An Analytical Solution to a MEMS Seek Time Model

Bo Hong
hongbo@cse.ucsc.edu

Scott A. Brandt
sbrandt@cs.ucsc.edu

Storage Systems Research Center
University of California Santa Cruz

1 Introduction

The huge disparity between memory access times and disk access times has been the subject of extensive research. CPU speed has been increasing rapidly but disk access latency has lagged behind—disk transfer rates have been increasing at 40% per year, while seek times and rotational latency have been increasing at less than 10% per year [4]. This disparity has created a performance bottleneck in computer systems. Many techniques based on limiting the seek and rotational latency of a disk drive have been developed to improve disk, and therefore system, performance [5, 7, 8, 10, 12].

A new class of secondary storage devices based on microelectromechanical systems (MEMS) [1, 6, 11] currently being developed promise seek times 10–20 times faster than hard drives, storage densities 10 times greater, and power consumption an order of magnitude lower. MEMS devices provide non-volatile storage using either physical [11] or magnetic [1] recording techniques to achieve extremely high density storage. In order to achieve these high densities, MEMS-based storage designs use a non-rotating storage device with storage media on one surface and a large array of read/write heads on another surface directly above the storage media. By moving the surfaces relative to each other using MEMS actuators, each read/write head can access a region of the surface.

Because of their high density, high parallelism, and rectilinear two–dimensional motion, MEMS–based storage devices have unique performance characteristics, compared with traditional rotating media. It requires that standard file system algorithms for disks, including scheduling and layout, must be revisited to determine their efficiency domain. Unfortunately, these devices do not yet exist. Therefore, system performance analysis must depend on accurate and tractable models.

This paper provides an analytical solution to the MEMS seek time model from Carnegie Mellon University (CMU) [2]. The CMU positioning model takes into account the external force (constant but bidirectional, $\pm F$), the spring force, and the initial and final access velocities, which are opposite for odd and even–indexed bit columns. However, Griffin *et al.* [2] used an iterative, rather than analytical, approach to solve the model, which is difficult to apply in practice.

2 Background

It is important to note that because MEMS-based storage devices are still in their infancy, many of the details are still uncertain. However, there are several proposed architectures [2, 6, 9, 11], and we have based the physical parameters of our seek time model on the specification from CMU [2, 9].

A MEMS-based storage device is comprised of two main components: groups of probe tips called *tip arrays* that are used to access data on a movable *media sled*. In a modern disk drive, data is accessed by

device capacity	3.2 GB
number of tips	6400
maximum concurrent tips	1280
sled mobility in x and y	100 μm
sled acceleration in x and y	803.6 m/s^2
sled access speed	28 mm/s
sled resonant frequency	739.0 Hz
spring factor	75%
media bit cell size	40×40 nm
bits per tip region (M×N)	2500×2500

Table 1: Default MEMS-based storage device parameters.

means of an arm that seeks in one dimension above a rotating platter. In a MEMS device, the entire media sled is positioned in the x and y directions by electrostatic forces while the heads remain stationary.[1] Another major difference between a MEMS-based storage device and a disk is that on a MEMS device, multiple tips can be active at the same time. Data can be then be striped across multiple tips, allowing a considerable amount of parallelism. However, power and heat considerations limit the number of probe tips that can be active simultaneously; it is estimated that 200 to 2000 probes will actually be active at once.

Figure 1 illustrates the low level data layout of a MEMS-based storage device. The media sled is logically broken into *tip regions*, defined by the area that is accessible by a single head, approximately 2000 by 2000 bits in size. Each tip in the MEMS device can only read the data in its own tip region; this limits the maximum sled movement to the dimensions of a single tip region. The smallest unit of data in a MEMS-based storage device is called a *tip sector*. Each tip sector, identified by the tuple $\langle x, y, tip \rangle$, has its own servo information for positioning as well as its own error correction information. The set of bits accessible to a single tip with the same x coordinate is called a *tip track*, and the set of all bits (under all tips) with the same x coordinate is referred to as a *cylinder*. Also, the set of tip sectors that can be accessed by simultaneously active tips is known as a *logical sector*. For faster access, disk sectors can be striped across logical sectors.

