

Competitive Algorithms for Block-Aware Caching

Christian Coester*

University of Sheffield

Sheffield, UK

christian.coester@gmail.com

Roie Levin[†]

Computer Science Department, Carnegie Mellon
University

Pittsburgh, PA, USA

roiel@cs.cmu.edu

Joseph (Seffi) Naor[†]

Computer Science Department, Technion

Haifa, Israel

naor@cs.technion.ac.il

Ohad Talmon[†]

Computer Science Department, Technion

Haifa, Israel

ohad@cs.technion.ac.il

ABSTRACT

Motivated by the design of real system storage hierarchies, we study the *block-aware caching problem*, a generalization of classic caching in which fetching (or evicting) pages from the same block incurs the same cost as fetching (or evicting) just one page from the block. Given a cache of size k , and a sequence of requests from n pages partitioned into given blocks of size $\beta \leq k$, the goal is to minimize the total cost of fetching to (or evicting from) cache. This problem captures generalized caching as a special case, which is already NP-hard offline. We show the following suite of results:

- For the eviction cost model, we show an $O(\log k)$ -approximate offline algorithm, a k -competitive deterministic online algorithm, and an $O(\log^2 k)$ -competitive randomized online algorithm.
- For the fetching cost model, we show an integrality gap of $\Omega(\beta)$ for the natural LP relaxation of the problem, and an $\Omega(\beta + \log k)$ lower bound for randomized online algorithms. The strategy of ignoring the block-structure and running a classical paging algorithm trivially achieves an $O(\beta)$ approximation and an $O(\beta \log k)$ competitive ratio respectively for the offline and online-randomized setting.

*Part of this research was done while the author was at Tel Aviv University, supported by the Israel Academy of Sciences and Humanities & Council for Higher Education Excellence Fellowship Program for International Postdoctoral Researchers.

[†]Supported in part by US-Israel BSF grant 2018352 and by ISF grant 2233/19 (2027511).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SPAA '22, July 11–14, 2022, Philadelphia, PA, USA.

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9146-7/22/07...\$15.00

<https://doi.org/10.1145/3490148.3538567>

- For both fetching and eviction models, we show improved bounds for the (h, k) -bicriteria version of the problem. In particular, when $k = 2h$, we match the performance of classical caching algorithms up to constant factors.

Our results establish a strong separation between the tractability of the fetching and eviction cost models, which is interesting since fetching/eviction costs are the same up to an additive term for the classic caching problem. Previous work of Beckmann et al. (SPAA 21) only studied online deterministic algorithms for the fetching cost model when $k > h$.

Our insight is to relax the block-aware caching problem to a submodular covering linear program. The main technical challenge is to maintain a competitive fractional solution to this LP, and to round it with bounded loss, as the constraints of this LP are revealed online. We hope that this framework is useful going forward for other problems that can be captured as submodular cover.

CCS CONCEPTS

• **Theory of computation** → **Caching and paging algorithms; Rounding techniques; Linear programming.**

KEYWORDS

Online Algorithms, Competitive Analysis, Caching

ACM Reference Format:

Christian Coester, Roie Levin, Joseph (Seffi) Naor[2], and Ohad Talmon[2]. 2022. Competitive Algorithms for Block-Aware Caching. In *Proceedings of the 34th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '22)*, July 11–14, 2022, Philadelphia, PA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3490148.3538567>

The full version of this paper [17] is available at <https://arxiv.org/abs/2205.12249>.

1 INTRODUCTION

Caching (also known as paging) has been extensively studied since the early days of online computation and competitive analysis, establishing itself as a cornerstone problem in this field, see e.g., [1, 3, 5, 6, 11, 12, 19, 25–27, 29, 30, 34, 36]. Recent years have witnessed increased activity on non-standard caching models, e.g., elastic caches [20], caching with time windows [21], caching with dynamic weights [18], caching with machine learning predictions [28], and writeback-aware caching [7, 8]. Many of the recent developments in competitive analysis, e.g., the online primal-dual method, projections, and mirror descent [13–15] are all rooted in online paging. We study here *block-aware caching*, a non-standard caching model studied recently, as well as its generalizations.

In the (classic) weighted paging problem there is a universe of n pages, a cache that can hold up to k pages, and each page is associated with a weight (fetch cost). At each time step a page is requested; if the requested page is already in cache then no cost is incurred, otherwise the page must be fetched into the cache, incurring a cost equal to its weight. The goal is to minimize the total cost incurred. This problem is well studied and understood, and we briefly mention the main results known for it.

Sleator and Tarjan [30], in their seminal paper on competitive analysis, showed that any deterministic algorithm is at least k -competitive, and that LRU (Least Recently Used) is precisely k -competitive for unweighted paging (i.e., all weights are equal). The k -competitive bound was later generalized to weighted paging as well [16, 35]. When randomization is allowed, Fiat et al [19] gave the elegant Randomized Marking algorithm for unweighted paging, which is $\Theta(\log k)$ -competitive against an oblivious adversary. For weighted paging, Bansal et al. [5] gave an $O(\log k)$ -competitive randomized algorithm using the online primal-dual framework [4, 15]. It uses a two-step approach. First, a deterministic competitive algorithm is designed for a fractional version of the problem. Then, a randomized online algorithm is obtained by *rounding* the deterministic fractional solution online.

Block-aware caching. Real storage systems operate by constructing a hierarchy of memory levels, starting from a very fast and small memory (e.g., an SRAM cache) to a very large and slow memory (e.g., flash or disk). The data items in each level are typically organized in *blocks*, and fetching (or evicting) data items from the same block incurs the same cost as fetching (or evicting) just a single item from the block. Using fetching costs models scenarios in which data is read-only; using eviction costs models scenarios in which data must be written to slow memory upon eviction, and the writing cost dominates the reading cost (see e.g. [7, 9]).

Thus, a natural question is how can one optimize cache performance by taking advantage of granularity changes across different storage hierarchy levels. This question was recently raised by Beckmann et al. [10], who defined the *block-aware caching problem*, generalizing the classic paging problem, as follows. Given a cache of size k , and a sequence of requests from n pages that are partitioned into given

blocks of size $\beta \leq k$, minimize the total cost of fetching (or evicting) from the cache so as to serve the requests¹.

Block-aware caching also arises in web and cloud settings, where data items can be aggregated into chunks (i.e., blocks) of data, such that accessing a whole chunk incurs the same cost as accessing just a single item. Consider a distributed cluster of servers, where a common cache of data items is maintained. One such example is the ZFS distributed file system that aggregates different devices into a single storage pool acting as an arbitrary data store. When accessing a server for a specific data item, the main cost paid (e.g., latency) is for accessing the server. The notion of a block of data in this setting corresponds to the largest chunk of data items that can be fetched from (or evicted to) a server, while maintaining that the cost of this operation is dominated by the cost of accessing the server. Web caching is another example of block-aware caching. Consider a content delivery network (CDN) that maintains a cache of data items and suppose the CDN connects to a website so as to access a data item (see, e.g., [24, 31]). Typically, TCP/IP provides a time window for connecting to the website and accessing the data item. Hence, it might be beneficial to fetch (or evict) many data items that belong to the website, and not just the particular data item that is currently accessed. Thus, the notion of a block of data in this setting corresponds to the maximum number of such data items that can be sent without increasing the travel time.

In the generalized caching problem [2, 6], pages are associated with both a size and a cost. At any point of time, the sum of the sizes of the pages in the cache cannot exceed the cache size. In the offline setting, generalized caching is known to be NP-hard, and in the online setting the known competitive factors for generalized caching [2, 6] match those of weighted caching. It is not hard to see that block-aware caching captures generalized caching as a special case. Replace a page p of size s by a block B of size s containing page p partitioned into unit size “slices”. The cost of accessing each slice is equal to the cost of p . Now, a request to page p is replaced by many requests to the slices in B . Thus, an optimal solution to the block-aware caching problem generated has to fetch the full block B into the cache.

Eviction and fetching costs. In classic paging, costs can be associated with either evicting or fetching pages. Clearly, for a given request sequence, optimal eviction and fetching costs of serving the requests can differ by at most an additive constant that only depends on the initial contents of the cache. However, this is not the case for block-aware caching, as optimal eviction and fetching costs can differ significantly, separating the two cost models. (We provide an example in Section 2.) As discussed above, the two cost models are practically motivated for block-aware caching, and we thus study both of them in this paper. We note that in the eviction cost model we are able to circumvent known lower bounds [10] that hold in the fetching cost model.

