

Locality-of-Reference Optimality of Cache-Oblivious Algorithms*

Peyman Afshani[†] John Iacono[‡] Varunkumar Jayapaul[§] Ben Karsin[¶]
Nodari Sitchinava^{||}

Abstract

The program performance on modern hardware is characterized by *locality of reference*, that is, it is faster to access data that is close in address space to data that has been accessed recently than data in a random location. This is due to many architectural features including caches, prefetching, virtual address translation and the physical properties of a hard disk drive; attempting to model all the components that constitute the performance of a modern machine is impossible, especially for general algorithm design purposes. What if one could prove an algorithm is asymptotically optimal on all systems that reward locality of reference, no matter how it manifests itself within reasonable limits? We show that this is possible, and that excluding some pathological cases, cache-oblivious algorithms that are asymptotically optimal in the ideal-cache model are asymptotically optimal in any reasonable setting that rewards locality of reference. This is surprising as the cache-oblivious framework envisions a particular architectural model involving blocked memory transfer into a multi-level hierarchy of caches of varying sizes, and was not designed to directly model locality-of-reference correlated performance.

1 Introduction

Modeling memory access time of modern computers is an important area of research that lies at the intersection of theoretical computer science, algorithm engineering, and practical aspects of computing. The reality is that modern computers are extremely complicated with numerous components that try to reduce the access time of the elements in the memory. Consequently, the access time varies by orders of magnitude depending on whether favorable conditions are met or not. Choices in the algorithm design can highly impact reaching those favorable conditions, which necessitates building good theoretical models of memory structures of modern computers.

The main direction of existing theoretical work has been on modeling of the memory hierarchy. This has been initiated by the Disk Access Model (DAM) [3], which is also known as the I/O model or the External Memory (EM) model. The DAM assumes a (fast) memory of size M and a disk (i.e., slow memory) of infinite size. The disk stores the input in blocks of size B and can be read or written to via the input-output (I/O) operations, where each I/O transfers one block of data at unit cost. The analysis then typically only considers the number of I/Os (read or written), ignoring any other computational cost. The justification is that since a disk is so much slower than internal memory, minimizing the number of block transfers and ignoring all else is a good model of runtime, however, sometimes this is not a realistic assumption. For example, when DAM is used to model cached memory (modeling cache misses as I/O operations, M as the size of the cache, and B as the size of the cache lines), the relative difference in the cost of a cache misses compared to the arithmetic operations is much smaller. Aside from this, the DAM model has additional limitations, e.g., it ignores the fact that accessing adjacent blocks on a disk is in practice much faster than two random blocks [12] and it models only two levels of memory.

Modeling more than two levels of the memory hierarchy is rather challenging. The big problem is that very precise models (e.g., the ones defining individual parameters for each level of memory hierarchy [13]) are often too complicated, making it hopelessly difficult to design and analyze algorithms. Other approaches, such as the

*The full version of the paper can be accessed at <https://arxiv.org/abs/1902.07928>

[†]Aarhus University, Denmark.

[‡]Université Libre de Bruxelles and New York University. Supported by Fonds de la Recherche Scientifique-FNRS under Grant no MISU F 6001 1 and NSF grant CCF-1911245.

[§]Université Libre de Bruxelles. Research partially completed while at Université Libre de Bruxelles, supported by Fonds de la Recherche Scientifique-FNRS under Grant no MISU F 6001 1.

[¶]NVIDIA. Research partially completed while at Université Libre de Bruxelles, supported by Fonds de la Recherche Scientifique-FNRS under Grant no MISU F 6001 1.

^{||}University of Hawaii at Manoa, USA. Partially supported by NSF grants CCF-1911245.




Figure 1: On the right, one can see a hierarchical model of modern computers, with multiple levels of permanent storage, look up tables for virtual memory and CPU cache. From top to bottom, the levels become smaller but faster. On the left, a sample access time is shown (not to scale). Here, it is assumed that each level contains a continuous portion of the elements from the level above it. As a result, the access time increases as we move away from the elements that are located in the fastest memory, as a complicated function of the distance. For caches, the access time dramatically increases when we hit an element that is stored in the next (slower) level. For mechanical devices (tapes, hard disks), the access time increases rather smoothly due to mechanical processes involved.

hierarchical memory model (HMM) [1, 2], model memory with variable access costs by assuming that the cost to access a memory address x is a non-decreasing function, $f(x)$, of the address itself. However, this does not accurately represent modern caches.

