Shutdown Policies for MEMS-Based Storage Devices – Analytical Models

Mohammed G. Khatib, Johan B.C. Engelen and Pieter H. Hartel
University of Twente
Enschede, the Netherlands
{m.g.khatib | j.b.c.engelen | p.hartel}@utwente.nl

Abstract

MEMS-based storage devices should be energy efficient for deployment in mobile systems. Since MEMS-based storage devices have a moving media sled, they should be shut down during periods of inactivity. However, shutdown costs energy, limiting the applicability of aggressive shutdown decisions.

The media sled in MEMS-based storage devices is suspended by springs. We introduce a policy that exploits the spring structure to reduce the shutdown energy. As a result, the aggressiveness of the shutdown decisions can be increased, reducing the energy consumption. This report devises analytical models of the shutdown time and energy of this policy.

1 Introduction

A new class of storage based on parallel-probe storage using Micro-Electro-Mechanical Systems (MEMS) has been proposed [1]. Enabled by high storage densities (> 1 Tb/in^2), MEMS technology promises to deliver small-form factor, high-capacity, and low-power storage devices. MEMS devices have potentially lower cost than flash, because MEMS devices have fewer requirements on the lithography. A MEMS-based storage device dissipates an order of magnitude less power than a disk drive. However, like disk drives, MEMS devices have a moving medium. For optimal energy saving, a MEMS device should be shut down (i.e., the sled is stopped) during periods of inactivity. Because of their micro-mechanical nature, MEMS devices lend themselves to more aggressive shut down policies than disk drives.

As the timeout decreases, the number of shutdowns increases, because more periods of inactivity are exploited. As a result, the energy consumed to shut down increases and so the does total energy, which limits the applicability of aggressive shutdown decisions. MEMS devices have a moving medium that is suspended by springs across the read/write probe (head) array as shown in Figure 1a. We show that the potential energy stored in the springs can be exploited at shutdown to accelerate the sled toward a resting position (i.e., the center). External energy is invested only to decelerate the sled, so that it stops at the center. Consequently, the shutdown energy is reduced to the deceleration energy, allowing to increase the aggressiveness of the shutdown decisions. We

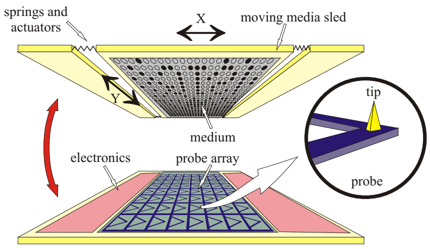

(a)

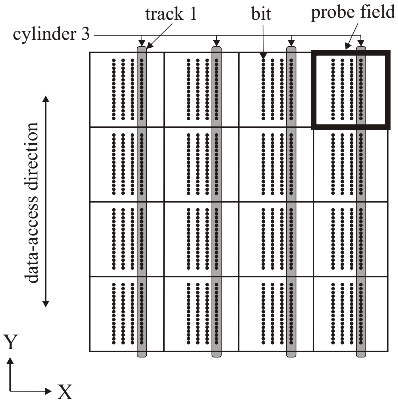

(b)

Figure 1: Three- and two-dimensional views of a MEMS-based storage device. *(a) Two layers facing each other where the media sled is attached to springs that suspend it across the probe array. (b) The storage area of a simplified MEMS-based storage device consisting of 4×4 probes. The storage area is logically divided into 16 storage fields each accessible by a single probe.*

call this policy the energy-efficient policy. This energy benefit, however, comes at a performance cost; that is the sled takes long time to reach the center, since it is not actively accelerated.

Another possible shutdown policy uses the actuators to accelerate the sled some distance and then decelerate it, so that it reaches the center in the shortest time possible, called the performance-efficient policy. When deploying this policy, the sled performs as if it seeks from its current position to the center with the exception that it stops on Y as well as on X at the center.