¹We note that Beckmann et al. [10] considered block-aware caching only in the fetching cost model.

1.1 Results and Techniques

Observe that if an algorithm is r -competitive for classical paging, then it is at most $\beta \cdot r$ -competitive for block-caching in both fetching/eviction cost models; the reason is simply that OPT can be simulated by a classical paging algorithm that performs any single batched fetch/eviction in at most β rounds. With this in mind, our goal in this work is to beat this trivial linear dependence on β .

Indeed, for the eviction cost model, we give the first set of algorithms avoiding a trivial multiplicative β overhead over their classical paging counterparts. We also give (h, k) -bicriteria² algorithms for both fetching and eviction cost models, which we in turn use to adapt the lower bound of [10] for the fetching cost model to randomized algorithms.

Eviction cost. We start in Section 3 with our main contributions: competitive algorithms for the eviction cost model. We show the following theorem.

Theorem 1.1. For the block-aware caching problem with eviction cost, there exist:

- a k -competitive deterministic online algorithm.
- an $O(\log^2 k)$ -competitive randomized (integral) online algorithm.
- an $O(\log k)$ -approximate randomized offline algorithm.

In fact, we study a more general version than the one introduced in [10] in which every block B may have a separate cost c_B . For this more general weighted setting we get competitive ratios of k , $O(\log k \log(k\Delta))$ and $O(\log(k\Delta))$ for the deterministic online, randomized online, and randomized offline settings respectively (where Δ is the aspect ratio, i.e., the maximum cost ratio between any two blocks).

A first main technical ingredient is a linear programming relaxation for block-caching. It is tempting to use a formulation with a variable x_p^t for each page p and time t indicating whether p is present in cache in step t . However, in this case the eviction cost becomes a complicated non-linear function of the x_p^t . Instead, we define variable ϕ_B^t for each block B and time step t indicating whether we evict B at t . This is reminiscent of the linear program for classical paging of [5] in which every variable represents whether a page is present in cache between two subsequent requests to a page, only that it may now be necessary to evict pages at any point between subsequent requests.

A naïve linear programming formulation has an integrality gap of β (see Appendix A.2): this is unsurprising since the naïve LP exhibits this gap even for the special case of generalized paging. To get around this, we express feasibility as the constraint that a particular sequence of monotone, submodular functions is maximized. We then make use of good (albeit exponential size) LP relaxations for these submodular set function constraints, which were discovered by Wolsey [33]. Our formulation may be viewed as a generalization of the strengthened LP relaxation due to [6] for generalized caching,

²In (h, k) paging, an online algorithm with cache size k competes against an offline cache of size h , where $k > h$.

which used the so-called knapsack cover (KC) inequalities. We first use the relaxation to give a k -competitive deterministic online algorithmic in Section 3.2.

Next, we develop an $O(\log k)$ -competitive fractional algorithm in Section 3.3, followed by an $O(\log \Delta k)$ -competitive online randomized rounding procedure in Section 3.4. Together these imply our $O(\log k \log \Delta k)$ -competitive randomized online algorithm. It is natural to try to adapt the continuous online primal-dual framework of [5, 6]; however, owing to the increased complexity of our LP, there are several technical roadblocks. For one, our formulation now has primal variables corresponding to the eviction of every block at every point in time, and a naïve adaptation of the continuous dynamics of [6] incurs loss that depends on the length of the request sequence. Nevertheless, we show how to carefully set the rate of increase of primal variables (with respect to the dual rate of increase) to construct a feasible solution with our claimed guarantee. For convenience, we do not present the fractional algorithm as online in the strict sense as we allow it to change decisions made in the past. However, it has the crucial property that it only increases LP variables, which suffices for the online rounding procedure.

For the rounding step, we forego maintaining an explicit distribution over cache states as in previous work [2, 5, 6], since it is unclear how to control the cost of the rebalancing stage when updating the distribution in each time step. Instead, we use the method of random rounding with alterations in a similar spirit to [7]. A key difference in our work is the added difficulty of working with the submodular cover formulation: this introduces additional challenges to the analysis of the rounding (as well as to the maintenance of the fractional solution). We make use of recent work on online submodular cover [23]; interestingly we are able to charge our alteration cost to the fractional *fetching* cost, even though this may be a factor β larger than the eviction cost.

Fetching cost. We turn in Section 4 to the fetching cost model, where we show strong lower bounds, implying that the integrality gap of $\Omega(\beta)$ of the natural LP formulation cannot be circumvented. We prove the following theorem for (h, k) block-aware caching³.

Theorem 1.2. When $k = O(h)$, no randomized online algorithm has competitive ratio better than $\Omega(\beta + \log k)$ for block aware caching with fetching costs.

Our main idea here is an online deterministic rounding procedure for fractional algorithms that incurs constant blowup in both cache usage and cost. This implies an online derandomization procedure for any randomized algorithm, which in turn strengthens the lower bounds for deterministic algorithms of [10] to apply to randomized algorithms as well. Our lower bound implies that beating the trivial linear dependence on β is *not* possible for the fetching cost model.

Our deterministic rounding procedure immediately implies improved bounds for the offline (h, k) block-aware caching problem. In particular, when $k = 2h$, we match the performance of classical caching algorithms up to constant factors.

³Beckmann et al. [10] showed several deterministic lower bounds on the competitive factor achievable for (h, k) block-aware caching, when $k \geq h + \beta - 1$.

2 MODEL AND PRELIMINARIES

2.1 Problem Definition

In the block-aware caching problem, there is a cache of size k and n pages which are partitioned into blocks. Let \mathcal{B} be the partition of the pages into blocks. Each block contains at most β pages, for some $\beta \in [k]$. For a block $B \in \mathcal{B}$, we denote by $c_B > 0$ its cost. At each time-step t , a page p_t is requested. To serve the request, the page p_t must be fetched into the cache if it is missing from the cache. The goal is to obtain a feasible cache policy while minimizing the total cost. We consider two different cost functions.

Eviction cost model. In this model, fetching into the cache is free, while evictions have a cost that can be aggregated: Evicting any subset A of a block B at a time-step has a cost of c_B . The goal is to minimize the total eviction cost.

Fetching cost model. In this model, evicting pages from the cache is free, while fetching of pages has a cost that can be aggregated. Fetching of any subset A of a block B at a time-step has a cost of c_B . The goal is to minimize the total fetching cost.

Unlike classic paging and its variants, the fetching cost and eviction cost models are not equivalent in block-aware caching. We show that the optimal fetching and eviction costs for the same request sequence may be off by a factor of β (in either direction!), and this bound is tight.

Claim 2.1. There exist instances of block-aware caching for which the optimal fetching cost is β larger than the optimal eviction cost, and there exist instances where the optimal eviction cost is β larger than the optimal fetching cost.

See Appendix A.1 for the proof.

For a page p , define $B(p)$ to be the block containing p , and let $r(p, t)$ be the time of the last request to p up until (and including) time t ; if there is no such request, then $r(p, t) := -\infty$. Define the aspect ratio $\Delta := c_{\max}/c_{\min}$ where $c_{\max} := \max_{B \in \mathcal{B}} c(B)$ and $c_{\min} := \min_{B \in \mathcal{B}} c(B)$. For convenience we will use the notation $[\ell] = \{1, \dots, \ell\}$ and $[\ell]_0 = \{0, 1, \dots, \ell\}$.

Our work relies on the theory of submodular functions which we introduce now for completeness.

Submodularity. We consider set functions of the form $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$, where \mathcal{N} is a set. For $A \subseteq B \subseteq \mathcal{N}$, let $f(A | B) := f(A \cup B) - f(B)$. For convenience, if $A = \{v\}$ is a singleton we also write $f(v | B) := f(\{v\} | B)$. We call f *submodular* if for all $v \in \mathcal{N}$, $A \subseteq B \subseteq \mathcal{N}$ we have $f(v | A) \geq f(v | B)$. A simple result is that if a set function f is submodular, then $f(\cdot | B)$ is also submodular for any $B \subseteq \mathcal{N}$. If for all $A \subseteq B \subseteq \mathcal{N}$ we have that $f(A) \leq f(B)$, then we say that f is *monotone*.

