

A RIGOROUS ANALYSIS FOR SET-UP TIME MODELS - A METRIC PERSPECTIVE

Eitan Bachmat

Tao Kai Lam

Avner Magen

Dept. of Computer Science
Ben Gurion University

EMC cooperation
Hopkinton, MA

Dept. of Computer Science
University of Toronto

Abstract. We consider model based estimates for set-up time. The general setting we are interested in is the following: given a disk and a sequence of read/write requests to certain locations, we would like to know the total time of transitions (set-up time) when these requests are served in an orderly fashion. The problem becomes nontrivial when we have, as is typically the case, only the counts of requests to each location rather than the whole input, in which case we can only hope to estimate the required time. Models that estimate the set-up time have been suggested and heavily used as far back as the sixties. However, not much theory exists to enable a qualitative understanding of such models. To this end we introduce several properties through which we can study different models such as (i) *super-additivity* which means that the set-up time estimate decreases as the input data is refined (ii) *monotonicity* which means that more activity produces more set-up time, and (iii) an approximation guarantee for the estimate with respect to the worst possible time.

We provide criteria for super-additivity and monotonicity to hold for popular models such as the independent reference model (IRM). The criteria show that the estimate produced by these models will be monotone for any reasonable system. We also show that the IRM based estimate functions, upto a factor of 2, as a worst case estimate to the actual set-up time.

To establish our theoretical results we use the theory of finite metric spaces, and en route show a result of independent interest in that theory, which is a strengthening of a theorem of Kelly [4] about the properties of metrics that are formed by concave functions on the line.

1 Introduction

Set-up times which are associated with moving a system from one state to another play a major role in the performance analysis of systems. Perhaps the most glaring example is provided by disk based storage systems in which the states correspond to locations on the disk. In this case the total duration of the movements of the disk's head (from one location to another or from one disk track to another), aka the *set-up time* is the dominant feature in the total service time, and hence a lot of effort is put in order to minimize it by means of reordering the disk's content. Interestingly enough, in this application as well as in other real world applications, the above task becomes a problem with *partial input*.

The reason is simple: to collect all transition information will be too costly and will render the original optimization useless as the set-up time will be second to the input collection time. Instead, the only information typically available is the state counts, ie the number of times that each state was requested. In graph terminology we want to know the length of a path in a weighted graph where we only know the number of times that each node was visited.

In order to estimate the set-up time, researchers have used stochastic *models*, ie stochastic processes with parameters that are inherited from the observed count. The simplest of these models, the *Independent Reference Model* (IRM) is very intuitive: the requests at any time are drawn (independently of the previous state) from a distribution proportional to the count vector. This simple model is the most popular model for the analysis of storage system performance; see for example [1, 3, 6–9] among many.

In this paper we consider new and basic properties of set-up time estimates and check whether they hold for the IRM model. In a full version of the paper we will consider other models such as the so called the *Partial Markov Model* (PMM). These properties relate the set-up time estimates to the worst case case and examine the changes in the estimate due to a different way of collecting the data. The applicability of these properties to various models is an evidence to their quality, and moreover they allow for a rigorous study of models that are heavily used, often with not enough underlying rationale. It is interesting to note that while the IRM is one of the oldest models of user access patterns, dating back to the sixties, the basic properties considered above have never been explored. What follows is a brief description of these properties.

Given time intervals $I \subset J$ it is obvious that a system suffers at least as much set-up time during J as it does during I . The *monotonicity* property simply says that the set-up time estimate of the model reflects that fact, ie it gives an estimate for J which is at least as big as the one for I . A model is said to be *super-additive* if the addition of input information (by means of higher resolution of measurements) does not increase the set-up time estimate. It is almost immediate that super-additivity implies monotonicity and that it applies to the worst case time which provides the largest possible set-up time consistent with a given input data. The last property compares the set-up time estimate with the worst case estimate (which is NP hard to compute). Showing that the estimate of a model does not deviate much from the worst case estimate is tantamount to showing that is not over optimistic.