This study devises analytical models of the shutdown time and energy of the energy-efficient policy. It offers a modification to the analytical seek model devised by Hong et al. [2] to model the shutdown time of the performance-efficient policy.

The next section provides some background on MEMS-based parallel-probe data storage. Section 3 introduces an energy-efficient policy and a performance-efficient policy to shut down the sled. Section 5 devises the analytical model of the shutdown time and energy of the energy-efficient policy. Section 6 offers the necessary modification of Hong's seek model [2] to model the performance-efficient policy.

2 MEMS-Based Storage

Several design models for MEMS-based storage have been proposed [1, 3, 4, 5]. Although these models adopt different storage and actuation techniques, they have a common architecture. A MEMS-based storage device consists of two distinct physical layers, one above the other, as shown in Figure 1a. The top layer, called the *media sled*, is suspended by springs above the bottom layer, where the Z distance is maintained by nanopositioners. The bottom layer is a two-dimensional array of read/write probes or heads, called the *probe array*. For example, an IBM prototype [1] has a 64×64 probe array. Probes can be clustered in groups to reduce the complexity of the circuitry.

The top layer is the media sled, where data bits are recorded. Bits can be recorded on a magnetic patterned medium as in µSPAM [4] and CMU MEMStore [3]; a polymer medium as in the IBM MEMS device [1]; or a phase-change medium as in the Nanochip MEMS device [5]. The sled moves independently in the X, Y, and Z directions relative to the probe array. In all design models, each probe sweeps over a bounded area of the media sled, called the *probe (storage) field* as sketched in Figure 1b. Consequently, seek times shorten and a relatively high (aggregate) data rate is attainable by operating many probes simultaneously, so that each probe accesses a small part of a sector, called a *subsector*.

To read or write data from/to the medium, the media sled moves along the Y direction, along which data tracks lie as shown in Figure 1b. While accessing data, the X actuators keep the sled still along the X direction on the accessed data track, counteracting the spring restoring force. When resting, the springs hold the sled at its resting position, where every probe faces the center position of its probe field.

As of early 2008, to our best knowledge, an IBM prototype [1] records a single bit in an area of 26 nm by 26 nm, whereas Nanochip [5] claims a 15 nm by 15 nm bit cell area with the potential to reach a scale of 2 nm by 2 nm. With such high densities, a single memory chip has a capacity of 1 TB per die. These devices have potentially low cost for three reasons. Firstly, they can be manufactured using the well-established batch MEMS fabrication technology [1]. Secondly, these devices can be manufactured using micron fabrication plants, because the minimum feature size for MEMS is two orders of magnitude higher than Flash chips (micrometers instead of nanometers). The equipment of micron plants were installed ten years ago and passed their break-even point, avoiding the need to build dedicated fabrication plants, unlike for flash memory. Thirdly, these plants can be used to make future generations of MEMS, since MEMS poses no requirements on the lithography process when increasing the density [5].

3 Two Shutdown Policies

To move the sled to the center, two different policies are possible. The first uses the potential energy stored in the springs to bring the sled as close as possible to the center (Figure 2 *top*) before the actuators starts to decelerate the sled in order to stop it at the center. This policy consumes less energy than the other policy, since it invests no energy in acceleration; we call it the *energy-efficient shutdown policy*.

The second policy uses the actuators to accelerate the sled some distance and then decelerate it to stop at the center. Here, energy is invested in acceleration as well as deceleration. Figure 2 *bottom* shows that because external energy is invested in acceleration, deceleration starts earlier compared to the previous policy in order to do enough counter work to bring the sled stationary at the center. The external energy invested speeds up the shutdown and therefore this