Submodular Cover. Wolsey [33] introduced the following problem known as *submodular cover*. Given a monotone, submodular function f over ground set \mathcal{N} , and cost function $c : \mathcal{N} \rightarrow \mathbb{R}^+$ on the ground set, output a minimum cost subset $S \subseteq \mathcal{N}$ such that

$f(S) \geq f(\mathcal{N})$. Wolsey gave the following LP relaxation:

$$\begin{array}{ll} \min & \sum_{v \in \mathcal{N}} c(v) \cdot x_v \\ \text{subject to} & \\ \forall S \subseteq \mathcal{N} : & \sum_{v \notin S} f(v | S) \cdot x_v \geq f(\mathcal{N}) - f(S) \\ \forall v \in \mathcal{N} : & x_v \geq 0 \end{array} \quad (2.1)$$

The constraints of this LP may be viewed as knapsack cover inequalities for a linearized version of the function f . Wolsey proved:

Claim 2.2 (Proposition 2 of [33]). A set S has $f(S) = f(\mathcal{N})$ if and only if χ_S , the characteristic vector of S , is a feasible integer solution to (2.1).

Furthermore, Wolsey showed that this LP has an integrality gap of $\log(\max_{v \in \mathcal{N}} f(v)) + 1$ when f is integer valued.

3 EVICTION COST

In this section we show our algorithmic results for block-aware caching with respect to eviction costs. Our proof uses the (online) primal-dual method and hence requires an LP for the eviction cost model. It is straightforward to write a simple LP relaxation; unfortunately, the naïve relaxation has an integrality gap of $\Omega(\beta)$ (see Appendix A.2), and recall that our goal is to beat the trivial algorithm's linear dependence on β .

3.1 Submodular Cover LP Formulation

To circumvent the naïve LP barrier, we strengthen the formulation using ideas from Wolsey's submodular cover LP. We start with some notation.

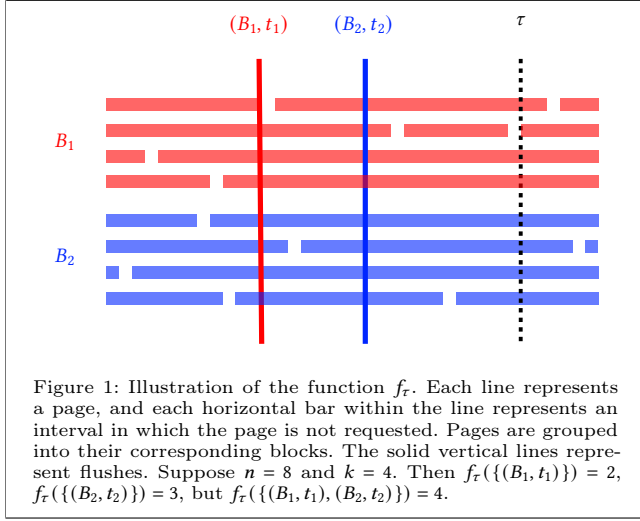
A *flush* is a tuple $(B, t) \in \mathcal{B} \times [T]_0$. The flush (B, t) corresponds to the event of evicting all cached pages of block B at time t . (There is no reason to only evict some of them, since they can be fetched back for free.) Let S be a set of flushes. We say that a page p is *missing* at time τ according to S if there exists $r(p, \tau) < t \leq \tau$ such that $(B(p), t) \in S$.⁴ Crucially, this definition ensures that the page p_t requested at time t is not missing at time t . We say that an algorithm is *induced* by a set of flushes S if the algorithm evicts all pages of block B (except p_t) at time t if and only if $(B, t) \in S$, and always loads p_t at time t . Let n_t be the number of pages requested up until time t .

We use the above to define a set function $f_\tau : 2^{\mathcal{B} \times [T]_0} \rightarrow \mathbb{Z}$ on sets of flushes:

$$f_\tau(S) := \min(n - k, |\{p : p \text{ is missing at time } \tau \text{ according to } S\}|)$$

In words, $f_\tau(S)$ is the number of pages that are outside of the cache at time τ for the algorithm induced by S , where this number is capped at $n - k$. The algorithm induced by a set of flushes S is feasible at time τ iff $f_\tau(S) \geq n - k$ for all τ .

⁴Adding to S all flushes of the form $(B, 0)$ ensures that also never-requested pages are missing by this definition.



We show the following simple fact in Appendix A.1:

Claim 3.1. For every τ , the function f_τ is submodular.

With the notation above, we can reformulate the block-aware caching problem with eviction cost as the solution to⁵

$$\begin{aligned} \min_{S \subseteq \mathcal{B} \times [T]_0} \quad & \sum_{\substack{(B,t) \in S \\ t \geq 1}} c_B \\ \text{subject to} \quad & \\ \forall \tau \in [T] : \quad & f_\tau(S) \geq n - k. \end{aligned} \quad (3.1)$$

Note that because $f_\tau(S)$ counts the number of pages evicted by S that are *not* p_τ , this single constraint captures both that the algorithm must flush at least $n - k$ pages in order to respect the cache size limit, and that the cache must contain page p_t at time t .

Note as well that we allow the algorithm to perform flushes at time 0, but only charge the cost for flushes performed after time 1. This conveniently allows the algorithm to clear the cache initially at no extra cost.

Finally, we are ready to write our LP, which is the intersection of the submodular cover LPs of (2.1) for the functions f_τ , across all

⁵This formulation is reminiscent of online and dynamic submodular cover problems [22, 23] in which the goal is also to maintain a feasible submodular cover while the underlying submodular function changes over time. However the cost models in these other works are different.

time steps $\tau \in [T]$.

$$\begin{aligned} \text{Primal} \\ \min \quad & \sum_{B,t \geq 1} c_B \cdot \phi_B^t \\ \text{subject to} \quad & \\ \forall S \subseteq \mathcal{B} \times [T]_0, \quad & \sum_{B,t} f_\tau((B,t) \mid S) \cdot \phi_B^t \\ \forall \tau \in [T] \quad & \geq n - k - f_\tau(S) \\ \forall B \in \mathcal{B}, t \in [T]_0 : \quad & \phi_B^t \geq 0 \end{aligned} \quad (P)$$

We will require the dual of this program, which is

$$\begin{aligned} \text{Dual} \\ \max \quad & \sum_{S,\tau} (n - k - f_\tau(S)) \cdot y_S^\tau \\ \text{subject to} \quad & \\ \forall B \in \mathcal{B}, t \in [T] : \quad & \sum_{S,\tau} f_\tau((B,t) \mid S) \cdot y_S^\tau \leq c_B \\ \forall S \subseteq \mathcal{B} \times [T], \quad & \forall \tau \in [T] : y_S^\tau \geq 0 \end{aligned} \quad (D)$$

That (P) is a valid relaxation of (3.1) follows from Claim 2.2.

Claim 3.2 (Corollary of Claim 2.2). A set S has $f_\tau(S) = n - k$ for all τ if and only if χ_S , the characteristic vector of S , is a feasible integer solution to (P).

We note for intuition's sake that even the constraints

$$\sum_{B,t} f_\tau((B,t)) \cdot \phi_B^t \geq n - k$$

alone already avoid the bad integrality gap example of Appendix A.2. One reason is that truncating f_τ at $n - k$ prevents the LP from overestimating how much space will be saved by evictions.

Given an LP solution ϕ , we also define the fractional value of a page p missing from cache at time t to be

$$x_p^t := \begin{cases} 1 & \text{if } r(p,t) = -\infty. \\ \min \left\{ 1, \sum_{u=r(p,t)+1}^t \phi_{B(p)}^u \right\} & \text{otherwise.} \end{cases} \quad (3.2)$$

Intuitively, whenever some fraction δ of a flush (B, t) is chosen, we imagine increasing the fractional amount by which each page in B is evicted to extent δ . On the other hand when a page p_t is requested at time t , we reset its fractional value to $x_{p_t}^t = 0$.

3.2 A k -Competitive Deterministic Online Algorithm

Our first algorithmic result is a k -competitive deterministic online algorithm, which beats the trivial $k\beta$ competitive ratio obtained by running the deterministic online algorithm for classical paging. Our deterministic algorithm is given in Algorithm 1. The algorithm

constructs simultaneously a primal solution x and a dual solution y to the LPs (P) and (D). We will ensure that these solutions satisfy all constraints known up to that time. We use $C(\tau)$ for the set of pages in cache at time τ , and we use S for the set of flushes performed by the algorithm so far. At the start of the algorithm, S is initialized as the set of all flushes of all blocks at time 0; this amounts to clearing the initial cache. The primal solution ϕ is set to the characteristic vector of S , and dual solutions y is initialized as the all-0-vector. At time τ , we first add the requested page p_τ to the cache. If this violates the cache constraint, we continuously increase the dual variable y_S^τ corresponding to the current set S and the current time τ until the dual constraint corresponding to some (B, t) with $f_\tau((B, t) | S) \geq 1$ becomes tight. Once this happens, we evict all pages of block B that are in cache (except p_τ , in case it belongs to this block) and update x and S to reflect that the flush (B, τ) has been performed.

