{N^2} \sum_{k=0}^{N-1} \sum_{j=0}^{N-1} |k - j| = \frac{2\Delta}{N^2} \sum_{k=0}^{N-1} \sum_{j=0}^{k-1} (k - j) = \frac{2\Delta}{N^2} \sum_{k=0}^{N-1} \left(k^2 - \frac{k(k-1)}{2} \right) = \frac{\Delta}{N^2} \sum_{k=0}^{N-1} (k^2 + k) \\
&= \frac{\Delta}{N^2} \left(\frac{(N-1)N(2N-1)}{6} + \frac{N(N-1)}{2} \right) = \frac{N\Delta}{3} \left(1 + \mathcal{O}\left(\frac{1}{N}\right) \right) = \frac{r_o - r_i}{3},
\end{aligned} \tag{3}$$

where $\mathcal{O}(\frac{1}{N})$ represents the terms of order $1/N$ or smaller. Since the number of tracks on an HDD disk surface is at the order of 10^6 , we can safely drop $1/N$ terms. Furthermore, since Δ is the separation of the tracks, $N\Delta$ was nothing but the radial distance spanned by the tracks, i.e., $r_o - r_i$.

If we look at the ratio of the mean distance to the maximum distance, $r_o - r_i$, we get:

$$\frac{\overline{\text{Radial Distance}}}{\text{Maximum Distance}} = \frac{1}{3}. \quad (4)$$

Now that we have shown in multiple ways that average radial distance is $1/3$ of the maximal distance, we will proceed to show that what we got is not exactly what we were looking for!

Sampling: tracks vs sectors

There was a non-trivial assumption in the previous section: “... *each r is equally likely to happen*”. We baked it into our math when we took the density as $\frac{1}{r_o - r_i}$ in the integrals, and when we took the density as $1/N$ in the summations. However, a closer inspection shows that this assumption is in conflict with uniformly distributed sectors across the surface. We can explicitly show that if we uniformly sample from tracks and uniformly sample from the locations on the tracks, i.e., the angles, the resulting distribution on the surface will **not** be uniform. Let’s do a simulation to confirm. We simply sample radius values in the range $[r_i, r_o]$ from a uniform distribution, and the angle values in the range $[0, 2\pi]$, and convert them into Cartesian coordinates x and y to plot the density of points on the disk.




Figure 2: The probability density of randomly selected points on the disk. The steep gradient shows that uniformly sampled radius and uniformly sampled angle do not yield a uniform distribution on the disk.

As shown in Figure 2, uniformly sampling radius and angle does not yield a uniform distribution on the disk. This intuitively follows from the fact that sampled points are spreading to a wider area at larger radii. In order to restore the uniform distribution on disk surface, we have to abandon the assumption that “... *each r is equally likely to happen*”.

The proper density

If we think a bit harder, we can fix this. Consider a small zone from r to $r + \delta r$, with $\delta r \ll r$. The area of this zone is basically the circumference times the width, that is $2\pi r \delta r$. So the

area is growing $\propto r$. Therefore the probability of landing on this band should grow $\propto r$, i.e., it is not constant. In fact it should be

$$f(r) = \frac{2r}{r_o^2 - r_i^2}, \quad (5)$$

so that it yields 1 when integrated from r_i to r_o . This also makes sense for HDDs. Assuming that the data is uniformly distributed on the disk surface, more data will be on the outer radius of the disk just because there is more area there. Therefore, as the head looks for random sectors, it will find them more on the outer radius, i.e., larger radii tracks will enter into the averaging with larger weights. We have to go back to square one and do the math all over again with the proper probability distribution for tracks. We plug this into the averaging integrals to get:

$$\begin{aligned} \overline{\text{Radial Distance}} &= \int_{r_i}^{r_o} \int_{r_i}^{r_o} dr_2 dr_1 f(r_1) f(r_2) |r_2 - r_1| = \frac{2}{(r_o^2 - r_i^2)^2} \int_{r_i}^{r_o} dr_2 \int_{r_i}^{r_2} dr_1 r_1 r_2 (r_2 - r_1) \\ &= \frac{4}{15} (r_o - r_i) \frac{r_o^2 + 3r_i r_o + r_i^2}{(r_o + r_i)^2} = \frac{4}{15} (r_o - r_i) \left(1 + \frac{r_i r_o}{(r_o + r_i)^2} \right) \\ &= \frac{4}{15} (r_o - r_i) \left(1 + \frac{\kappa}{(1 + \kappa)^2} \right). \end{aligned} \quad (6)$$

The result suggests that the ratio of the mean distance to the maximum distance is

$$\frac{\overline{\text{Radial Distance}}}{\text{Maximum Distance}} = \frac{4}{15} \left(1 + \frac{\kappa}{(1 + \kappa)^2} \right) \text{ where } \kappa \equiv \frac{r_o}{r_i}, \quad (7)$$

which is shown in Figure 3. This is not as elegant as a simple fixed factor of $1/3$, but, nevertheless it is the correct one as we will support further below.