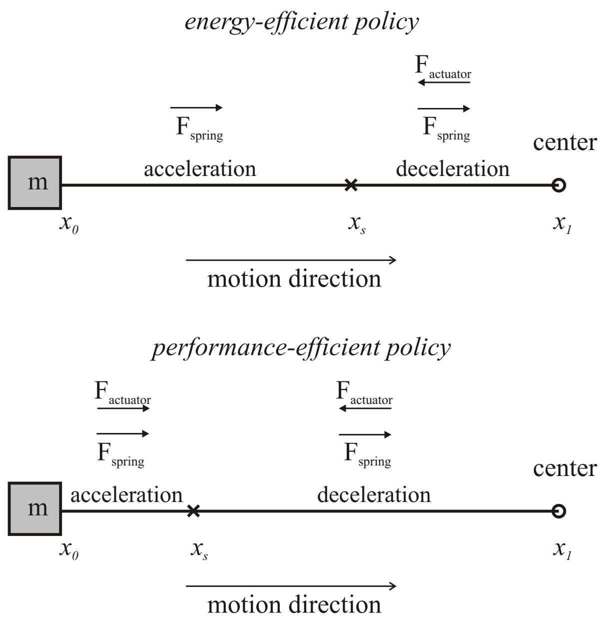

Figure 2: The energy-efficient shutdown policy versus the performance-efficient one. Unlike the former policy, the latter uses the actuators for acceleration.

policy is called the *performance-efficient shutdown policy*.

4 Modeling

Madhyastha and Yang [6] compare three models of the motion of the media sled in MEMS-based storage devices. The models are: (1) a simple single-parameter model, (2) a spring model, and (3) an optimal control model. Madhyastha and Yang show that the optimal control model is more accurate than the simple and spring models, because it assumes no constant acceleration nor constant force during motion. In their simulations against real workloads, Madhyastha and Yang show that the simple and spring models overestimate the optimal seek time by approximately a factor of two and five, respectively.

The optimal control model is called the bang-bang model in optimal control theory, because it switches abruptly between two modes, namely acceleration and deceleration. It captures the dynamics of the system and factors in all forces during the sled motion [6,2]. Our analytical models are based on the bang-bang model to derive accurate shutdown time and thus energy. The bang-bang model is extended to account for non-zero initial velocity. Table 1 provides descriptions of the parameters of the model.

5 The Energy-Efficient Policy

In MEMS-based storage devices the media sled is suspended by springs across the probe array. If no external force is applied, the springs pull the medium

Table 1: List of the parameters of the analytical model

Parameter	Description	Unit
k	spring constant	kg
m	mass	kg
F	force due to actuators	N
a	acceleration due to actuators	m/s^2
a_s	acceleration due to springs	m/s^2
v_a	data access speed	m/s

back to its rest position, so that every probe faces the center of its storage area, the probe field.

To save investing external energy (i.e., from the host system), the energy-efficient shutdown policy makes use of the potential energy stored in the springs to accelerate the sled toward the center position. To bring the sled stationary at the center, we determine a *switching point*, on each direction, at which a counter force is applied by the actuators to decelerate the sled.

Thus, when deploying the energy-efficient policy, energy (from the host system) is invested only in deceleration and not in acceleration. In the following, we derive the analytical model for the motion along the X and Y directions separately; we start with X followed by Y.

5.1 Motion along X

We assume that the sled is at position x_0 and moves to the center $x_1 = 0$ and the switching point is denoted by x_s. The switching point is determined using the energy conservation equation as follows:

$$\begin{aligned} \sum E = 0 \quad &\Rightarrow \quad \frac{k_x}{2} x_0^2 - (x_s - x_1)F = 0 \\ &\Rightarrow \quad x_s = \frac{k_x}{2F} x_0^2, \end{aligned} \tag{1}$$

where F is the counter force applied by the actuators to stop the the sled at the center. The spring motion toward the center is split in two phases; the acceleration and the deceleration phase. The acceleration phase is between x_0 and x_s, whereas the deceleration phase is between x_s and x_1 (as depicted in Figure 2a).