Table 1 shows default MEMS–based storage device parameters. The physical constants used in our analytical solution are given or can be easily derived from these design parameters.

3 Modeling Seek Time of a MEMS–Based Storage Device

In our work, we use the positioning model and physical parameters from CMU [2, 3]. The CMU positioning model takes into account the external force (constant but bidirectional, $\pm F$), the spring force, and the initial and final access velocities, which are opposite for odd and even–indexed bit columns. Griffin *et al.* [2] used an iterative approach to solve the model, which is unlikely to be employed in practice. In this section, we propose an analytic solution to the CMU model.

Because the actuation mechanisms and control loops for x and y positioning are independent in MEMS–based storage devices, positioning in the x and y dimensions can proceed in parallel. Therefore,

$$t_{seek} = \max(t_x, t_y), \tag{1}$$

where t_{seek} is the seek time and t_x and t_y are the seek times in the x and y dimensions. A seek in the x and y dimensions consists of a base seek plus a settling time in the x dimension and turnaround times in

[1] Some MEMS storage device designs, like the IBM Millipede, fix the sled and move the heads. The effect is the same—the heads move relative to the media.

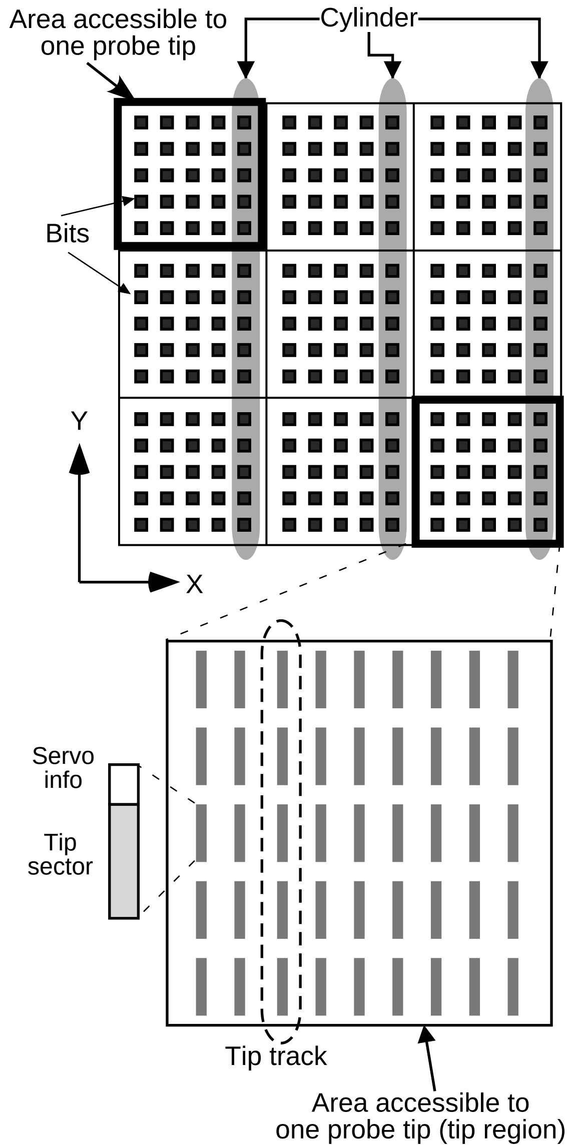

Figure 1: Data layout on a MEMS device.

the y dimension. Therefore, t_x consists of a base seek plus a settling time, t_{settle}, a function of the resonant frequency of the system, and t_y consists of a base seek plus a turnaround time, $t_{turnaround}$, a function of the actuator and spring forces. t_{settle} and $t_{turnaround}$ are given in Equations 2 and 3:

$$t_{settle} = \frac{1}{2\pi f}, \tag{2}$$

$$t_{turnaround} = 2\frac{v_{access}}{a_{actuator} \pm a_{spring}}, \tag{3}$$

where t_{settle} is the settling time, $t_{turnaround}$ is the turnaround time, f is the resonant frequency, v_{access} is the access velocity, and $a_{actuator}$ and a_{spring} are accelerations by the actuators and springs.