Our results: We show that monotonicity applies to the IRM , *regardless of the metric involved*. We further show that IRM set-up time estimate is a $1/2$ approximation to the worst case. Our results concerning super-additivity have the following curious feature: Super additivity holds in the IRM model *provided* that the “time-metric”, ie the times associated with the transition times between pair of states, belongs to the well studied class of metric spaces known as *negative type metrics*. Not all metric spaces belong to this class, but as we show, the time metrics that come from motion of disk drives in fact do, owing to the physical features of the system. Therefore IRM is indeed super additive with

respect to these I/O systems. These results show that the IRM can be used to produce reliably conservative estimates which are easy to calculate and that easily lend themselves to compactness-of-input/accuracy tradeoff. Following these observations the first and second authors used the IRM set-up time estimate as a central ingredient in a commercially available application which dynamically reconfigures data in a disk array. Details of the application and successful results from real production environments are to be presented elsewhere.

Techniques: Naturally, much of the notions and proofs come from and use the theory of metric spaces. The classes of interest in this discussion are ℓ_1 -metrics and negative type metrics, as well as the general class of metrics. In the process of establishing our results we extend a result of Kelly on the properties of invariant metrics on the real line coming from concave functions.

Organization: The rest of the paper is organized as follows. Section 2 Introduces set-up times and discusses some basic definitions and facts from the theory of metric spaces relevant to our discussion. Section 3 describes the basic models which we will study and introduces the concepts of monotonicity, super additivity, dominance and approximation. In section 4 we prove criteria for monotonicity and super additivity of the IRM estimate in terms of metric properties of the set-up time function. Finally, Section 5 discusses properties of metric arising from the seek times in disk drives.

2 Preliminaries

2.1 Set-up time

Throughout the paper we let X represent the states of a system. In this section we let n denote the number of states in X . Following [1] section 6.2, we let the function $d : X \times X \rightarrow \mathbf{R}^+$, be the set-up time function; namely, for $i, j \in X$, $d(x_i, x_j)$ represents the amount of time which is required to switch the system from state i to state j .

The abstract notion of a state can acquire many different meanings in different applications. For example, the states can refer to different tasks that the system needs to accomplish as in production systems and processors, or, to physical locations where tasks should be conducted as in storage systems. We assume that there is some process which generates a sequence of requests for the states of X .

Given a time interval I let $\mathbf{x}^I = \mathbf{x} = x_1, \dots, x_m$ be the sequence of requests for states of X during I . The *Total set-up time* during time interval I is simply the sum of the set-up times between consecutive requests

$$T(\mathbf{x}) = \sum_{j=1}^{m-1} d(x_j, x_{j+1})$$

In some cases we are not given the sequence of requests (a trace) but rather some partial information about the sequence \mathbf{x} . We wish to estimate the total set-up time of the sequence using the information available to us. In this paper we shall assume that the partial information available to us is the activity vector $\mathbf{a} = \mathbf{a}_I = (a_1, \dots, a_n)$, where a_i is the number of requests for state i during time interval I . We will assume that in general \mathbf{a} can be any vector with integer nonnegative entries. We let $a = \sum_i a_i$ be the total number of requests.

2.2 Metric Spaces

The theory of finite metric spaces will be used in the statements and proofs of our results. The following section provides some basic definitions and facts about metric spaces which will be needed later on.

We continue with a few standard definitions. A pair (X, d) where X is a set and d is a function $d : X \times X \rightarrow \mathbf{R}^+$ is called a *metric-space* if (i) $d(x, x) = 0$ for all $x \in X$ and $d(x, y) > 0$ for $x \neq y$, (ii) $d(x, y) = d(y, x)$ for all $x, y \in X$ and (iii) $d(x, y) + d(y, z) \geq d(x, z)$ for all $x, y, z \in X$. If instead of property (i) we only require that $d(x, x) = 0$ we say that (X, d) form a *semi-metric*. If we do not require the symmetry property, we say that (X, d) form a *Pseudometric*. One can “symmetrize” such an object by taking $d^*(x, y) = (d(x, y) + d(y, x))/2$. It can be easily seen that d^* satisfies 1’ and 3’ if d does. Set-up time functions can be reasonably assumed to satisfy the triangle inequality since one way to switch from state x to state z is to first switch from x to y and then from y to z . Set-up time functions cannot always be assumed to be symmetric as can be seen from rotational latency in disk drives.