Acceleration phase — According to Newton's second law, the motion of the sled in the acceleration phase is described as follows:

$$\begin{aligned} m_x \ddot{x} &= -k_x x \\ \Rightarrow \quad \ddot{x} &= -\frac{k_x}{m_x} x\,. \end{aligned}$$

This is an ordinary differential equation, whose general solution has the form:

$$x(t) = C_1 \cos\left(\sqrt{\frac{k_x}{m_x}} t\right) + C_2 \sin\left(\sqrt{\frac{k_x}{m_x}} t\right)$$

Since $x(t=0)=x_0$ and $\dot{x}(t=0)=0$, we get $C_1=x_0$ and $C_2=0$, respectively. Thus:

$$x(t)=x_0\cos(\sqrt{\frac{k_x}{m_x}}t) \tag{2}$$

Rearranging Equation (2), we get the acceleration time, t_{xa}, between x_0 and x_s:

$$t_{xa}=\sqrt{\frac{m_x}{k_x}}\arccos(\frac{x_s}{x_0})\,. \tag{3}$$

Deceleration phase — In the deceleration phase the actuators exert a force to stop the sled at the center, which works against the spring force. The motion is described as follows:

$$m_x\ddot{x}=-m_xa_x-k_xx$$
$$\Rightarrow\quad \ddot{x}=-a_x-\frac{k_x}{m_x}x$$

Here a_x is the acceleration due to the actuators and the minus sign before it indicates that it opposes the motion direction of the sled. This motion equation has the following general solution:

$$x(t)=C_1\cos(\sqrt{\frac{k_x}{m_x}}t)+C_2\sin(\sqrt{\frac{k_x}{m_x}}t)-\frac{m_xa_x}{k_x}$$

Since $x(t=0)=x_s$, we get:

$$C_1=x_s+\frac{m_xa_x}{k_x}\,.$$

Two boundary conditions exist, namely $x(t=t_{xd})=x_1$ and $\dot{x}(t=t_{xd})=0$, which we use to calculate the deceleration time, t_{xd}, and then C_2:

$$t_{xd}=\sqrt{\frac{m_x}{k_x}}\arccos(\frac{x_s+\frac{m_xa_x}{k_x}}{x_1+\frac{m_xa_x}{k_x}}), \tag{4}$$

and

$$C_2=C_1\tan(\sqrt{\frac{k_x}{m_x}}t_{x_d})\,.$$

The solution given in Equation (4) holds for $x_s<x_1$. When $x_s>x_1$, as in our case, since $x_1=0$ (the center), we multiply every x by -1, so that $-x_0<-x_1$.

Total shutdown time along X — The total shutdown time along X is:

$$t_{\text{x-total}}=t_{xa}+t_{xd}\,. \tag{5}$$

5.2 Motion along Y

We assume that the sled is at position y_0 and moves to the center $y_1=0$ and the switching point is denoted by y_s. Unlike the X direction, on which the sled stays still, the sled moves along Y to access data. Here, we distinguish two states of the sled when a decision is made to shut it down: (1) the sled moves along Y in the direction of the center (briefly it moves inward) or the sled moves toward the borders and thus outward.

If the sled moves inward, then an additional initial speed should be considered, namely its data access speed (v_a), whereas if it moves outward, it should brake first and then accelerates toward the center starting at a velocity of zero. Thus, when moving outward the sled motion is identical to that discussed for X and thus the shutdown time is calculated similarly to which an additional brake time is added:

$$t_{\text{brake}} = \frac{v_a}{a_s}\,. \tag{6}$$

In the following we study the first case, where the sled accelerates toward the center with an initial velocity v_a. The switching point (y_s) is:

$$\frac{k_y}{2}y_0^2 - (y_s - y_1)F - \frac{m_y}{2}v_a^2 = 0$$
$$\Rightarrow \quad y_s = \frac{k_y y_0^2 - m_y v_a^2}{2F} \tag{7}$$

Here F is the actuation force exerted by the actuators between y_s and y_1. The motion along Y is also split into an acceleration and deceleration phase.

