Data Placement Based on the Seek Time Analysis of a MEMS-based Storage Devices

Zachary N. J. Peterson
zachary@cse.ucsc.edu

Computer Systems Laboratory
Jack Baskin School of Engineering
University of California, Santa Cruz
1156 High Street
Santa Cruz, CA 95064

Abstract

Reducing access times to secondary I/O devices has long been the focus of many systems researchers. With traditional disk drives, access time is the composition of seek time and rotational latency, and so many techniques to order I/O requests or place data to minimize these factors have been developed. MEMS-based storage devices are seen by many as a replacement or an augmentation for modern disk drives, but algorithms for reducing access time for MEMS-based devices are still poorly understood. These devices, based on MicroElectroMechanical systems (MEMS), use thousands of active read/write heads working in parallel on a two-dimensional non-rotating magnetic substrate, eliminating rotational latency variable from the access time equation. This leaves seek time as the dominant variable. Therefore, new data layout techniques based on minimizing the unique seek time characteristics of a MEMS-based storage device must be developed. This paper examines the qualities of a MEMS-based storage device, and based on experimental simulation, develops an understanding of the seek time characteristics on such a device. These characteristics then allow us to identify equivalent regions in which to place data for improved access.

1 Introduction

Modern disk drives can no longer keep up with the performance trends of IC technology. RAM is growing, in terms of capacity and speed, 50% faster then disk, creating a performance bottleneck in computer systems. To compensate for this behavior, many techniques based on limiting the seek and rotational latency of a disk drive have been developed to improve disk, and therefore, system performance. These techniques include placing data on disk, based on workload, to reduce the time to access that data. However, because data layout on modern disk drives is a one-dimensional problem, applying analogous techniques to MEMS-based storage, an intrinsically two-dimensional device that incurs no rotational delay, suggests that these techniques will not take full advantage of the parallelism or additional degree of locality inherent in the device. For example, data written in contiguous sectors, the preferred allocation for sequential access on rotating media, will not be able to take advantage of the use of multiple read/write heads in this device. Griffin *et al.* have shown that simply applying traditional one-dimensional disk drive techniques to MEMS storage is possible, but rudimentary [3].

To better develop an understanding of the access time characteristics of a MEMS-device we have created a device simulator that generates seek times from any two points on the device. This information has enabled us to observe equivalence regions, or areas of media that share the same seek times from a fixed point, where like data can be stored to improve access times. We present results that are contradictory to those presented in the literature, and propose an enhanced data layout technique based on seek characteristics.

2 MEMS Background

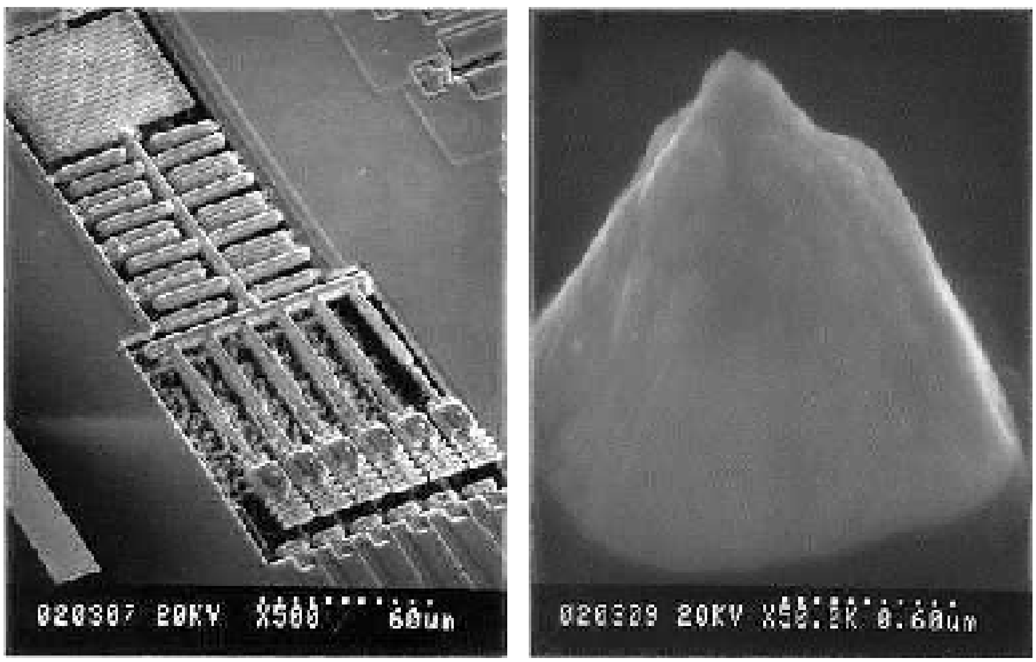

Figure 1: A grouping of probe tips on a cantilevered beam. Adjacent, a magnification of a single probe tip.