Assume that the initial and final positions are x_0 and x_1, respectively, and the initial and final velocities are v_0 and v_1, respectively. In a base seek, we also assume that $x_0 < x_1$ and $v_0 = v_1 = 0$ in the x dimension and $v_0 = v_1 > 0$ in the y dimension. A base seek consists of two phases: acceleration and deceleration.

The actuators accelerate the sled toward the destination in the acceleration phase and reverse polarity and decelerate the sled to its final destination and velocity in the deceleration phase. In addition to the actuator force, the sled springs constantly pull the sled toward its center-most position. It is therefore essential to know when and where to reverse the polarity of the actuators.

Because the kinetic energy of the sled is unchanged at the beginning and the end of a base seek, we have

$$0 = (x_m - x_0)F - (x_1 - x_m)F - \tfrac{k}{2}(x_1^2 - x_0^2),$$

where F is the absolute force by actuators, k is the spring constant, and x_m is the position at which actuators reverse polarity, from positive to negative. Therefore, x_m is given in Equation 4:

$$x_m = \frac{x_0 + x_1}{2} + \frac{k}{4F}(x_1^2 - x_0^2). \tag{4}$$

3.1 Seek Time in the X Dimension

Because the initial and final velocities in the x dimension are zero, no turnaround time needs to be considered. However, due to the rapid acceleration and deceleration of the sled, the spring–sled system momentarily oscillate in the x dimension before damping to $v_x = 0$. Therefore, extra settling time must be taken into account. Assume that the initial and final positions are x_0 and x_1, respectively. We also assume that $x_0 < x_1$ (if $x_0 > x_1$, let $x_0 = -x_0$ and $x_1 = -x_1$).

Because a base seek is divided into an acceleration and an deceleration phase,

$$t_x \quad = \quad t_{xa} + t_{xd} + t_{settle}, \tag{5}$$

where t_{xa} and t_{xd} are the times elapsed in the acceleration and deceleration phase, respectively.

The equation of describing the movement in the acceleration phase is as follows:

$$\ddot{x} \quad = \quad a - \frac{kx}{m},$$

where a is the acceleration by actuators and m is the mass of the sled. Because $x(0) = x_0$ and $\dot{x}(0) = 0$, we have

$$x(t) \quad = \quad (x_0 - \frac{ma}{k})\cos(\sqrt{\frac{k}{m}}t) + \frac{ma}{k}. \tag{6}$$

Because the position where the acceleration phase ends, x_m, is given by Equations (4),

$$x_m \quad = \quad \frac{x_0 + x_1}{2} + \frac{k}{4ma}(x_1^2 - x_0^2). \tag{7}$$

Therefore, we have

$$\begin{aligned} x_m &= x(t_{xa}) \\ &= (x_0 - \frac{ma}{k})\cos(\sqrt{\frac{k}{m}}t_{xa}) + \frac{ma}{k}. \end{aligned} \tag{8}$$

Combining Equations (7) and (8), we have

$$t_{xa} \quad = \quad \sqrt{\frac{m}{k}}\arccos(\frac{x_m - \frac{ma}{k}}{x_0 - \frac{ma}{k}}). \tag{9}$$

The equation of describing the movement in the deceleration phase is as follows:

$$\ddot{x} \quad = \quad -a - \frac{kx}{m},$$

whose general solution is

$$x(t) \quad = \quad c_1 \cos\sqrt{\frac{k}{m}}t + c_2 \sin\sqrt{\frac{k}{m}}t - \frac{ma}{k}. \tag{10}$$

Because $x(0) = x_m$, we have

$$c_1 \quad = \quad x_m + \frac{ma}{k}. \tag{11}$$

Because $x(t_{xd}) = x_1$ and $\dot{x}(t_{xd}) = 0$, we have

$$c_1 \quad = \quad (\frac{ma}{k} + x_1)\cos\sqrt{\frac{k}{m}}t_{xd}. \tag{12}$$

Combining Equations (11) and (12), we have

$$t_{xd} \quad = \quad \sqrt{\frac{m}{k}}\arccos(\frac{x_m + \frac{ma}{k}}{x_1 + \frac{ma}{k}}). \tag{13}$$

Combining Equations (2), (9), and (13), we have

$$t_x \quad = \quad \sqrt{\frac{m}{k}}\arccos(\frac{x_m - \frac{ma}{k}}{x_0 - \frac{ma}{k}}) + \sqrt{\frac{m}{k}}\arccos(\frac{x_m + \frac{ma}{k}}{x_1 + \frac{ma}{k}}) + \frac{1}{2\pi f}. \tag{14}$$

3.2 Seek Time in the Y Dimension

Because the initial and final access velocities of the sled in the y dimension are not zero, the sled may turn around once or twice to the specified velocity in the right direction in addition to position itself at the right spot. Therefore, we must take into account turnaround times when modeling MEMS seek time in the y dimension.