The most successful attempt at analyzing cache misses in multi-level cache hierarchies is probably the *cache-oblivious* framework [8]. It surprisingly avoids the complexity of modeling memory hierarchies by completely avoiding it: the algorithms are designed “obliviously to M and B ”, i.e., in the classical (single-level) RAM model. However, if such algorithms are analyzed in the two-level DAM with the best cache-management policy, also known as the ideal-cache model, and they happen to be efficient in terms of cache misses with respect to the two levels of memory, then they are also efficient for all levels of multi-level memory hierarchy. Moreover, it has been shown [11] that such algorithms are also efficient for many reasonable cache management policies, e.g., least-recently used (LRU) policy, typically implemented by hardware in practice. However, up to now, cache-oblivious algorithms haven’t been shown to optimize anything beyond cache misses.

Capturing locality of reference. In a real hardware, cache utilization is only one aspect that affects program runtime. For example, Jurkiewicz and Mehlhorn [9] show that the time it takes to perform address translations for virtual memory noticeably affects the runtime of programs on real hardware. Figure 1 illustrates the complicated ways various hardware features affect the memory access time. It shows a memory hierarchy, where cheaper and larger but slower memories are placed at the top. The functions that describe the access times of the components may have different behaviors, e.g., for a magnetic tape the access time is basically a linear function of the physical distance, whereas for hard disks it is a more complicated function.

Locality of reference is a fundamental principle of computing that heavily impacts both hardware and algorithm design [12]. In his widely-cited article [7], Peter Denning provides an overview of the history of the concept and how it became very popular in almost all aspects of computing. To quote him, “The locality was adopted as an idea almost immediately by operating systems, database, and hardware architects.”

The DAM algorithm and the cache-oblivious algorithms try to capture *spatial locality* and *temporal locality* – the two fundamental components of the notion of “locality of reference” – from an algorithmic point of view. However, given the complexity of modern hardware, the main question is if other, possibly more complex hardware features, can be modeled simply enough to facilitate design and analysis of algorithms. The approach of the HMM model [1, 2] to model the cost of access as a function of memory address is one way to keep the modeling complexity at bay. For instance the authors [1, 2] show that if the cost of accessing address x is $\log(x)$ then sorting n elements can be done in $O(n \log n \log \log n)$ time. But is $\log(x)$ the correct function for modern (and ever-changing) hardware?

1.1 Our Results. We propose to pick up where the previous attempts have left off by following a holistic approach. We present the *locality of reference (LoR)* model, a computational model that looks at memory in a new way: the cost of a memory access is based on the proximity from prior accesses via what we call a *locality function*.

More specifically, we consider the machine as having an infinite memory with a linear *address space*, i.e., memory cells are numbered with the set of natural numbers. Let $E = \{e_1, \dots, e_{|E|}\}$ be the sequence of memory addresses accessed by an algorithm A while running on a given input. A simple locality function ℓ can define the cost of accessing address e_i as a function of the *distance* from the address e_{i-1} of the preceding access, for example, $\log(|e_i - e_{i-1}|)$, $\sqrt{|e_i - e_{i-1}|}$, or any other arbitrarily function of $|e_i - e_{i-1}|$.

A specific locality function can capture the (complicated) cost of accessing the data on the hardware running the algorithm. For example, in the classical random access memory (RAM) model, the sequence E simply takes $O(|E|)$ time, so setting ℓ to a constant function captures the RAM model. As we show later, by setting ℓ to a logarithmic function, we can model the cost of the TLB in virtual memory translation.

The goal of this paper is not to define locality functions for various models of computation. Instead, our results show that cache-oblivious algorithms go beyond minimizing the number of cache misses. In particular, we show that optimal cache-oblivious algorithms are *locality-of-reference optimal*, meaning, they are asymptotically optimal with respect to *any* choice of locality function, subject to some mild constraints (needed to ensure that these functions reward locality of reference). In Section 3 we present our result in a simplified setting, in which we focus on the algorithms that do not benefit from large cache sizes. In Section 4 we generalize the result to more general algorithms that do utilize the full cache.

1.2 Example application: optimality of van Emde Boas layout. We now demonstrate an example application of our results which will appear as Theorem 3.1: namely, that the van Emde Boas layout—a layout for implicit static search trees that is optimal for cache-oblivious searching [8]—is also optimal for address translation on modern virtual memory architectures. This cost of address translation is non-negligible and has been observed to impact the performance of fundamental algorithms such as sorting and permuting in practice [9].