In this section we provide a brief physical description of a MEMS-based storage device. It is important to note that because this device is still in its developmental infancy, many of the details are still uncertain. We have based our experimental model on the specification from Carnegie Mellon University, described in “Modeling and Performance of MEMS-Based Storage Devices” [2] and “Designing Computer Systems with MEMS-based Storage” [1].

A MEMS-based storage is comprised of two main components: groups of probe tips called *tip arrays*, and a *media sled*. Probe tips are miniature read/write heads on the order of 1μm large. These tips are positioned at the end of a cantilevered beam. Grouping many beams together we create a *tip system*, which are similarly grouped into tip arrays. Figure 1 shows a picture of a tip system along side a single probe tip. Tip arrays are positioned over a movable sled coated in a magnetic substrate, to which multiple probe tips [1] can read or write bits in parallel using orthogonal magnetic recording techniques. The sled moves in the x and y directions, translated by electrostatic and spring forces, providing a two-dimensional surface for data placement.

2.1 Low Level Layout

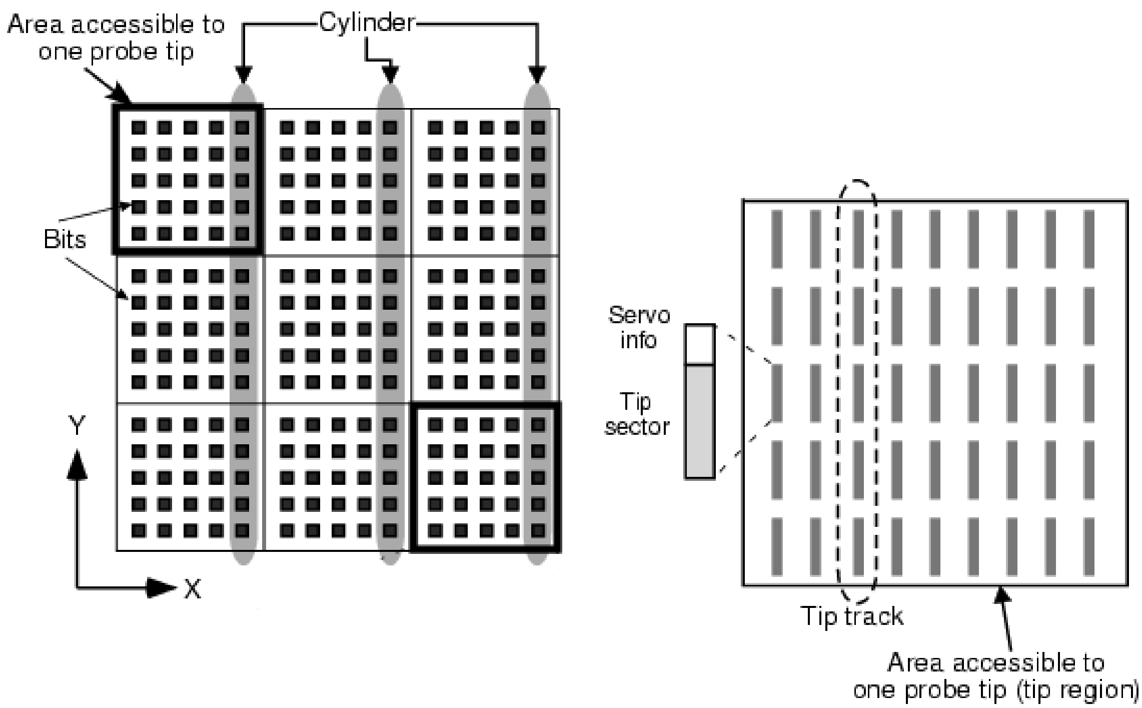

Figure 2: The low level data layout on a MEMS device.

Figure 2 illustrates the low level data layout of MEMS-based storage device. The media sled is logically broken into *tip regions*, defined by the area that is accessible by a single probe tip, approximately 2000 by 2000 bits in size. Each tip region is then separated into *tip tracks*, or the full stride of single tip. The tip tracks are further separated into *tip sectors*, the smallest accessible unit, analogous to disk drive sectors. Tip sectors are indexed by the tuple $< x, y, tip >$, where x and y are some distance coordinate and tip is a tip number. Therefore, we define a tip track as all tip sectors who share the same x and tip, and a *cylinder* as all tip tracks who share the same x. To access data, the sled is positioned over the appropriate sectors and accessed in parallel by moving in the y direction.