Table 2 summarizes the number of turnarounds needed in different conditions, which depends on the direction of the sled movement and the directions of initial and final access velocities. For a base seek in the y dimension (i.e., $y_0 < y_1$, $v_0 > 0$, and $v_1 > 0$), no turnaround is needed. We can calculate turnaround time using Equation 3.

Let us consider the base seek in the y dimension. Assume that the initial and final positions are y_0 and y_1, respectively, and the initial and final access velocities are v_0 and v_1, respectively. For a base seek, we also assume that $y_0 < y_1$, $v_0 > 0$, and $v_1 > 0$ (if $y_0 > y_1$, let $y_0 = -y_0$, $x_1 = -x_1$, $v_0 = -v_0$, and $v_1 = -v_1$).

Because a base seek is divided into an acceleration and an deceleration phase,

$$t_y \quad = \quad t_{ya} + t_{yd} + t_{turnaround}, \tag{15}$$

where t_{ya} and t_{yd} are the times elapsed in the acceleration and deceleration phase, respectively.

The equation of describing the movement in the acceleration phase is as follows:

$$\ddot{y} \quad = \quad a - \frac{ky}{m},$$

Number of Turnarounds	Conditions		
	$y_0 < y_1$	$v_0 > 0$	$v_1 > 0$
0	Y	Y	Y
1	Y	Y	N
1	Y	N	Y
2	Y	N	N
2	N	Y	Y
1	N	Y	N
1	N	N	Y
0	N	N	N

Table 2: Number of turnarounds in different conditions.

Because $y(0) = y_0$ and $\dot{y}(0) = v_0$, we have

$$y(t) = (y_0 - \frac{ma}{k})\cos(\sqrt{\frac{k}{m}}t) + v_0\sqrt{\frac{m}{k}}\sin(\sqrt{\frac{k}{m}}t) + \frac{ma}{k}. \quad (16)$$

Because the position where the acceleration phase ends, y_m, is given by Equations (4),

$$y_m = \frac{y_0 + y_1}{2} + \frac{k}{4ma}(y_1^2 - y_0^2). \quad (17)$$

Therefore, we have

$$\begin{aligned} y_m &= y(t_{ya}) \\ &= (y_0 - \frac{ma}{k})\cos(\sqrt{\frac{k}{m}}t_{ya}) + v_0\sqrt{\frac{m}{k}}\sin(\sqrt{\frac{k}{m}}t_{ya}) + \frac{ma}{k}. \end{aligned} \quad (18)$$

Combining Equations (17) and (18), we have

$$t_{ya} = \sqrt{\frac{m}{k}} \times (\arcsin(\frac{y_m - \frac{ma}{k}}{\sqrt{(y_0 - \frac{ma}{k})^2 + (v_0\sqrt{\frac{m}{k}})^2}}) - \arcsin(\frac{y_0 - \frac{ma}{k}}{\sqrt{(y_0 - \frac{ma}{k})^2 + (v_0\sqrt{\frac{m}{k}})^2}})). \quad (19)$$

The equation of describing the movement in the deceleration phase is as follows:

$$\ddot{y} = -a - \frac{ky}{m},$$

whose general solution is

$$y(t) = c_1\cos\sqrt{\frac{k}{m}}t + c_2\sin\sqrt{\frac{k}{m}}t - \frac{ma}{k}. \quad (20)$$

Because $y(0) = y_m$, we have

$$c_1 = y_m + \frac{ma}{k}. \quad (21)$$

Because $y(t_{yd}) = y_1$ and $\dot{y}(t_{yd}) = v_1$, we have

$$c_1 = (y_1 + \frac{ma}{k})\cos(\sqrt{\frac{k}{m}}t_{yd}) - v_1\sqrt{\frac{m}{k}}\sin(\sqrt{\frac{k}{m}}t_{yd}). \quad (22)$$