Acceleration phase — The motion equation for this phase is:

$$\ddot{y} = -\frac{k_y}{m_y}y$$

The solution for this equation has the form:

$$y(t) = C_1 \cos(\sqrt{\frac{k_y}{m_y}}t) + C_2 \sin(\sqrt{\frac{k_y}{m_y}}t)\,.$$

Since $y(t=0) = y_0$ and $\dot{y}(t=0) = v_a$, we get $C_1 = y_0$ and $C_2 = v_a$, respectively. Thus:

$$y(t) = y_0 \cos(\sqrt{\frac{k_y}{m_y}}t) + v_a \sin(\sqrt{\frac{k_y}{m_y}}t) \tag{8}$$

Rearranging Equation (8), we get the acceleration time along Y, t_{ya}:

$$t_{ya} = \sqrt{\frac{m_y}{k_y}}(\arcsin(\frac{y_s}{\sqrt{y_0^2 + (v_a\sqrt{\frac{m_y}{k_y}})^2}}) - \arcsin(\frac{y_0}{\sqrt{y_0^2 + (v_a\sqrt{\frac{m_y}{k_y}})^2}}))\,. \tag{9}$$

We multiply every y by -1, so that $-y_s < -y_1$.

Deceleration phase — The deceleration phase along Y has the same initial (it starts at y_s) and boundary conditions ($y(t = t_{yd}) = y_1$ and $\dot{y}(t = t_{yd}) = 0$) like that along X. Thus, the deceleration time is:

$$t_{yd} = \sqrt{\frac{m_y}{k_y}}\arccos(\frac{y_s + \frac{m_y a_y}{k_y}}{x_1 + \frac{m_y a_y}{k_y}})\,. \tag{10}$$

We multiply every y by -1 so that $-y_s < -y_1$.

Total shutdown time along Y — The total shutdown time along Y is:

$$t_{\text{y-total}} = t_{ya} + t_{yd} + \begin{cases} t_{\text{brake}} & ; \text{if outward} \\ 0 & ; \text{if inward} \end{cases} \tag{11}$$

If the sled is moving outward at the shutdown decision (case 2), t_{ya} and t_{yd} are calculated using Equations (3) and (4), respectively, after replacing every parameter of X with its Y counterpart. Otherwise, if the sled is moving inward (case 1), we use Equations (9) respectively (10) instead. Also, y_s is calculated using Equation (1) or (7), respectively.

5.3 Shutdown Time and Energy

The motions along X and Y directions are independent, because of the independent actuators and their structure. Thus, the time to shut down is the maximum of the shutdown times along X and Y:

$$t_{\text{shutdown}} = \max(t_{\text{x-total}}, t_{\text{y-total}})\,. \tag{12}$$

Remember that no energy is invested to brake the sled if it moves outward nor energy is invested to accelerate the sled; we use the potential energy stored in the springs instead. The total energy is, therefore, the sum of the energy invested only to decelerate the sled along X and along Y. We drive the actuators with the maximum allowed current (thus the maximum power) to shorten the deceleration time. Thus, the total energy is:

$$E_{\text{shutdown}} = P_{\max}(t_{xd} + t_{yd})\,. \tag{13}$$

Again, we use Equation (4) to calculate t_{yd} if the sled is moving outward. If it is moving inward, we use Equation (10).

6 The Performance-Efficient Policy

The performance-efficient policy allows the sled to reach the center in the shortest time possible. Therefore, we use the actuators for acceleration and deceleration, consuming external energy in both phases. When deploying the performance-efficient policy, the sled acts as if it seeks from the current position to the center position. An exception is when shutting down the sled has to stop at the center not only on X, but also on Y. Hong et al. [2] devise an analytical bang-bang model of the seek time for MEMS-based storage devices. The model of the motion along X holds completely for the shutdown time along X; the acceleration time is:

$$t_{xa} = \sqrt{\frac{m_x}{k_x}} \arccos\left(\frac{x_s - \frac{m_x a_x}{k_x}}{x_0 - \frac{m_x a_x}{k_x}}\right) \tag{14}$$