[1] Because of power and heat considerations, not all tips will be able to be active at the same time. Instead, 200 to 2000 simultaneously active tips are expected.

2.2 MEMS Model

For our experiments we use a MEMS-device simulator to generate the seek time from one sector to another. In our simulator we implemented a device model developed at the University of California, Santa Cruz (UCSC) by Yang [8]. In brief, Yang models the classical physical and mechanical forces that affect the positioning and settling times of the sled using the standard physics differential equation:

$$m\,\ddot{x} + \lambda\,\dot{x} + kx = F(t)$$

that has the solution:

$$\begin{aligned} x_t &= C_3 e^{-rt}\cos\omega t + C_4 e^{-rt}\sin\omega t + \frac{F}{k} \qquad (1) \\ C_3 &= -\frac{F}{k} \\ C_4 &= \frac{rC_3}{\omega} \end{aligned}$$

Equation 1 allows us, in particular, to model the spring like behavior exhibited when positioning the media sled.

2.3 Related Work

Optimizing data layout for secondary storage devices based on access times has long been the focus of many research projects. For instance, Worthington *et al.* produced algorithms that have been optimized to order requests based on the seek time of the disk drive head in relation to the rotating media under it [7]. Additionally, similar studies have been done in relation to rotational latency by Jacobson and Wilkes [4].

Beyond the research being performed at UCSC in MEMS-based storage, much of the software and hardware development is being done at Carnegie Mellon University [1–3, 5]. The important difference in their research effort involves they're access time model, which assumes not only an ideal acceleration but also ignores dampening and restoring spring forces. Using they're polynomial seek time equations, they have looked at the applied aspects of the integration of such a device in memory hierarchy, as well as the overall software design and performance. IBM has developed a prototype device [6], that unlike the CMU model, has a media sled that moves in the z direction. This enables data to be written using tiny physical marks on the media, as opposed to magnetic recording used in the CMU model. Further hardware research is also being done at HP labs and the Sandia National Laboratory.

3 Seek Time Analysis

To develop an understanding of the seek time characteristics of a MEMS-based device, we conducted experiments using the model described in Section 2.2. By fixing a starting sector, we could then determine the physical distance and seek time to all other sectors on the device. We choose to examine three sectors in particular: the top-left, the center -left, and dead-center. Because the device is symmetrical in nature, these position map to all corners of the devices.

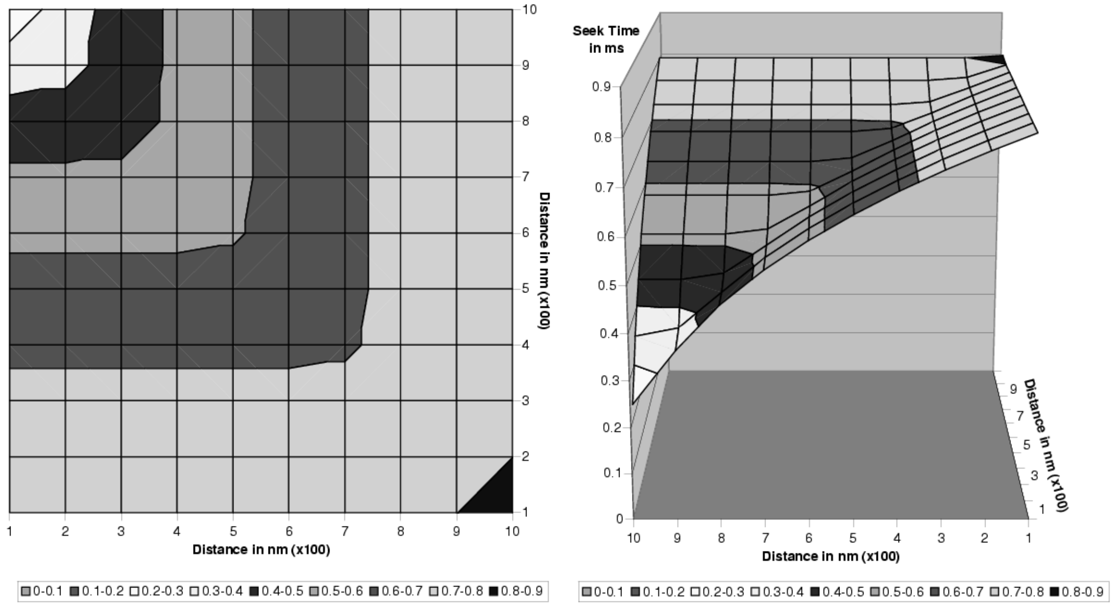

Figure 3: Equivalence regions for the top-left sector.