Certain metric spaces are induced by norms. The ℓ_p norm on \mathbf{R}^n is $\|\mathbf{x}\|_p = (\sum_{i=1}^n |x_i|^p)^{\frac{1}{p}}$ where $\mathbf{x} = (x_1, \dots, x_n)$. A metric space (X, d) is called an ℓ_p -metric if there exists a mapping $\phi : X \rightarrow \mathbf{R}^n$ such that $d(x, y) = \|\phi(x) - \phi(y)\|_p$ for all $x, y \in X$. We sometimes say *Euclidean metric* instead of ℓ_2 -metric. A space (X, d) is *negative type* if (X, \sqrt{d}) is Euclidean.

Some basic facts about metric spaces Assume (X, d) is a finite metric space, $X = \{x_1, \dots, x_n\}$. There are two classical criteria for it to be Euclidean.

- Schoenberg’s criterion: (X, d) is Euclidean if and only if for all n real numbers v_1, \dots, v_n with $\sum_i v_i = 0$ we have $\sum_{i,j} v_i v_j d^2(x_i, x_j) \leq 0$. (This criterion is the reason for the name *negative type*, as by definition, d is Euclidean iff d^2 is negative type.)
- Cayley’s criterion: Consider the order $n - 1$ matrix M with entries $M_{i,j} = d^2(x_i, x_n) + d^2(x_j, x_n) - d^2(x_i, x_j)$, $i, j = 1, \dots, n - 1$. Then (X, d) is Euclidean if and only if the matrix M is positive semi definite, ie, all of its eigenvalues are nonnegative.

We say that a metric (X, d) is L_1 if there exist functions f_x , $x \in X$ such that $d(x, y) = \int_{\mathbf{R}} |f_x(t) - f_y(t)| dt$. It is known that a finite metric space is L_1 iff it is

ℓ_1 . Another well known fact we later use is that every ℓ_1 -metric is negative type [5]. Negative type distances do not necessarily satisfy the triangle inequality.

A distance function can be defined on the line given a real positive function F with certain properties. We define the distance d_F between i and j as $d_F(i, j) = F(|i - j|)$. utilized in this paper. We note that if F is convex then d_F satisfies the triangle inequality and thus provides a metric.

3 Models and their properties

Recall that our input is an activity vector, that is the count of requests to the different states; however, in order to know the total set-up time we need to know the actual sequence of requests. In the absence of the actual sequence we use models for estimating set-up time. A model for estimating set-up time is an interpretation of an activity vector as a distribution over sequences, and the resulting estimate is then the expected set-up time for a random sequence drawn from this distribution. For example some models will interpret an activity vector $(100, 100)$ as a uniform distribution of sequences that visit either location 1 or 2, while other will consider the distribution in which either all first 100 requests are for the first location or all of them were for the other; clearly the two different models in the above example will produce very different time estimates.

3.1 Examples of models and estimates

We now describe a few models M and their associated set-up time estimates.

The IRM (independent reference model) The IRM models independent random requests to states in X , taking into account that the different states are not uniformly popular. The model is parameterized by a probability distribution $p = p_i$ on the set of states X . The model itself is then given by the product measure on X^a . The product measure reflects an underlying assumption that requests are generated independently of each other. To be compatible with the observed activity vector we set the request probability for state i to be $p_i = a_i/a$ and the length of the generated sequence to be a . For this model the expected total set-up time is

$$T(\mathbf{a}, d; IRM) = a \sum_{i,j} p_i p_j d(x_i, x_j) = \frac{1}{a} \sum_{i,j} a_i a_j d(x_i, x_j)$$

We will refer to $T(\mathbf{a}, d; IRM)$ as the IRM estimate. For ease of notation we will sometimes use $T(\mathbf{a}, d)$ instead of $T(\mathbf{a}, d; IRM)$.

The next model is not discussed in details in this extended abstract, and our results about it will be presented in the full version of the paper.

The PMM (Partial Markov models) r_i of not moving to another state, and in the event of a move, the next state is j with probability q_j , independent of the current requested state. Consequently, the transition probabilities of moving from i to j are $p_{i,j} = (1 - r_i)q_j$ for $i \neq j$ and $p_{i,i} = r_i + (1 - r_i)q_i$. Here $0 \leq r_i, q_i \leq 1$. We call the vector $\mathbf{r} = (r_i)$ the *locality* vector of the model. Given a locality vector \mathbf{r} and an observed activity vector \mathbf{a} for some time interval I there exists a unique partial Markov model P which is compatible with \mathbf{r} and \mathbf{a} . By compatibility we mean that \mathbf{r} is the locality vector of P and \mathbf{a}/a is the stationary distribution of P which expresses the expected reference probabilities in the model P . Fix the vector $\mathbf{r} = (r_i)$. We let $P^{\mathbf{r}}$ denote the partial Markov model which for each interval I uses the model P compatible with \mathbf{r} and \mathbf{a}_I to model the request stream during I (note that $P^{\mathbf{0}}$ is simply the IRM). The $P^{\mathbf{r}}$ estimate is