The deceleration time is:

$$t_{xd} = \sqrt{\frac{m_x}{k_x}} \arccos\left(\frac{x_s + \frac{m_x a_x}{k_x}}{x_1 + \frac{m_x a_x}{k_x}}\right), \tag{15}$$

Unlike the X seek model, the Y seek model demands modification to account for the stop on Y at the center. The acceleration time along Y becomes:

$$t_{ya} = \sqrt{\frac{m_y}{k_y}} \arcsin(\frac{y_s - \frac{m_y a_y}{k_y}}{\sqrt{(y_0 - \frac{m_y a_y}{k_y})^2 + (v_a\sqrt{\frac{m_y}{k_y}})^2}}) - \sqrt{\frac{m_y}{k_y}} \arcsin(\frac{y_0 - \frac{m_y a_y}{k_y}}{\sqrt{(y_0 - \frac{m_y a_y}{k_y})^2 + (v_a\sqrt{\frac{m_y}{k_y}})^2}}), \quad (16)$$

and the deceleration time is:

$$t_{yd} = \sqrt{\frac{m_y}{k_y}} \arccos(\frac{y_s + \frac{m_y a_y}{k_y}}{y_1 + \frac{m_y a_y}{k_y}}) . \quad (17)$$

Every x and y is multiplied by -1, since $x_0 > x_1$ and $y_0 > y_1$, respectively.

7 Implementation in DiskSim

The MEMS model of DiskSim assumes an instantaneous shutdown time of the media sled in MEMS devices. As a result, it assumes that the shutdown operation consumes no energy. Since MEMS devices shut down very often, an accurate modeling of the shutdown time and energy becomes necessary.

We, therefore, implement the two shutdown policies presented in this report in the MEMS model of DiskSim. Simulation can be run with either policy by adjusting the shutdown policy parameter. For better modeling of a real MEMS device, we keep track of the position of the sled during shutdown. Doing so, a new request can interrupt shutdown, if it arrives during shutdown.

8 Summary

We devise an analytical model of an energy-efficient policy to shut down the media sled in MEMS-based storage devices. This policy exploits the spring structure of these devices to do part of the shutdown without active actuation. Thus, it reduces the shutdown energy.

As a result of the reduction in the shutdown energy, shutdown decisions can be made more aggressively than what would be possible with other policies that do not exploit the spring structure. Increasing the aggressiveness yields further savings on the active energy during periods of inactivity, since the sled is stopped earlier, increasing the energy efficiency of MEMS-based storage devices.

9 Acknowledgements

This research is supported by the Technology Foundation STW, applied science division of NWO and the technology programme of the Ministry of Economic Affairs under project number TES.06369.

References

[1] Mark A. Lantz, Hugo E. Rothuizen, Ute Drechsler, Walter Häberle, and Michel Despont. A Vibration Resistant Nonopositioner for Mobile Parallel-Probe Storage Applications. *Journal of Microelectromechanical Systems*, 16(1):130–139, February 2007.

[2] Bo Hong and Scott A. Brandt. An analytical solution to a MEMS seek time model. Technical Report UCSC-CRL-02-31, Storage Systems Research Center, University of California, Santa Cruz, September 2002.

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

[4] L. Abelmann, T. Bolhuis, A. M. Hoexum, G. J. M. Krijnen, and J. C. Lodder. Large capacity probe recording using storage robots. *IEE Proceedings: Science, Measurement and Technology*, 150(5):218–221, 2003.

[5] Nanochip Inc. `http://nanochipinc.com/tech.htm`. Accessed in November 2007.

[6] Tara Madhyastha and Katherine Pu Yang. Physical modeling of probe-based storage. In *Proceedings of the 18th IEEE Symposium on Mass Storage Systems and Technologies*, pages 207–224, April 2001.