Let us review modern virtual memory design. Consider a machine that uses U bits for addressing memory. Virtual memory is implemented as a trie of degree 2^b , for some parameter b , where the translation process translates b bits at a time, starting from the most significant bits of the address. In the worst-case, one translation necessitates U/b lookups, but this cost is often much lower when TLB caching is used. In particular, if two addresses e_i and e_{i-1} share kb most significant bits, then after translating the first memory address, the first k steps of the second translation are cached in the TLB. As a result, the cost function associated with virtual memory translation is essentially $\ell(|e_i - e_{i-1}|) = \log_b(|e_i - e_{i-1}|)$. This function is clearly non-negative, non-decreasing and concave, thus, satisfying the requirements of Theorem 3.1.

The van Emde Boas layout of an implicit complete binary search tree (for brevity we'll call it *vEB tree*) is defined as follows. Given a complete binary search tree on n vertices of height $h = \lceil \log n \rceil$, let T_0 be the top subtree of height $\lceil h/2 \rceil$ and let $T_1, \dots, T_{\sqrt{n}}$ be the subtrees rooted at the children of the leaves of T_0 . Then, vEB tree is defined by placing the subtree T_0 in a contiguous portion of an array, immediately followed by $T_1, T_2, \dots, T_{\sqrt{n}}$, with each subtree T_i laid out recursively. Search on vEB tree is known to incur $O(\log_B n)$ cache misses, where B is the block size of the DAM model [8].

Consider a machine that is equipped with virtual memory with parameters $U = O(\log n)$ and $b = O(1)$. The root-to-leaf traversal of vEB tree recursively traverses T_0 , jumps by at most $n - \sqrt{n}$ addresses from a leaf of T_0 to the root of the appropriate subtree T_i , and recursively traverses T_i . Thus, the cost $Q_\ell^{\text{LoR}}(n)$ of the traversal using the above locality of reference function ℓ for virtual memory can be computed using the following recurrence:

$$Q_\ell^{\text{LoR}}(n) = \begin{cases} 2 \cdot Q_\ell^{\text{LoR}}(\sqrt{n}) + O(\log n) & \text{if } n \geq 2 \\ O(1) & \text{otherwise} \end{cases}$$

This solves to $Q_\ell^{\text{LoR}}(n) = O(\log n \log \log n)$. And Theorem 3.1 implies that this is an asymptotically optimal LoR cost for virtual address translation for the above function ℓ .

1.3 Related work. The closest work to our presentation here is the Hierarchical Memory Model [1]. In this model, accessing memory location x takes time $f(x)$. This was extended to a blocked version where accessing k memory locations takes times $f(x) + k$ [2]. In particular the case where $f(x) = \log x$ was studied and optimality obtained for a number of problems. This model, through its use of the memory cost function f , bears a number of

similarities to ours, and it is meant to represent a multi-level cache where the user manually controls the movement of data from slow to fast memory. However, while it is able to capture temporal coherence well, even in the blocked version it does not capture fully the idea of spatio-temporal locality of reference, where an access is fast because it was close to something accessed recently.

Another model that proposed analyzing algorithm performance on a multi-level memory hierarchy is the Uniform Memory Hierarchy model (UMH) [4]. The UMH model is a multi-level variation of the DAM that simplifies analysis by assuming that the block size increases by a fixed factor at each level and that there is a uniform ratio between block and memory size at each level of the hierarchy. Unfortunately, this assumption is quite restrictive and doesn't hold in practice on modern hardware.

2 Preliminaries

2.1 Models of computation Let P be a problem, $\mathcal{I}^P = \{I_1, I_2, \dots\}$ be a set of valid instances (input sequences) for which problem P can be solved, and $\mathcal{I}_n^P = \{I \in P : |I| = n\}$ be the subset of instances of P with input of size n . Let $E(A, I) = \{e_1, e_2, \dots\}$ be the sequence of accesses (reads and writes) to memory locations that arises by executing algorithm A on instance I , and let $\mathcal{A}_P = \{A_1, A_2, \dots\}$ denote the set of all algorithms that correctly solve P , i.e., generate a correct output for every instance in \mathcal{I}^P .