$$T(\mathbf{a}, d; P^{\mathbf{r}}) = a \left(\sum_{i,j} (a_i/a) P_{i,j}^{\mathbf{r}} d(x_i, x_j) \right)$$

Partial Markov models are useful in capturing locality of reference phenomenon, [1, 2], which means that a request to state i is likely to be followed by another request to state i within a short time span. Many applications naturally exhibit this type of behavior. The larger the entries of the locality vector \mathbf{r} , the more likely states are to repeat in succession. In the partial Markov model the number of repetitive successions is distributed geometrically.

The worst case (supremum) model In the worst case model W we assume that the sequence of states during time interval I was the sequence which maximizes the total set-up time among all sequences which are consistent with the vector \mathbf{a} . The measure is thus a δ measure on the worst case sequence. Consequently,

$$T(\mathbf{a}, d; W) = \max \sum_{i=1}^a d(x_i, x_{i+1})$$

where the maximum is over all sequences of states in X , of length a that agree with the frequency vector \mathbf{a} and $x_1 = x_{a+1}$. We refer to $T(\mathbf{a}, d; W)$ as the worst case estimate.

3.2 Properties of models

We introduce notions which will allow us to examine the behavior of model based estimates with regards to changes in the input data and to compare estimates for different models.

super additivity Let I be a time interval and let I_1, \dots, I_k be a subdivision of I into subintervals. Accordingly, we have $\mathbf{a}_I = \sum_{j=1}^k \mathbf{a}_{I_j}$. A model M is said to be *super additive* with respect to a set-up time function d if the inequality

$$T(\mathbf{a}_I, d; M) \geq \sum_{j=1}^k T(\mathbf{a}_{I_j}, d; M) \tag{1}$$

always holds. Super additivity may be interpreted as stating that the addition of input information, namely, \mathbf{a}_I , instead of \mathbf{a} , never increases the estimate.

monotonicity We say that a vector $\mathbf{a} = (a_i)$ dominates a vector $\mathbf{b} = (b_i)$ if for all i , $a_i \geq b_i$. We use the notation $\mathbf{a} \geq \mathbf{b}$ to denote dominance. A model M is said to be *monotone* with respect to d if for any pair of time intervals $I \subset J$ we have $T(\mathbf{a}_I, d; M) \leq T(\mathbf{a}_J, d; M)$, or stated otherwise, for any pair of vectors \mathbf{a}, \mathbf{b} with nonnegative entries and such that $\mathbf{a} \geq \mathbf{b}$ we have

$$T(\mathbf{a}, d; M) \geq T(\mathbf{b}, d; M) \quad (2)$$

approximation Let $0 < \alpha < 1$. Given a set up function d , a model M_1 is said to provide an α *approximation* to a model M_2 (and vice versa) if for any activity vector \mathbf{a} we have

$$\alpha \leq \frac{T(\mathbf{a}, d; M_1)}{T(\mathbf{a}, d; M_2)} \leq \frac{1}{\alpha} \quad (3)$$

We say that a model M is *conservative* if it α approximates the worst case model W for some $\alpha > 0$.

4 Metric space criteria for properties of models

In this section we establish criteria for monotonicity and super additivity of the IRM estimates in terms of metric properties of the set-up time function d . We also establish a criterion for the IRM estimate to be a 1/2 approximation to the worst case estimate.