Algorithm 1 Deterministic Online Algorithm

```

1:  $S \leftarrow \{(B, 0) : B \in \mathcal{B}\}$ 
2:  $\phi \leftarrow \chi_S, y \leftarrow \vec{0}$ 
3: for time  $\tau = 1, \dots, T$  do
4:    $C(\tau) \leftarrow C(\tau - 1) \cup \{p_\tau\}$ 
5:   if  $|C(\tau)| > k$  then
6:     Increase  $y_S^\tau$  until the dual constraint corresponding to
       some  $(B, t)$  for which  $f_\tau((B, t) | S) \geq 1$  is tight.
7:      $C(\tau) \leftarrow C(\tau) \setminus (B \setminus \{p_\tau\})$ .
8:      $\phi_B^\tau \leftarrow 1$ .
9:      $S \leftarrow S \cup \{(B, \tau)\}$ .
```

We show that both primal and dual solutions are feasible, and that the cost of the primal is at most k times the cost of the dual. By weak duality, this implies:

Theorem 3.3. Algorithm 1 is k -competitive.

We begin by showing feasibility.

Lemma 3.4. Algorithm 1 terminates. Upon termination, ϕ is feasible for (P) and y is feasible for (D).

Proof. We first show that the algorithm terminates and the primal is feasible. Assume by induction that the algorithm maintains a feasible cache for every time step strictly less than τ (it is trivially feasible at time 0). Since exactly one page is requested per time step, if the cache is not feasible at the beginning of time step τ , then $|C(\tau)| = k + 1$. If this is the case, then there must exist some (B, t) such that $f_\tau((B, t) | S) \geq 1$, in which case the dual constraint corresponding to such a (B, t) will become tight after y_S^τ is increased sufficiently (in particular, the increase of y_S^τ terminates). Note that $t \leq \tau$ in this case. Then $f_\tau((B, t) | S) \geq f_t((B, t) | S) \geq 1$, so at least one page is evicted upon performing the flush (B, τ) , and thus feasibility is restored at time τ . Hence the algorithm maintains a feasible cache state for every time τ , and by Claim 3.2 it follows that the primal is feasible for (P).

We now show feasibility of the dual. Clearly $y_S^\tau \geq 0$. Increasing y_S^τ could lead to a violation of the dual constraint corresponding to (B, t) only if $f_\tau((B, t) | S) \geq 1$, but in this case we stop once the constraint becomes tight. Thus, dual constraints are never violated. Note that dual variables y_S^τ corresponding to future time steps τ are 0, so it does not matter that the coefficients $f_\tau((B, t) | S)$ of future time steps are not known yet. \square

Finally, we relate the primal and dual costs.

Lemma 3.5. The cost of the primal is at most k times the cost of the dual.

Proof. The algorithm sets $\phi_B^\tau = 1$ if and only if $(B, \tau) \in S$, so the primal cost is

$$P = \sum_{B, \tau} c_B \cdot \phi_B^\tau = \sum_{(B, \tau) \in S} c_B.$$

The algorithm adds (B, τ) to S only if a constraint corresponding to some $(B, t) = (B, t_\tau)$ with $f_\tau((B, t_\tau) | S) \geq 1$ becomes tight at time τ . Thus,

$$c_B = \sum_{S', u} f_u((B, t_\tau) | S') \cdot y_{S'}^u. \quad (3.3)$$

To account for our primal cost, for every flush $(B, \tau) \in S$, we charge $f_u((B, t_\tau) | S') \cdot y_{S'}^u$ of the cost of the flush to the dual variable $y_{S'}^u$. Since every non-zero dual variable $y_{S'}^u$ in our final solution has its coefficient in the objective $n - k - f_u(S') \geq 1$ (otherwise it would never have been increased), it suffices to argue that each dual variable $y_{S'}^u$ receives a total charge of at most $k \cdot y_{S'}^u$.

To see this, note that (B, τ) only charges its cost to variables $y_{S'}^u$ for which $u \in [t_\tau, \tau]$. Indeed, $f_u((B, t_\tau) | S') > 0$ only if $t_\tau \leq u$; moreover, when (3.3) becomes tight at time τ we have $y_{S'}^u = 0$ for $u > \tau$ and all S' , and such $y_{S'}^u$ could subsequently increase only if $f_u((B, t_\tau) | S') = 0$ since otherwise the dual constraint would become violated. After (B, τ) is added to S , for all $t' \leq \tau$ and all $\tau' \geq \tau$ the multiplier $f_{\tau'}((B, t') | S) = 0$, and so $y_{S'}^u$ is charged at most once by every block B . Furthermore, if flush (B, τ) charges dual variable $y_{S'}^u$, then the coefficient $f_u((B, t) | S')$ is at most the number of pages from block B that were in cache at the end of time step $u - 1$. Since the total number of pages in cache at the end of time step $u - 1$ was at most k , the total amount charged to $y_{S'}^u$ is at most $k \cdot y_{S'}^u$. \square

3.3 An $O(\log k)$ -Competitive Monotone-Incremental Fractional Algorithm

In this section we give a competitive fractional algorithm for block caching with eviction cost. For simplicity of presentation, our algorithm is not online in the strict sense, as at time τ we allow it to change the value of ϕ_B^t for $t < \tau$. However, it has the crucial property that it only increases LP variables ϕ_B^t . We call an algorithm with this property *monotone-incremental*. This property suffices for

our rounding procedure in Section 3.4 to yield an online algorithm. We prove:

Theorem 3.6. Algorithm 2 is an $O(\log k)$ -competitive monotone-incremental fractional algorithm for block-aware caching with eviction cost.

To describe the algorithm, we define a flush (B, t) to be *alive* at time τ if $t = r(p, \tau) + 1$ for some $p \in B$. Intuitively, for an offline algorithm it is most beneficial to flush a block B only at time steps directly after some page from B was requested. Accordingly, our fractional algorithm will increase ϕ_B^t only for (B, t) that are alive. The algorithm is given in Algorithm 2. It starts by initializing S as the set of all flushes at time 0, ϕ as the corresponding characteristic vector, and y as the all-0-vector. At time τ , when some constraint (S', τ) is violated, we will show that this will also be the case for some $S' \supseteq S$, so that the condition of the while-loop will be true. We then increase the corresponding dual variable as well as primal variables corresponding to all alive flushes according to (3.4). While doing so, we occasionally add new flushes to the set S . As we will show later, all flushes (B, t) added to the set S will satisfy $\phi_B^t = 1$, i.e., they are chosen integrally by the fractional algorithm.

Algorithm 2 $O(\log k)$ -competitive monotone-incremental fractional algorithm

```

1:  $S \leftarrow \{(B, 0) : B \in \mathcal{B}\}$ .
2:  $\phi \leftarrow \chi_S$ ,  $y \leftarrow \vec{0}$ .
3: for time  $\tau = 1, \dots, T$  do
4:   while  $\exists S' \supseteq S$  s.t. primal constraint of  $(S', \tau)$  violated do
5:     Increase  $y_{S'}^\tau$  continuously, and meanwhile for every
       alive  $(B, t)$  increase  $\phi_B^t$  at rate
       
$$\frac{d\phi_B^t}{dy_{S'}^\tau} = \begin{bmatrix} \frac{\ln(k \cdot \beta + 1)}{c_B} \cdot f_\tau((B, t) | S') \\ \cdot \left( \phi_B^t + \frac{1}{k \cdot \beta} \right) \end{bmatrix} \quad (3.4)$$

       until the dual constraint corresponding to some
       alive  $(B_0, t_0)$  for which  $f_\tau((B_0, t_0) | S') \geq 1$  is
       tight.
6:    $S \leftarrow S \cup \{(B_0, t_0)\}$ .
```

Lemma 3.7. Algorithm 2 terminates.