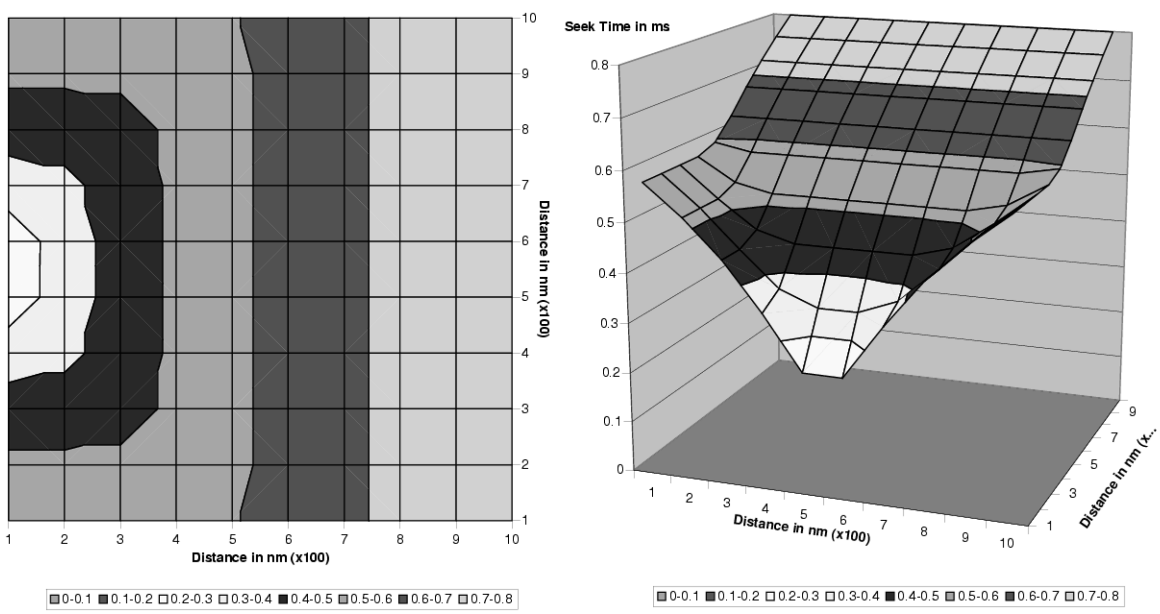

Figure 4: Equivalence regions for the center-left sector.

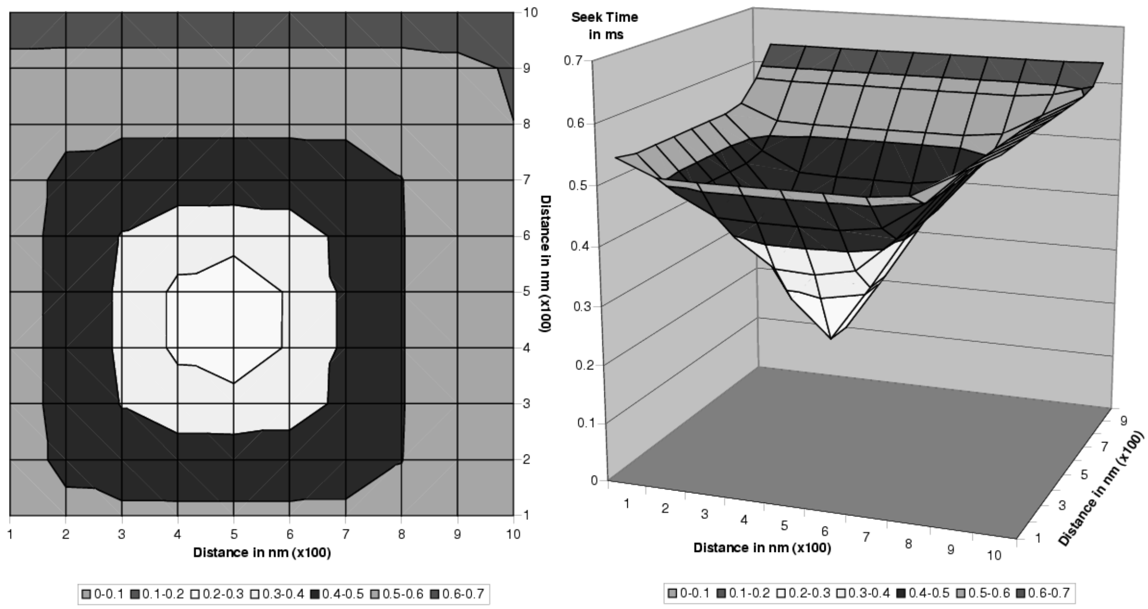

Figure 5: Equivalence regions for a near-center sector.