Figure 3: The ratio of the mean distance to the maximum distance vs $\kappa = \frac{r_o}{r_i}$. The marked κ value of 2.43 corresponds to typical disk sizes of outer radius $r_o = 1.7''$ and inner radius $r_i = 0.7''$.

Brute force solution

We can always use Monte Carlo simulations to compute the results. We can select x and y in the range $[-r_o, r_o]$ from uniform distributions, and reject the x, y pairs that fall outside of the disk. That will certainly result in a uniform coverage of the disk as shown in Figure 4.

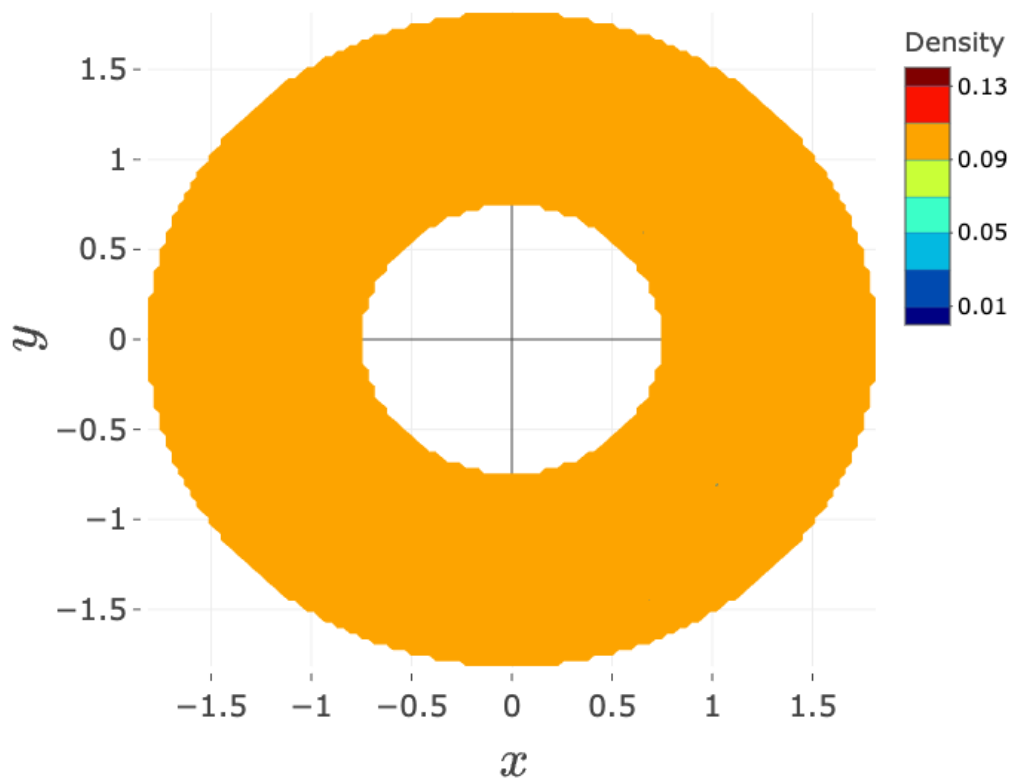


Figure 4: The probability density of points constructed from x and y values selected from uniform distribution. Once the points that fall out of the disk area are rejected, the disk is uniformly covered. We use $r_o = 1.7''$ and inner radius $r_i = 0.7''$ in the simulation.

We can then select a pair of points \vec{r}_1 and \vec{r}_2 , and build the histogram of $|\vec{r}_2 - \vec{r}_1|$, as shown in Figure 5.




Figure 5: The distribution of the difference of two randomly chosen radii r_2 and r_1 . The “model” line is explained in the next section.

Computing the mean value of the distribution and dividing it by the maximum distance, we get approximately 0.27 to be compared to 0.27 as predicted in Eq. 7.

Shut up and calculate!

There is a famous saying among quantum physicists: in most cases quantum behavior becomes so bizarre that there is no intuition that can help you, hence you have to just “shut up and calculate”. I sometimes feel the same way with functions of random variables. In this particular example, we do have an intuitive expectation on the probability distributions and finding out that the distribution we originally used was not uniform on the disk saved the day. However, we would have been in deep trouble if this were a three dimensional problem. It would have been hard to visualize the distributions and/or our intuition might have just failed us. Can’t we just shut up and calculate it formally from the first principles of random variables?