Proof. If the condition of the while-loop is true, then $f_\tau(S') < n - k$, and thus there exists some alive (B_0, t_0) with $f_\tau((B_0, t_0) | S') \geq 1$. Increasing $y_{S'}^\tau$ sufficiently will eventually tighten a corresponding constraint, so each iteration of the while-loop terminates. When a new element is added to S at the end of an iteration, $f_\tau(S)$ increases by at least 1. When $f_\tau(S)$ has reached value $n - k$ (or earlier), the condition of the while-loop cannot be true any more. \square

Lemma 3.8. At the end of Algorithm 2, ϕ is feasible for (P) and y is feasible for (D).

Before proving Lemma 3.8, we need the following claim, whose proof we leave for Appendix A.1.

Definition 3.9. Given a fractional solution ϕ , we say that the constraint (S, τ) is maximal-integral if for any (B, t) such that $\phi_B^t = 1$ it holds that $(B, t) \in S$.

Claim 3.10. If a fractional solution ϕ to (P) has no violated maximal-integral constraints, then ϕ is feasible.

Proof of Lemma 3.8. We first show feasibility of the dual. Suppose the dual constraint corresponding to some (B, t) gets violated when $y_{S'}^\tau$ is increased. Let $t_0 \leq t$ be maximal such that (B, t_0) is alive at time τ . (If no such t_0 exists, then any pages evicted by the flush (B, t) are requested again in $[t, \tau]$; but then $f_\tau((B, t) | S') = 0$, so increasing $y_{S'}^\tau$ would not have led to a violation of the constraint corresponding to (B, t) .) Since no pages of B are requested at times in $[t_0, t)$, we have $f_\tau((B, t_0) | S'') = f_\tau((B, t) | S'')$ for any S'' , meaning that the dual constraint of (B, t_0) would become violated at the same time during the increase of $y_{S'}^\tau$. But then $f_\tau((B, t_0) | S') \geq 1$ (otherwise, increasing $y_{S'}^\tau$ would not increase the left-hand side of constraint (B, t_0)) and therefore we would have stopped increasing $y_{S'}^\tau$ when the constraint got tight.

To see that the primal is feasible, we will show that for all $(B, t) \in S$ we have $\phi_B^t = 1$. It then follows from Claim 3.10 in the appendix that if a primal constraint (S', τ) is infeasible, then this is also the case for some $S' \supseteq S$; thus, the algorithm would not have terminated.

Consider the differential equation $dz/dy = \eta \cdot (z + \delta)$ for some constants $\eta \geq 0$ and $\delta > 0$. When y increases from a to b , we have

$$\ln(z(b) + \delta) - \ln(z(a) + \delta) = \eta \cdot (b - a). \quad (3.5)$$

For some (B, t) that eventually gets added to S , consider the dynamics of ϕ_B^t . It starts at 0, and increases with every $y_{S'}^\tau$ according to (3.4). Applying (3.5) for every such $y_{S'}^\tau$ and summing, we have

$$\begin{aligned} & \ln\left(\phi_B^t + \frac{1}{k \cdot \beta}\right) - \ln\left(\frac{1}{k \cdot \beta}\right) \\ &= \sum_{S', \tau} \frac{\ln(k \cdot \beta + 1)}{c_B} \cdot f_\tau((B, t) | S') \cdot y_{S'}^\tau. \end{aligned}$$

Taking exponents and solving, we have that

$$\phi_B^t = \frac{1}{k \cdot \beta} \cdot \left(\exp\left(\frac{\ln(k \cdot \beta + 1)}{c_B} \sum_{S', \tau} f_\tau((B, t) | S') \cdot y_{S'}^\tau\right) - 1 \right).$$

In particular, when constraint (B, t) becomes tight and is added to S , the value of ϕ_B^t is 1. \square

Lemma 3.11. The cost of the primal is at most $O(\log k)$ times the cost of the dual.

Proof. Consider the algorithm at a fixed time τ during a step in which $y_{S'}^\tau$ is increased by an infinitesimal amount $dy_{S'}^\tau$. The dual profit is $dy_{S'}^\tau \cdot (n - k - f_\tau(S'))$, and so it suffices bound the corresponding increase in primal cost. The primal cost increase is:

$$\begin{aligned} & \sum_{B,t} c_B \cdot \frac{1}{c_B} \ln(k \cdot \beta + 1) \cdot f_\tau((B,t) | S') \cdot \left(\phi_B^t + \frac{1}{k \cdot \beta} \right) \cdot dy_{S'}^\tau \\ &= \sum_{B,t} \ln(k \cdot \beta + 1) \cdot f_\tau((B,t) | S') \cdot \left(\phi_B^t + \frac{1}{k \cdot \beta} \right) \cdot dy_{S'}^\tau. \end{aligned}$$

Since we only increase $y_{S'}^\tau$ if the corresponding constraint in the primal is not satisfied, we have

$$\sum_{B,t} f_\tau((B,t) | S') \cdot \phi_B^t < n - k - f_\tau(S'). \quad (3.6)$$

We claim that

$$\sum_{(B,t) \text{ alive}} \frac{f_\tau((B,t) | S')}{k \cdot \beta} \leq n - k - f_\tau(S'). \quad (3.7)$$

Inequalities (3.6) and (3.7) together imply that the increase in the primal cost is at most $2 \ln(k \cdot \beta + 1) \cdot (n - k - f_\tau(S')) \cdot dy_{S'}^\tau$, which in turn implies the lemma statement since $\beta \leq k$.

To prove (3.7), we first show that

$$\sum_B \frac{f_\tau((B,\tau) | S')}{k} \leq n - k - f_\tau(S'). \quad (3.8)$$

Note that $\sum_B f_\tau((B,\tau) | S') \leq n - f_\tau(S') - 1$. In the case that $n - f_\tau(S') \geq k + 1$, (3.8) holds by the fact that $(z - 1)/k \leq z - k$ for every $z \geq k + 1$. Otherwise, when $n - f_\tau(S') < k + 1$, then $f_\tau(S') = n - k$ (since f_τ is integer valued and truncated at $n - k$). Then (3.8) holds with both sides equal to 0.

We now obtain (3.7) via

$$\begin{aligned} \sum_{(B,t) \text{ alive}} \frac{f_\tau((B,t) | S')}{k \cdot \beta} &\leq \sum_{(B,t) \text{ alive}} \frac{f_\tau((B,\tau) | S')}{k \cdot \beta} \\ &\leq \sum_B \frac{f_\tau((B,\tau) | S')}{k} \\ &\leq n - k - f_\tau(S') \end{aligned}$$

where the second inequality uses that there are at most β flushes alive at any time, and the last inequality is (3.8). \square

3.4 An $O(\log \Delta k)$ -Competitive Online Randomized Rounding Scheme

Finally we show an online $O(\log \Delta k)$ randomized rounding scheme for our block caching LP (P). Recall the definition of the aspect ratio Δ from Section 2 and note that in the standard unweighted setting, $\Delta = 1$.

At time t , the algorithm evicts block B with probability $\gamma \cdot \phi_B^t$, where $\gamma = O(\log k \Delta)$. If the cache is still infeasible at time t , evict an arbitrary block so long as at least one of its pages has $x_p^t > 0$ (recall from the definition (3.2) that x_p^t is the amount missing from page p at time t). For clarity of exposition, the rounding procedure is written as if the underlying fractional solution (x, ϕ)

is computed online. However the procedure can be carried out so long as the solution is monotone-incremental; at time τ , if the fractional solution increases any ϕ_B^t for $t < \tau$ by some amount δ_t , we can evict B at time τ with probability $\min(1, \gamma \cdot (\phi_B^\tau + \sum_{t < \tau} \delta_t))$.

Thus together with Theorem 3.6, our rounding scheme implies:

Theorem 3.12. For block-aware caching with eviction cost, there exists an $O(\log k \log(k\Delta))$ -competitive algorithm.

Furthermore, using the round-or-separate procedure of [23], one can simultaneously solve and round the SUBMODULARCOVER LP (2.1) offline in polynomial time. Using the analysis of this section, this implies:

Theorem 3.13. For block-aware caching with eviction cost, there exists an $O(\log(k\Delta))$ -approximation algorithm.

We perform the rounding assuming a few key properties of our fractional solution which we show we can assume (online) without changing our asymptotic guarantees.

Lemma 3.14. Let (x, ϕ) be a fractional solution for LP (P). For an additional multiplicative constant factor to the competitive ratio, we can assume that (x, ϕ) has the following properties:

- For every time t , every page p has $x_p^t \in [0, 1/2] \cup \{1\}$.
- Every nonzero coordinate has $\phi_B^t \geq 1/4k^2$.