Figure 3 illustrates the seek time verses distance results gathered by setting the fixed sector to the top-left corner, and then seeking to all other sectors on the device. The various shades and, in the case of the 3D graph, levels, represent equivalent seek time regions based on a 0.1 ms granularity. There is a clear rectangular stratification, where layers of equivalent seek time regions exist on the media sled. Intuitively, seek times increase as the destination sector moves away from the starting sector. It is interesting to note that the regions grow thicker proportional to the distance away from the starting sector.

Figure 4 shows the surface and 3D seek times to all sectors, starting from the center-left sector. This data shows the same stratification as seen in Figure 3. What is more apparent in this surface graph, however, is the very rectangular relationship the equivalence regions have to the starting sector. Said another way, there exists a ratio of width and heighth where data benefits from being place within that region. Figure 5 is a near-center representation of the same seek time analysis, and clearly shows the very rectangular nature of the equivalence regions. These results are contradictory to those in the literature, as the device does not exhibit any benefit from a linear layout, similar to modern disk drives.

3.1 Data Layout

Based on the results from Section 3, it is clear that a direct application of current disk drive data layout will be insufficient, and possibly detrimental to the performance of a MEMS-based storage device. It is our affinity to suggest a data layout policy the prefers rectangular placement over track and cylinder placement. By developing a heuristic for the cost of placing data in the x direction versus the y, as previous thought, we believe we can achieve greater performance from this two dimensional device.

4 Future Work

To further back our claim of rectangular stratification, we will need to perform similar seek time analysis on future generations of MEMS-device models as more information about the actual device characteristics are revealed. Additionally, we are working on an algorithm that will base its placement decision on the data presented in this paper. Furthermore, results comparing that algorithm to standard layout techniques will be necessary.

5 Conclusion

By developing seek time characteristics of a MEMS-based storage device, we were able to show that data placement is not as simple as applying current techniques to this new and powerful device. Through device simulation we observed a rectangular stratification of seek times, or a physical regions of the device that shared similar seek times. From this information, we can conclude that applying traditional data layout techniques to this device is non-optimal, and that new policies must be developed. We recommend a two dimensional algorithm that uses heuristics based on this research to make more intelligent decision when placing data.

Acknowledgments

Thank you to Feng Weng for his help in understanding the MEMS-device simulator. Also thanks to the CMU PDL who have been gracious in sharing their data and hardware specifications with us. Finally, thanks to Dr. Scott Brandt who provided the opportunity and insight to develop this idea.

References

[1] L. R. Carley, Gregory Ganger, and David Nagle. Designing computer systems with MEMS-based storage. *Communications of the ACM*, 43(11), November 2000.

[2] John L. Griffin, Steven W. Schlosser, Gregory R. Ganger, and David F. Nagle. Modeling and performance of MEMS-based storage devices. In *Proceedings of the 2000 SIGMETRICS Conference*, pages 56–65, June 2000.

[3] John L. Griffin, Steven W. Schlosser, Gregory R. Ganger, and David F. Nagle. Operating system management of MEMS-based storage devices. In *Proceedings of the 4th Symposium on Operating Systems Design and Implementation (OSDI)*, pages 227–242, October 2000.

[4] David M. Jacobson and John Wilkes. Disk scheduling algorithms based on rotational position. Technical Report HPL-CSP-91-7rev1, Hewlett-Packard Laboratories, Concurrent Systems Project, March 1992.

[5] S. Schlosser, J. Griffin, Nagle, and G. Ganger. Designing computer systems with MEMS-based storage. In *Proceedings of the 9th Interational Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 1–12, 2000.

[6] P. Vettiger, M. Despont, U. Drechsler, U. Urig, W. Aberle, M. Lutwyche, H. Rothuizen, R. Stutz, R. Widmer, and G. Binnig. The ”millipede”– more than one thousand tips for future afm data storage. *IBM Journal of Research and Developmen*, 44(3):323–340, 2000.

[7] Bruce Worthington, Gregory Ganger, and Yale Patt. Scheduling algorithms for modern disk drives. In *Proceedings of the 1994 SIGMETRICS Conference*, pages 241–251, May 1994.

[8] Pu Yang. Modeling probe-based data storage. Master’s thesis, University of California, Santa Cruz, June 2000.