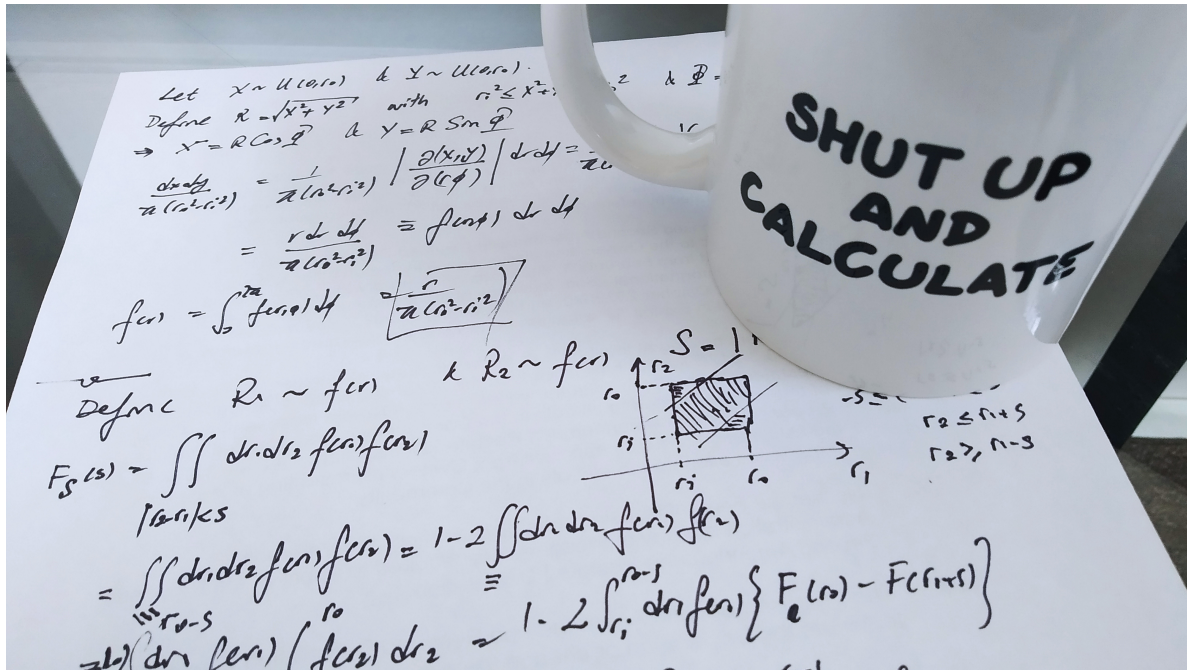


Figure 6: Sometimes it is best to shut up and calculate. Follow the first principles of the theory of random variables and they will lead you to the answer, assuming you don't screw up the math on the way.

Let us first define the cumulative distribution function $F(r, \phi)$. We want it to cover the disk uniformly: its increase, let's call it dF , should be with a uniform rate as we add small portions of area, $dA = r dr d\phi$. That is:

$$\frac{dF}{dA} = \frac{dF}{r dr d\phi} = c, \quad (8)$$

where c is the normalization constant. It is simply inverse of the disk area:

$$c = \frac{1}{\iint_{\text{disk}} dA} = \frac{1}{\pi(r_o^2 - r_i^2)} \quad (9)$$

Since there is no ϕ dependence, we can integrate it out and define the radial probability density function:

$$f(r) \equiv \frac{dF}{dr} = \int_0^{2\pi} d\phi r c = \frac{2r}{r_o^2 - r_i^2}, \quad (10)$$

which is identical to Eq. 5.

We can derive the same density function starting from our trusted uniform distribution. Consider two random variables X and Y uniformly distributed in the domain defined by $r_i^2 < x^2 + y^2 < r_o^2$ with the density $\frac{1}{\pi(r_o^2 - r_i^2)}$. We can define two new random variables $R = (X^2 + Y^2)^{1/2}$ and $\Phi = \text{sign}(Y) \arccos\left(\frac{X}{(X^2 + Y^2)^{1/2}}\right)$. The measure of the integral will transform with the Jacobian matrix

$$\frac{dx dy}{\pi(r_o^2 - r_i^2)} = \frac{1}{\pi(r_o^2 - r_i^2)} \left| \frac{d(x, y)}{d(r, \phi)} \right| dr d\phi = \frac{1}{\pi(r_o^2 - r_i^2)} \begin{vmatrix} \cos \phi & -r \sin \phi \\ \sin \phi & r \cos \phi \end{vmatrix} dr d\phi = \frac{r}{\pi(r_o^2 - r_i^2)} dr d\phi. \quad (11)$$

Upon integrating out ϕ , we pick up a factor of 2π , and get the same expression for $f(r)$ as we had in Eqs. 10 and 5.