We defer the proof to Appendix A.1.

A key technical tool in this section is a lemma from [23], which in turn relies on a relationship between continuous extensions of submodular functions proven by [32].

Lemma 3.15 (Lemma 2.5 of [23]). Let $x \in [0, 1]^n$ be a feasible solution to (2.1). Let R be a set obtained by performing randomized rounding according to $\min(1, \gamma \cdot x)$. Then:

$$\mathbb{E}_R[f(R)] \geq f(\mathcal{N}) - e^{-\gamma} f(\mathcal{N}).$$

We can now present our rounding scheme.

Algorithm 3 $O(\log k)$ -Approximate Rounding

```

1: for time  $\tau \in [T]$  do
2:   For every block  $B$  evict the set  $\{p \in B \mid x_p^\tau > 0\}$  with
     probability  $\min(1, \gamma \cdot \phi_B^\tau)$ .
3:   Fetch  $p_\tau$  if it is missing from the cache.
4:   while the cache is infeasible do
5:     Let  $B$  be an arbitrary block in the cache that has a page  $p$ 
       with  $x_p^\tau > 0$ , evict the set of pages  $\{p \in B \mid x_p^\tau > 0\}$ .

```

Note that we assume that the fractional solution on which Algorithm 3 executes is one that has the properties given by Lemma 3.14.

We now prove our main rounding lemma.

Lemma 3.16. For $\gamma = \log(4k^2\beta\Delta)$, given a feasible fractional solution (x, ϕ) with cost $c(\phi)$, Algorithm 3 produces a feasible integral cache policy of cost $O(\log k\Delta) \cdot c(\phi)$.

To prove Lemma 3.16, we charge the cost of the algorithm to the fetching cost of the fractional solution. To relate this fractional fetching cost to the fractional eviction cost, we need a claim which we prove in Appendix A.1.

Claim 3.17. Let $c_{\text{Fetch}}(z)$ be the fetching cost of a fractional solution z . Then

$$c_{\text{Fetch}}(z) \leq \beta \left(c(z) + \sum_{B \in \mathcal{B}} c_B \right).$$

Proof of Lemma 3.16. Let x, ϕ be a fractional solution given by Lemma 3.14, and let S be the set of flushes performed by our algorithm.

The algorithm produces a feasible cache policy by construction, as we always fetch p_t and we always run the eviction loop in Lines 4 and 5 until the cache is feasible. Note that there is always a block to evict with a page p that has $x_p > 0$, otherwise x is integral, and is the characteristic vector of the pages the algorithm has in cache, in which case the algorithm's cache is already feasible since the fractional solution is feasible. Furthermore, the expected cost of the evictions due to the randomized rounding step at Line 2 is at most $\gamma \cdot c(\phi) = O(\log k\Delta) \cdot c(\phi)$.

It remains to show that the total cost due to alterations in the eviction loop in Lines 4 and 5 is bounded. We now show that it is at most $O(c(\phi))$.

Let Λ be the set of times τ such that p_τ is not already fully in the fractional cache. Our algorithm maintains the invariant that if $x_p^t = 0$, then it is also fully in cache of the integral solution produced by our algorithm at time t . This means that at times $\tau \notin \Lambda$, neither the fractional solution nor the rounding algorithm incur a cost increase. Hence we focus on the case where $\tau \in \Lambda$.

For every τ , the solution ϕ is feasible for the LP (2.1) with the function f^τ , so by Lemma 3.15

$$\mathbb{E}[f_\tau(S)] \geq n - k - \frac{1}{4k^2\beta\Delta}.$$

In particular, this holds for all $\tau \in \Lambda$. In words, the expected number of pages in cache is bounded by $k + \frac{1}{4k^2\beta\Delta}$. Since every eviction due to Line 5 costs at most c_{\max} and evicts at least one page, the expected cost of the alteration while loop at time τ is bounded by $c_{\max}/(4k^2\beta\Delta) = c_{\min}/4k^2\beta$.

On the other hand, since $\tau \in \Lambda$, the page p_τ is not fully in cache, and since by Lemma 3.14 the fractional solution evicts pages in increments of at least $1/(4k^2)$, it holds that $x_{p_\tau}^\tau \geq 1/(4k^2)$. This means that the fetching cost of the fractional solution at time τ is at least $c_{\min}/4k^2$.

Hence the expected cost of the alteration step in time τ is at most the fractional fetching cost at time τ , divided by β . Summing this inequality over time, the total cost paid by the algorithm over all all time due to Line 5 is at most $c_{\text{Fetch}}(\phi)/\beta$. By Claim 3.17, the fetching cost $c_{\text{Fetch}}(\phi) \leq \beta(c(\phi) + \sum_{B \in \mathcal{B}} c_B)$, and hence the total cost of alterations is at most $c(\phi) + \sum_{B \in \mathcal{B}} c_B$. This completes the proof. \square

4 FETCHING COST

We present our $\Omega(\beta)$ lower bound against randomized algorithms for online block-aware caching with fetching costs. We first present a bicriteria rounding algorithm for the naïve LP of Appendix A.2. We then argue that this procedure can be used to derandomize any randomized algorithm for block-aware caching with fetching costs. Together with the lower bound against deterministic algorithms given by [10], this implies a lower bound against randomized algorithms.

4.1 Bicriteria Online Rounding Algorithm

Consider the following deterministic online rounding scheme. For every page p , evict p from the cache at time t if $x_p^t > 1/2$. If a page p_t is not in cache upon request at time t , then at time t fetch all pages from $B(p_t)$ such that $x_p^t \leq 1/2$.

Theorem 4.1. Given a feasible fractional solution x to the block-aware caching problem, the procedure above produces an integral solution that uses at most $2k$ cache space at any point in time, and whose fetching cost is at most twice the fetching cost of x .

Proof. The procedure produces a feasible solution by construction, since $x_{p_t}^t = 0 \leq 1/2$. It also violates the cache size constraint by at most a factor of 2, since no page is present in the integral cache unless $x_p^t \leq 1/2$, meaning the fractional cache usage is at least half the integral cache usage.

Finally, to justify that the integral solution has cost at most twice the fractional cost, charge the cost of integrally loading $B(p_t)$ to the fractional decrease of x_B^t since the last time t' at which $B(p_t)$ was loaded. Since p_t had $x_p^{t'} > 1/2$ (otherwise we would have loaded it earlier), the fractional cost incurred since time t' was at least $1/2 \cdot c_{B(p)}$. \square

Corollary 4.2. When $k = 2h$, there is a 2-competitive offline algorithm for block-aware caching with fetching cost.

We mention briefly that a similar rounding procedure produces a cache policy that is 2-competitive with the eviction cost of the fractional solution, and also uses at most a factor 2 more space. If p_t is not in cache at time t , fetch it. On the other hand if any page in cache at time t has fractional value $x_p^t < 1/2$, evict all of $B(p)$.

4.2 Lower Bounds for Randomized Algorithms

Finally we turn to showing our lower bound. Our starting point is the lower bound of [10] against deterministic algorithms.

Theorem 4.3 (Theorem 4.1 of [10]). The competitive ratio of any deterministic online policy for block-aware caching with fetching costs is at least

$$\frac{k + (B - 1)(h - 1)}{k - h + 1}$$

for $h \leq k - B + 1$.

We now show how to use the online deterministic rounding procedure of Section 4.1 to derandomize any online algorithm for Block-Aware caching with fetching costs. This proves the main claim of this section:

Theorem 4.4. The competitive ratio of any randomized policy for block-aware caching with fetching costs is at least

$$\frac{2k + (B - 1)(h - 1)}{4k - 2h + 2}$$

for $h \leq k - B + 1$.

Proof. Suppose there is a randomized online algorithm \mathcal{R} for (h, k) -block-aware caching with fetching costs with expected cost $c_{\mathcal{R}}$. Then we can convert this randomized cache policy online to a fractional solution x . To do so, set x_p^t be the expected value of the indicator of whether page p is loaded at time t . Note that these expectations can be computed using only the sequence of requests up to and including time t . This solution x is feasible to the simple fetching cost LP (A.1), and furthermore has LP cost $c_{\mathcal{R}}$.

Applying Theorem 4.1 to the fractional solution x produces an integral cache policy cost at most $2 \cdot c_{\mathcal{R}}$ and space $2k$. The claim follows by using the lower bound on the cost of any such policy given by Theorem 4.3, and solving for $c_{\mathcal{R}}$. \square

Combining this with the well known $\Omega(\log k)$ lower bound for randomized algorithms for classical paging, we obtain the following consequence.