Theorem 1 (*A criterion for Super additivity*) *The IRM estimate is super additive with respect to d if and only if d is negative type.*

Proof. It is enough to establish super additivity for a subdivision of I into two subintervals, that is to show that for all nonnegative vectors $\mathbf{a} = (a_i), \mathbf{b} = (b_i)$,

$$T(\mathbf{a} + \mathbf{b}, d) \geq T(\mathbf{a}, d) + T(\mathbf{b}, d) \quad (4)$$

Let $a = \sum_i a_i$ and $b = \sum_i b_i$. Then

$$\begin{aligned} & T(\mathbf{a} + \mathbf{b}, d) - T(\mathbf{a}, d) - T(\mathbf{b}, d) \\ &= \sum_{i \neq j} \frac{(a_i + b_i)(a_j + b_j)d(x_i, x_j)}{a + b} - \sum_{i \neq j} \frac{a_i a_j d(x_i, x_j)}{a} - \sum_{i \neq j} \frac{b_i b_j d(x_i, x_j)}{b} \\ &= \frac{1}{ab(a + b)} \sum_{i \neq j} d(x_i, x_j) (a_i b_j ab + a_j b_i ab - a_i a_j b^2 - b_i b_j a^2) \\ &= \frac{1}{ab(a + b)} \sum_{i \neq j} d(x_i, x_j) (a_i b - b_i a) (b_j a - a_j b) \\ &= -\frac{ab}{a + b} \sum_{i \neq j} d(x_i, x_j) \left(\frac{a_i}{a} - \frac{b_i}{b} \right) \left(\frac{a_j}{a} - \frac{b_j}{b} \right). \end{aligned}$$

Setting $v_i = \frac{a_i}{a} - \frac{b_i}{b}$, we get

$$T(\mathbf{a} + \mathbf{b}, d) - T(\mathbf{a}, d) - T(\mathbf{b}, d) = -\frac{ab}{a+b} \sum_{i \neq j} v_i v_j d(x_i, x_j). \quad (5)$$

We note that $\sum_i v_i = 0$, hence by Schoenberg's criterion the IRM estimate is super additive if d is negative type. Conversely if the IRM estimate is super additive then

$$\sum_{i,j} v_i v_j d(x_i, x_j) \leq 0$$

for all \mathbf{v} of the form $\mathbf{a}/a - \mathbf{b}/b$ where \mathbf{a}, \mathbf{b} are vectors with integer non negative entries. After scaling we may deduce that the property holds whenever \mathbf{a}, \mathbf{b} have rational non negative entries and by density of the rationals for all \mathbf{a}, \mathbf{b} with non negative entries. Every vector $\mathbf{v} = (v_1, \dots, v_h)$ such that $\sum_i v_i = 0$ has a multiple of the form $\frac{1}{a}\mathbf{a} - \frac{1}{b}\mathbf{b}$, where \mathbf{a}, \mathbf{b} have non negative entries. Indeed if $a_i = \max\{v_i, 0\}$ and $b_i = \max\{-v_i, 0\}$, then $a = b$ and $\frac{1}{a}\mathbf{a} - \frac{1}{b}\mathbf{b} = \frac{1}{a}\mathbf{v}$, hence Schoenberg's criterion holds and d is negative type.

Theorem 2 (*criteria for monotonicity*) *The IRM estimate is monotone with respect to d if and only for every choice of k , the matrix $B(k, d)_{i,j} = d(x_i, x_k) + d(x_k, x_j) - d(x_i, x_j)$ defines a nonnegative quadratic form when restricted to vectors with nonnegative entries. In particular, if d is a pseudo metric or negative type then the IRM estimate is monotone with respect to d .*

Proof. We check the sign of the partial derivatives of $T(\mathbf{a}, d)$ with respect to a_k (where $k \in \{1, \dots, n\}$ is an arbitrary element).

$$\begin{aligned} & \frac{\partial}{\partial a_k} T(\mathbf{a}, d) \\ &= \frac{a(\sum_i a_i d(x_i, x_k) + \sum_j a_j d(x_j, x_k)) - \sum_{i,j} a_i a_j d(x_i, x_j)}{a^2} \\ &= \frac{1}{a^2} \sum_{i,j} a_i a_j (d(x_i, x_k) + d(x_j, x_k) - d(x_i, x_j)) = \frac{1}{a^2} \mathbf{a} B \mathbf{a}^t \end{aligned}$$

where $B = B(k, d)$ is the matrix with ij entry $d(x_i, x_k) + d(x_j, x_k) - d(x_i, x_j)$. Assume that for all k , $B(k, d)$ is positive semi definite on vectors with nonnegative entries then $\frac{\partial}{\partial a_k} T(\mathbf{a}, d) \geq 0$ for all k and all activity vectors \mathbf{a} . It follows from the Mean-value Theorem that if $\mathbf{a} \geq \mathbf{b}$ then $T(\mathbf{a}, d) \geq T(\mathbf{b}, d)$. Conversely if there are $\mathbf{a} \geq 0$ and k such that $\mathbf{a} B(k, d) \mathbf{a}^t < 0$ then taking \mathbf{b} which is identical to \mathbf{a} except that b_k is slightly smaller than a_k we get $T(\mathbf{a}, d) < T(\mathbf{b}, d)$, which proves the first statement of part 3.