Combining Equations (21) and (22), we have

$$t_{yd} = \sqrt{\frac{m}{k}} \times (\arcsin(\frac{y_1 + \frac{ma}{k}}{\sqrt{(y_1 + \frac{ma}{k})^2 + (v_1\sqrt{\frac{m}{k}})^2}}) - \arcsin(\frac{y_m + \frac{ma}{k}}{\sqrt{(y_1 + \frac{ma}{k})^2 + (v_1\sqrt{\frac{m}{k}})^2}})). \quad (23)$$

Combining Equations (3), (19), and (23), we have

$$\begin{aligned} t_y = & \sqrt{\frac{m}{k}} \arcsin(\frac{y_m - \frac{ma}{k}}{\sqrt{(y_0 - \frac{ma}{k})^2 + (v_0\sqrt{\frac{m}{k}})^2}}) - \sqrt{\frac{m}{k}} \arcsin(\frac{y_0 - \frac{ma}{k}}{\sqrt{(y_0 - \frac{ma}{k})^2 + (v_0\sqrt{\frac{m}{k}})^2}}) \\ & + \sqrt{\frac{m}{k}} \arcsin(\frac{y_1 + \frac{ma}{k}}{\sqrt{(y_1 + \frac{ma}{k})^2 + (v_1\sqrt{\frac{m}{k}})^2}}) - \sqrt{\frac{m}{k}} \arcsin(\frac{y_m + \frac{ma}{k}}{\sqrt{(y_1 + \frac{ma}{k})^2 + (v_1\sqrt{\frac{m}{k}})^2}}) \\ & + t_{turnaround}. \end{aligned} \quad (24)$$

4 Conclusions

In this paper, we provided an analytical solution of the CMU positioning model of the MEMS–based storage device. Because MEMS–based storage devices do not yet exist, models that bridge the gap between the physical and performance characteristics of the device provide important feedback to hardware and software designers. Our solution of the CMU positioning model provides a solid starting point for the MEMS system design and performance evaluation.

References

[1] L. Carley, J. Bain, G. Fedder, D. Greve, D. Guillou, M. Lu, T. Mukherjee, S. Santhanam, L. Abelmann, and S. Min. Single-chip computers with microelectromechanical systems-based magnetic memory. *Journal of Applied Physics*, 87(9):6680–6685, May 2000.

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

[3] J. L. Griffin, S. W. Schlosser, G. R. Ganger, and D. 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, Oct. 2000.

[4] J. L. Hennessy and D. A. Patterson. *Computer Architecture A Quantitative Approach*. Morgan Kaufmann Publishers, 3rd edition, 2003.

[5] D. M. Jacobson and J. Wilkes. Disk scheduling algorithms based on rotational position. Technical Report HPL-CSP-91-7rev1, Hewlett-Packard Laboratories, Concurrent Systems Project, Mar. 1992.

[6] Nanochip Inc. Nanochip: Array nanoprobe mass storage IC. http://www.nanochip.com/, 1999.

[7] J. Peacock. The counterpoint fast file system. In *Proceedings of the Winter 1988 USENIX Technical Conference*, pages 243–249. USENIX Association, Jan. 1988.

[8] M. Rosenblum and J. K. Ousterhout. The design and implementation of a log-structured file system. *ACM Transactions on Computer Systems*, 10(1):26–52, Feb. 1992.

[9] S. W. Schlosser, J. L. Griffin, D. F. Nagle, and G. R. 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, Cambridge, MA, Nov. 2000.

[10] E. Shriver, E. Gabber, and L. H. C. A. Stein. Storage management for web proxies. In *Proceedings of the 2001 USENIX Annual Technical Conference*, pages 203–216. USENIX, June 2001.

[11] 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 Development*, 44(3):323–340, 2000.

[12] B. Worthington, G. Ganger, and Y. Patt. Scheduling algorithms for modern disk drives. In *Proceedings of the 1994 SIGMETRICS Conference*, pages 241–251, May 1994.