Corollary 4.5. When $k = O(h)$, no randomized algorithm has competitive ratio better than $\Omega(B + \log k)$.

REFERENCES

- [1] Dimitris Achlioptas, Marek Chrobak, and John Noga. 2000. Competitive analysis of randomized paging algorithms. *Theoretical Computer Science* 234, 1-2 (2000), 203–218.
- [2] Anna Adamaszek, Artur Czumaj, Matthias Englert, and Harald Räcke. 2019. An $O(\log k)$ -Competitive Algorithm for Generalized Caching. *ACM Trans. Algorithms* 15, 1 (2019), 6:1–6:18.
- [3] Susanne Albers, Sanjeev Arora, and Sanjeev Khanna. 1999. Page replacement for general caching problems. In *Symposium on Discrete Algorithms*. 31–40.
- [4] N. Alon, B. Awerbuch, Y. Azar, N. Buchbinder, and J. Naor. 2003. The online set cover problem. In *Symposium on the Theory of Computation*. 100–105.
- [5] Nikhil Bansal, Niv Buchbinder, and Joseph Naor. 2012. A Primal-Dual Randomized Algorithm for Weighted Paging. *J. ACM* 59, 4 (2012), 19.
- [6] Nikhil Bansal, Niv Buchbinder, and Joseph Naor. 2012. Randomized Competitive Algorithms for Generalized Caching. *SIAM J. Comput.* 41, 2 (2012), 391–414.
- [7] Nikhil Bansal, Joseph Naor, and Ohad Talmon. 2021. Efficient Online Weighted Multi-Level Paging. In *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures*. 94–104.
- [8] Nathan Beckmann, Phillip B. Gibbons, Bernhard Haeupler, and Charles McGuffey. 2020. Writeback-aware caching. In *Symposium on Algorithmic Principles of Computer Systems*. SIAM, 1–15.
- [9] Nathan Beckmann, Phillip B. Gibbons, Bernhard Haeupler, and Charles McGuffey. 2020. Writeback-Aware Caching. In *Algorithmic Principles of Computer Systems, APOCS@SODA, 2020*. 1–15.
- [10] Nathan Beckmann, Phillip B. Gibbons, and Charles McGuffey. 2021. Brief Announcement: Block-Granularity-Aware Caching. (2021).
- [11] Avrim Blum, Carl Burch, and Adam Kalai. 1999. Finely-Competitive Paging. In *Symposium on Foundations of Computer Science (FOCS)*. 450.
- [12] A. Blum, M. Furst, and A. Tomkins. 1996. What to do with your free time: algorithms for infrequent requests and randomized weighted caching.
- [13] Sébastien Bubeck, Michael B. Cohen, Yin Tat Lee, James R. Lee, and Aleksander Mądry. 2018. K-server via multiscale entropic regularization. In *symposium on theory of computing*. 3–16.
- [14] Niv Buchbinder, Shahar Chen, and Joseph Naor. 2014. Competitive analysis via regularization. In *Symposium on Discrete Algorithms*. 436–444.
- [15] Niv Buchbinder, Joseph Seffi Naor, et al. 2009. The design of competitive online algorithms via a primal-dual approach. *Foundations and Trends® in Theoretical Computer Science* 3, 2–3 (2009), 93–263.
- [16] Marek Chrobak and Lawrence L. Lamore. 1991. An Optimal On-Line Algorithm for k-Servers on Trees. *SIAM J. Comput.* 20, 1 (1991), 144–148.
- [17] Christian Coester, Roie Levin, Ohad Talmon, et al. 2022. Competitive Algorithms for Block-Aware Caching. *arXiv preprint arXiv:2205.12249* (2022).
- [18] Guy Even, Moti Medina, and Dror Rawitz. 2018. Online Generalized Caching with Varying Weights and Costs. In *Symposium on Parallelism in Algorithms and Architectures, SPAA*. 205–212.
- [19] Amos Fiat, Richard M. Karp, Michael Luby, Lyle A. McGeoch, Daniel Dominic Sleator, and Neal E. Young. 1991. Competitive Paging Algorithms. *Journal of Algorithms* 12, 4 (1991), 685–699.
- [20] Anupam Gupta, Ravishankar Krishnaswamy, Amit Kumar, and Debmalya Panigrahi. 2019. Elastic Caching. In *Symposium on Discrete Algorithms, SODA*. 143–156.
- [21] Anupam Gupta, Amit Kumar, and Debmalya Panigrahi. 2020. Caching with Time Windows. In *Symposium on Theory of Computing*.
- [22] Anupam Gupta and Roie Levin. 2020. Fully-Dynamic Submodular Cover with Bounded Recourse. In *2020 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE.
- [23] Anupam Gupta and Roie Levin. 2020. The online submodular cover problem. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1525–1537.
- [24] Syed Hasan, Sergey Gorinsky, Constantine Dovrolis, and Ramesh K. Sitaraman. 2014. Trade-offs in optimizing the cache deployments of CDNs. In *2014 IEEE Conference on Computer Communications, INFOCOM 2014, Toronto, Canada, April 27 - May 2, 2014*. IEEE, 460–468.
- [25] Sandy Irani. 1996. Competitive analysis of paging: A survey. In *Proc. of the Dagstuhl Seminar on Online Algorithms*.
- [26] Sandy Irani. 1997. Page replacement with multi-size pages and applications to Web caching. In *Symposium on Theory of computing*. 701–710.
- [27] S. Irani. 2002. Randomized weighted caching with two page weights. *Algorithmica* 32, 4 (2002), 624–640.
- [28] Thodoris Lykouris and Sergei Vassilvitskii. 2018. Competitive Caching with Machine Learned Advice. In *International Conference on Machine Learning, ICML, 2018, Vol. 80*. 3302–3311.
- [29] Lyle A. McGeoch and Daniel D. Sleator. 1991. A strongly competitive randomized paging algorithm. *Algorithmica* 6, 6 (1991), 816–825.
- [30] Daniel D. Sleator and Robert E. Tarjan. 1985. Amortized efficiency of list update and paging rules. *Commun. ACM* 28, 2 (1985), 202–208.
- [31] Zhenyu Song, Daniel S. Berger, Kai Li, and Wyatt Lloyd. 2020. Learning Relaxed Belady for Content Distribution Network Caching. In *17th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2020, Santa Clara, CA, USA, February 25-27, 2020*, Ranjita Bhagwan and George Porter (Eds.). USENIX Association, 529–544.
- [32] Jan Vondrák. 2007. Submodularity in combinatorial optimization. (2007).
- [33] L. A. Wolsey. 1982. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica* 2, 4 (01 Dec 1982), 385–393. <https://doi.org/10.1007/BF02579435>
- [34] Neal Young. 1991. On-line caching as cache size varies. In *Symposium on Discrete Algorithms*. 241–250.
- [35] Neal E. Young. 1994. The k-Server Dual and Loose Competitiveness for Paging. *Algorithmica* 11, 6 (1994), 525–541.
- [36] Neal E. Young. 1998. On-line file caching. In *Symposium on Discrete Algorithms*. 82–86.

A APPENDIX

A.1 Deferred Proofs

Claim 2.1. There exist instances of block-aware caching for which the optimal fetching cost is β larger than the optimal

eviction cost, and there exist instances where the optimal eviction cost is β larger than the optimal fetching cost.

Proof. Consider the following instance. For any β , let $n = 2\beta^2$ pages be organized into 2β blocks of size β . Let P be the first β blocks and Q be the second β blocks. We set $k = \beta^2$, and fill the cache initially with all the pages of the P blocks. The request sequence consists of rounds. For $i = 1, \dots, \beta$, in round i request the first $\beta - i$ pages of each P block, and the first i of the Q blocks in their entirety, and repeat this sequence L times within the round. For sufficiently large constant L , the optimal solution must have precisely the requested pages of a round in its cache. Thus, in round i it evicts the $(\beta - i + 1)^{th}$ page of each P block and fetches the i^{th} Q block in its entirety. The fetching cost of this solution is β , while the eviction cost is β^2 .

To see the other direction, observe that if we instead start the cache with the pages of the Q blocks, and for $i = 1, \dots, \beta$ we make round i request precisely the pages not requested in round i above (and once again repeat this sequence L times), the optimal solution will always evict one Q block in its entirety and fetch a single page from each P block in each round. Here the fetching cost is β^2 and the eviction cost is β . \square

Claim 3.1. For every τ , the function f_τ is submodular.