If d is a semi-metric then B has nonnegative entries and so $\mathbf{a} B(k, d) \mathbf{a}^t \geq 0$ and if d is negative type then by Cayley's criterion $\mathbf{a} B(k, d) \mathbf{a}^t \geq 0$ which completes the proof.

Theorem 3 (*Comparison of the IRM estimate and worst case estimate*) *If d satisfies the triangle inequality then for all activity vectors \mathbf{a} we have*

$$2T(\mathbf{a}, d; IRM) \geq T(\mathbf{a}, d; W) \quad (6)$$

where W is the worst case model.

Proof. Assume first that the activity vector is the vector $(1, 1, \dots, 1)$. The IRM estimate here is $\frac{1}{n} \sum_{i,j} d(x_i, x_j)$, while the worst case estimate is the length of the longest Hamiltonian cycle in the complete graph on X with edge weights given by d . Assume without loss of generality that the longest Hamiltonian path in X is $1, 2, \dots, n$. Since d satisfies the triangle inequality we have for $1 \leq i < n$ and for $j \in X$ $d(x_i, x_{i+1}) \leq d(x_i, x_j) + d(x_j, x_{i+1})$ (the $n+1$ point coincides with the first point). Summing over all i, j we get

$$n \sum_{i=1}^n d(x_i, x_{i+1}) \leq 2 \sum_{i,j} d(x_i, x_j).$$

Therefore $2T(\mathbf{a}, d; IRM) \geq T(\mathbf{a}, d; W)$. To complete the proof we need to consider a general activity vector (a_1, \dots, a_n) . Let X' be the metric space with a points that is composed of groups of a_j points of type j . Given d on X we induce a metric on X' by letting the distance between a point of type i and a point of type j be $d(x_i, x_j)$. Clearly X' also satisfies the triangle inequality. We have thus reduced the problem to the case of the activity vector $(1, 1, \dots, 1)$ and are done.

5 Set-up Time Functions of a Disk

In this section we show that the radial seek time function of a disk drive, which is the standard set-up function in storage system research is an ℓ_1 -metric and in particular is negative type. From this we conclude that the IRM estimates are super additive when applied to disk seek times. Data on disk drives resides on *tracks* which form concentric circles of varying radii r around the center of a platter. To get from a track at radius r_1 to another track at radius r_2 the head of the device performs a radial motion. The time it takes the disk head to perform this radial motion is known as (radial) seek time. Seek time is translation invariant Furthermore, the acceleration and deceleration of the head dictate that the seek time from r_1 to r_2 has the form

$$d_F(r_1, r_2) = F(|r_1 - r_2|)$$

where F is a concave non decreasing function.

If we let X be the set of data locations on the disk then a theorem of Kelly proved in [4] can be interpreted as stating that (X, d_F) is negative type. We prove a stronger result of independent interest using a much simpler proof.

Theorem 4 Let F be a concave nondecreasing function with $F(0) = 0$ and let $X \subset \mathbf{R}$. Then (X, d_F) is an ℓ_1 metric space.

Proof. Let $X = \{x_1, \dots, x_n\}$. Consider

$$Y = \{|x_i - x_j| : 1 \leq i, j \leq n\}$$

the set of possible distances in X , and order the elements of Y as $0 = y_0 < y_1 < y_2 < \dots < y_m$. let G be the piecewise linear function which

- (i) coincides with F on Y
- (ii) is linear on all intervals $[y_i, y_{i+1}]$ and
- (iii) is constant on $[y_m, \infty)$ (that is, gets the value $F(y_m)$ there).