We have shown in three different ways that the radius, R , will be a random variable with density $f(r)$. Then we take two such random radii and calculate their radial separation: $S \equiv |R_1 - R_2|$, which is yet another random variable. Figure 5 shows the simulation results for S , and it looks deceptively like a line. However, we do know it can't be an exact line because then it would have made it a triangle and a triangle has its center of mass at $1/3$ of the distance from the edge. But we have shown that the real answer deviates from $1/3$. So we should expect some higher order terms in play.

If we really wanted we could calculate the exact form of the density function of the random variable S . It is best to start from the cumulative distribution:

$$F_S(s) = P(|r_1 - r_2| < s) = \iint_{|r_2 - r_1| < s} dr_1 dr_2 f(r_1) f(r_2). \quad (12)$$

We need to figure out the domain for which $|r_1 - r_2| < s$ is satisfied. It is the green shaded area in Figure 7.




Figure 7: The domain of interest for integration. In the green shaded area $|r_2 - r_1| < s$ is satisfied.

Therefore, the cumulative probability function of the difference can be written as

$$\begin{aligned}
 F_S(s) &= \iint_{|r_2 - r_1| < s} dr_1 dr_2 f_R(r_1) f_R(r_2) = \iint_{\text{green}} dr_1 dr_2 f_R(r_1) f_R(r_2) \\
 &= 1 - \iint_{\text{red}} dr_1 dr_2 f_R(r_1) f_R(r_2) = 1 - 2 \int_{r_i}^{r_o - s} dr_1 f_R(r_1) \int_{r_1 + s}^{r_o} f_R(r_2) dr_2 \\
 &= 1 - 2 \int_{r_i}^{r_o - s} dr_1 f_R(r_1) [F_R(r_o) - F_R(r_1 + s)], \tag{13}
 \end{aligned}$$

from which we can get the probability density by differentiating with respect to s :

$$\begin{aligned}
 f_S(s) &= \frac{\partial}{\partial s} F_S(s) = 2 \int_{r_i}^{r_o - s} dr_1 f_R(r_1) f_R(r_1 + s) \\
 &= \frac{4}{3(r_o^2 - r_i^2)^2} [2(r_o^3 - r_i^3) - 3(r_o^2 + r_i^2)s + s^3]. \tag{14}
 \end{aligned}$$

This is the “model” line in Figure 5.

The distribution does differ from a line with a cubic term. $f_S(s)$ completely defines the statistics of the radial distance of two randomly selected sectors. We can verify that it reproduces the average value by computing $\frac{1}{r_o - r_i} \int_0^{r_o - r_i} s f(s) ds$, which indeed gives the result in Eq. 7.

Seek time

This post is titled “HDD average seek time”, but so far we have been computing distances. The relation between the distance and the time would have been in a one to one relation if the head moved with constant speed on the radial direction, none of which is exactly true. The head rotates around a pivot point that sits outside the disk, and follows an arc of a circle. Furthermore, it does not move with constant speed as it starts from zero speed, accelerates, slows down and settles at the destination track. Therefore the seek time includes higher order terms. However, we can use the ratio of average distance to the maximum distance as a proxy to estimate the ratio of average seek time to the time it would take the head to move between the opposite edges of the disk, i.e., the maximum seek time. I will create a separate post to compute the seek time in more detail. Stay tuned!

Conclusions

HDDs are complicated beasts and in this post we over-simplified few things. An important simplification we made was related to the assumption that the data is uniformly spread on the disk. This is not completely accurate. The capacity density of HDD surfaces has some nontrivial radial dependence [2]. This is an artifact of the skew of the read-write head with respect to the track direction. There are many other effects, which are beyond the scope of this post, that will cause nonuniform distribution of capacity and, in turn the data. In order to address such effects one needs to define $f(r)$ to incorporate the variations and do the math accordingly.

The bottom line is that it is fair to say that average seek time is about 1/3 of the full seek time with certain small and nontrivial deviations.

- [1] R. H. Arpaci-Dusseau and A. C. Arpaci-Dusseau, *Operating Systems: Three Easy Pieces*, 1.00 ed. <https://pages.cs.wisc.edu/~remzi/OSTEP/file-disks.pdf>; Arpaci-Dusseau Books, 2018.
- [2] M. A. Cordle *et al.*, “Radius and skew effects in an HAMR hard disk drive,” *IEEE Transactions on Magnetics*, vol. 52, no. 2, pp. 1–7, 2016, doi: [10.1109/TMAG.2015.2478115](https://doi.org/10.1109/TMAG.2015.2478115).