Proof. Consider the function g^τ , where

$$g^\tau(S) := |\{p : p \text{ is missing at time } \tau \text{ according to } S\}| \\ = \left| \bigcup_{\phi \in S} \{p : p \text{ is missing at time } \tau \text{ according to } \phi\} \right|.$$

g^τ is a coverage function, and hence it is submodular. The function f^τ is the minimum of g^τ and the constant function $n - k$, and so f^τ is also submodular. \square

Claim 3.10. If a fractional solution ϕ to (P) has no violated maximal-integral constraints, then ϕ is feasible.

Proof. It suffices to show that if ϕ violates a constraint (S, τ) , and (B_0, t_0) is such that $\phi_{B_0}^{t_0} = 1$, then x also violates $(S \cup \{(B_0, t_0)\}, \tau)$. Since x is violated,

$$n - k - f_\tau(S) > \sum_{B,t} f_\tau((B, t) | S) \cdot \phi_B^t \\ = \sum_{(B,t) \neq (B_0, t_0)} f_\tau((B, t) | S) \cdot \phi_B^t + f_\tau((B_0, t_0) | S) \\ \geq \sum_{(B,t) \neq (B_0, t_0)} f_\tau((B, t) | S \cup \{(B_0, t_0)\}) \cdot \phi_B^t \\ + f_\tau((B_0, t_0) | S)$$

where the second inequality above used submodularity. Rearranging gives that

$$\sum_{(B,t) \neq (B_0, t_0)} f_\tau((B, t) | S \cup \{(B_0, t_0)\}) \cdot \phi_B^t < n - k - f_\tau(S \cup \{(B_0, t_0)\}). \quad \square$$

Lemma 3.14. Let (x, ϕ) be a fractional solution for LP (P). For an additional multiplicative constant factor to the competitive ratio, we can assume that (x, ϕ) has the following properties:

- For every time t , every page p has $x_p^t \in [0, 1/2] \cup \{1\}$.
- Every nonzero coordinate has $\phi_B^t \geq 1/4k^2$.

Proof. To guarantee the first property, every time a page from a block B is evicted to extent $1/2$, evict the entire block B for a cost of c_B . Charge this eviction to the evictions that caused this page to go from 0 to $1/2$, which cost at least $c_B/2$.

To ensure the second property, consider the following online algorithm.

Algorithm 4 Structure Solution

- 1: Define solution \tilde{x} such that $\tilde{\phi}_B^t = \phi_B^t + \mathbb{1}\{\phi_B^t \geq 1/2\} \cdot (1 - \phi_B^t)$.
 - 2: for block B do
 - 3: Set $t_B \leftarrow 0$.
 - 4: for time $t \in [T]$ do
 - 5: Let $\Delta = \sum_{t'=t_B+1}^t \tilde{\phi}_B^{t'}$.
 - 6: if $\Delta \geq 1/4k^2$ then
 - 7: $\phi_B^t \leftarrow \Delta$.
 - 8: $t_B \leftarrow t$.
 - 9: Output $\tilde{\varphi} = \min(2 \cdot \varphi, 1)$.
-

The structural guarantee that every nonzero coordinate has $\varphi_B^t \geq 1/4k^2$ holds by construction. The cost is also less than $2 \cdot c(\phi)$ by construction.

It remains to show φ is feasible. Consider any constraint of the form:

$$\sum_{B,t} f_\tau((B, t) | S) \cdot \phi_B^t \geq n - k - f_\tau(S)$$

For a block B , let τ_B be the last time before τ that t_B was set to in Line 8. By construction

$$\sum_{t=\tau_B+1}^{\tau} \phi_B^t \leq \frac{1}{4k^2}.$$

Since $f((B, t)|S) \leq k$, and by the property that every p has $x_p^t \in [0, 1/2] \cup \{1\}$, there are at most $2k$ blocks with nonzero pages in cache. This also means

$$\sum_B \sum_{t=\tau_B+1}^{\tau} f((B, t) | S) \cdot \phi_B^t \leq \frac{1}{2}$$

and hence

$$\sum_{B,t} f_\tau((B, t) | S) \cdot \phi_B^t \geq n - k - f_\tau(S) - \frac{1}{2}.$$

If $n - k - f_\tau(S) = 0$, then the constraint is also trivially satisfied by φ . Else $n - k - f_\tau(S) \geq 1$. To conclude, note that since S is maximal-integral, and ϕ has no coordinates $\phi_B^t \in (1/2, 1)$:

$$\begin{aligned} \sum_{t \leq \tau_0} \sum_B f((B, t) | S) \cdot \varphi_B^t &= 2 \sum_{t \leq \tau_0} \sum_B f((B, t) | S) \cdot \phi_B^t \\ &\geq 2 \left(n - k - f_\tau(S) - \frac{1}{2} \right) \\ &\geq n - k - f_\tau(S) \end{aligned}$$

If φ satisfies all integral-maximal constraints, Claim 3.10 implies it also satisfies all other constraints, and the lemma statement follows. \square

Claim 3.17. Let $c_{\text{Fetch}}(z)$ be the fetching cost of a fractional solution z . Then

$$c_{\text{Fetch}}(z) \leq \beta \left(c(z) + \sum_{B \in \mathcal{B}} c_B \right).$$

Proof. Let $\bar{c}(z)$ and $\bar{c}_{\text{Fetch}}(z)$ be the classic paging eviction/ fetching cost of z , i.e. the cost if page fetches/evictions cannot be batched in blocks. The difference between the total fetching cost and eviction cost paid for a single page is at most $c(B(p))$, and hence $\bar{c}(z) = \bar{c}_{\text{Fetch}}(z) \pm \sum_{B \in \mathcal{B}} c_B \cdot \beta$. On the other hand, $\bar{c}(z) \leq \beta \cdot c(z)$. Combining these observations:

$$c_{\text{Fetch}}(z) \leq \bar{c}_{\text{Fetch}}(z) \leq \bar{c}(z) + \beta \cdot \sum_{B \in \mathcal{B}} c_B \leq \beta \cdot \left(c(z) + \sum_{B \in \mathcal{B}} c_B \right). \quad \square$$

A.2 The Natural LP has $\Omega(\beta)$ Integrality Gap

Consider the following simple LP formulation, where $\sigma \in \{-1, 1\}$ is a fixed constant. We use x_p^t for the fraction of page p missing from the cache at time t .

$$\begin{aligned} &\min_{\phi, x} \sum_{B, t} c_B \cdot \phi_B^t \\ &\text{subject to} \\ &\forall t \in [T] : x_{p(t)}^t = 0 \\ &\forall t \in [T], \forall B \in \mathcal{B}, \forall p \in B : \phi_B^t \geq \sigma (x_p^t - x_p^{t-1}) \\ &\forall t \in [T] : \sum_p x_p^t \geq n - k \\ &\forall t \in [T], \forall B \in \mathcal{B} : \phi_B^t \in [0, 1] \\ &\forall t \in [T], \forall p \in [n] : x_p^t \in [0, 1] \end{aligned} \quad (\text{A.1})$$

If $\sigma = 1$, this is the eviction cost model and ϕ_B^t denote the fractional extent to which B is evicted at time t ; if $\sigma = -1$, this is the fetching cost model and ϕ_B^t denote the fractional extent to which B is fetched at time t .

Unfortunately, for both fetching and eviction cost models, this LP has an integrality gap of $\Omega(\beta)$. Consider the following instance in

which $n = 2\beta$ pages are divided into two blocks B_1 and B_2 . The cache is of size $k = 2\beta - 1$, and is initially empty. The request sequence repeats for several rounds. In each round, it requests first all pages from B_1 and then all pages from B_2 .

The integral algorithm must pay at least 1 per round, since 2β pages are requested and the cache is of size $2\beta - 1$. On the other hand, the fractional solution begins by loading both blocks to extent $(\beta - 1)/\beta$. Subsequently, when B_i is requested for $i \in \{1, 2\}$, it loads B_i to extent 1 and B_j (where $j \neq i$) to extent $(\beta - 1)/\beta$. Hence both the fractional fetching and eviction costs per round are $2/\beta$. We summarize this observation in the following theorem.

Theorem A.1. The simple LP relaxation for block-aware caching with both fetching/eviction cost models has an integrality gap of $\Omega(\beta)$.