Obviously $(X, d_F) = (X, d_G)$ since $F = G$ on the set of all relevant values Y , so it is enough to prove the claim for G , which is also non decreasing and concave. We now define functions $H_{s,t}$ as follows.

$$H_{s,t}(x) = sx \text{ if } x < t \text{ and } st \text{ otherwise.}$$

We also let $s_i = \frac{G(y_i) - G(y_{i-1})}{y_i - y_{i-1}}$ be the sequence of slopes of G . We now claim that G is a convex combination of functions of the form $H_{s,t}$.

The proof proceeds by induction on m . If $m = 0$ then $G = H_{1,0} = 0$. For $m > 0$, look at the function $\tilde{G} = G - H_{s_m, y_m}$. It is not hard to see that $\tilde{G}(0) = 0$, \tilde{G} is constant beyond y_{m-1} and is piecewise linear with breakpoints y_1, \dots, y_{m-1} . A piecewise linear function is concave and nondecreasing if and only if its slopes are decreasing and nonnegative, and so $s_1 \geq s_2 \geq \dots \geq s_m \geq 0$, and similarly $s_1 - s_m \geq s_2 - s_m \geq \dots \geq s_{m-1} - s_m \geq 0$. But, these are the slopes of \tilde{G} and it is therefore concave and nondecreasing. We may now apply the induction hypothesis to \tilde{G} and this proves the claim.

Since a sum of ℓ_1 -metrics is also an ℓ_1 -metric, we are left with the task of showing that for a function $F = H_{s,y}$, the resulting metric d_F is an ℓ_1 -metric. Notice that $d_F(i, j) = s \cdot \min\{|i - j|, y\}$. Let $f_i = \frac{1}{2s} \chi_{[x_i, x_i + y]}$ be the function whose value is $\frac{1}{2s}$ on the interval $[x_i, x_i + y]$ and zero otherwise. It is easy to see that for any $i, j \in \mathbf{R}$

$$d_F(i, j) = s \cdot \min\{|i - j|, y\} = \int_{\mathbf{R}} |f_i(x) - f_j(x)| dx$$

This shows that d_F is an L_1 metric and hence l_1 .

Combining theorem 1 with theorem 4 we get

Theorem 5 The IRM estimate is super additive with respect to the seek time function d_F for any physical disk drive.

6 Conclusions and future work

We have introduced several natural properties of set-up time estimates and studied them for the IRM. We have shown that the IRM estimate satisfies monotonicity which is a “sanity check” for set-up time estimates, and further that the IRM is an easily computable approximation to the worst case estimate. In the specific but important context of seek functions in disk drives we showed that the IRM shares another formal property that holds for worst case estimates namely super additivity. It would be interesting to explore monotonicity, super additivity and various approximation relations among other models. One interesting class of examples are the renewal models which were suggested by Opderbeck and Chu in [6]. The IRM is a special case of such models where the renewal model is based on exponential inter-arrival times. It would be interesting to investigate other cases such as hyperexponential, gamma or Pareto bounded heavy tail distributions. Such an investigation will likely require refined definitions for properties such as monotonicity and super additivity since the associated models are not Markovian.

Acknowledgments: We would like to thank Timothy Chow for helpful discussions regarding a preliminary version of this paper.

References

1. Aven O.I., Coffman E.G. and Kogan Y.A. *Stochastic analysis of computer storage*, D.Reidel publishing, 1987.
2. Coffman E.G. and Denning P.J. *Operating systems theory*, Prentice Hall, 1973.
3. Grossman D.D. and Silverman H.F. Placement of records on a secondary storage device to minimize access time *J. of the ACM* 20, 429-438, 1973.
4. Kelly J.B. Hypermetric spaces and metric transforms, in *Inequalities III*, edited by O.Shisha, 149-158, 1972.
5. Lovasz L., Pyber L., Welsh D.J.A. and Ziegler G.M. Combinatorics in pure mathematics, in *Handbook of combinatorics*, MIT press, North Holland 2039-2082, 1995.
6. Opderbeck H. and Chu W.W., The renewal model for program behavior, *SIAM J. of computing* 4, 356-374, 1975.
7. Vanichpun S. and Makowski A., The output of a cache under the independent reference model - where did all the locality of reference go?, *Proceedings of SIGMETRICS*, 2004.
8. Wong C. K. *Algorithmic Studies in mass storage systems*, computer science press, 1983.
9. Yue P.C. and Wong C.K., On the optimality of the probability ranking scheme in storage applications, *J. of the ACM* 20, 624-633, 1973.