The general ideal-cache and LRU-cache models incorporate memory size M and block size B when computing the cost of an execution sequence. The $\frac{M}{B}$ blocks stored in internal memory make up the *working set*, and we define $\mathcal{W}_{M,B}^{\mathbb{M}}(E, i)$ to be the working set after the i -th access to the execution sequence E in model \mathbb{M} . In the ideal-cache model (and the DAM model) the evictions from the working set $\mathcal{W}_{M,B}^{\text{OPT}}$ are selected such that the total cost of executing E is minimized [8], while in the LRU-cache model the evictions from the working set $\mathcal{W}_{M,B}^{\text{LRU}}$ follow the *least recently used* policy [12]. Let $Q_x^{\mathbb{M}}(E(A, I))$ denote the cost of executing algorithm A on instance I in model \mathbb{M} with model parameters x . Then the cost of algorithm A on instance I in the ideal-cache or LRU-cache models with cache parameters M and B is

$$Q_{M,B}^{\mathbb{M}}(E(A, I)) = \sum_{i=1}^{|E(A,I)|} \begin{cases} 0 & \text{if } e_i \in \mathcal{W}_{M,B}^{\mathbb{M}}(E(A, I), i-1) \\ 1 & \text{otherwise} \end{cases}$$

for $\mathbb{M} \in \{\text{CO}, \text{LRU}\}$. If $E(A, I)$ only depends on the input instance I and not on any of the parameters of the model (e.g., M or B) then the algorithm is *cache-oblivious*. Effectively, a cache-oblivious algorithm is one that runs in the classical RAM (with one level of memory). On the other hand, if the algorithm explicitly uses multiple levels of memory, or if the access pattern $E(A, I)$ depends on the particular values of M or B , then it is called *cache-aware*. A more rigorous and formal definition of the working set, cache replacement policies, and the cache-oblivious and LRU costs is included in the full version of the paper. In Sections 3 and 4 we also define the cost $Q_\ell^{\text{LoR}}(E(A, I))$ of an algorithm on instance I in our memoryless and general *Locality of Reference (LoR)* models with locality function ℓ .

Similarly, we define $W_x^{\mathbb{M}}(P, A, n) = \max_{I \in \mathcal{I}_n^P} \{Q_x^{\mathbb{M}}(E(A, I))\}$ as the worst-case cost of algorithm A on problem instances of size n for problem P on model \mathbb{M} with parameters x .

Analysis of both cache-oblivious and cache-aware algorithms relies on additional constraints defining the relationship between the cache parameters M and B . For example, the analysis of the cache-aware sorting algorithms [3] assumes $M \geq 2B$, while the analysis of the cache-oblivious sorting algorithms [8] and cache-oblivious sparse matrix dense vector (SpMV) multiplication [5] assume, respectively, $M \geq B^2$ and $M \geq B^{1+\epsilon}$ – the so-called *tall-cache assumption*. Therefore, let $\mathcal{M}_P(B)$ denote the set of all values of M (as a function of B) that satisfy such constraints for problem P .

To prove our results, we need to rigorously define algorithm optimality. Otherwise, as there are multiple parameters involved, the order of the quantifiers would be unclear and ambiguous.

DEFINITION 2.1. *A cache-oblivious algorithm A for problem P is asymptotically CO-optimal in the ideal-cache model with the cache parameters B and $M \in \mathcal{M}_P(B)$ iff¹:*

$$\exists c > 0, n_0 > 0 \forall n > n_0 \forall B \geq 1 \forall M \in \mathcal{M}_P(B) \forall A' \in \mathcal{A}_P [W_{M,B}^{\text{CO}}(P, A, n) \leq c \cdot W_{M,B}^{\text{CO}}(P, A', n)]$$

¹A standard and understandable reaction to this definition and the math to come is to sneer at the chain of quantifiers and declare “just write this using asymptotic notation!” Unfortunately, such notation is not flexible enough to say exactly what we need to say, and thus we make the quantification explicit.

To implement the cache replacement policy of the ideal-cache model requires the knowledge of what an algorithm will do in the future. Instead, modern hardware caches implement an approximation of the LRU-cache model, each time evicting the block that has been accessed least recently. While it is often easier to analyze the cache misses of an algorithm in the ideal-cache model, in this work we are able to work directly in the LRU-cache model.

DEFINITION 2.2. *A cache-oblivious algorithm A for problem P is asymptotically LRU-optimal in the LRU-cache model with the cache parameters B and $M \in \mathcal{M}_P(B)$ iff:*

$$\exists_{c>0, n_0>0} \forall_{n>n_0} \forall_{B \geq 1} \forall_{M \in \mathcal{M}_P(B)} \forall_{A' \in \mathcal{A}_P} \left[W_{M,B}^{\text{LRU}}(P, A, n) \leq c \cdot W_{M,B}^{\text{LRU}}(P, A', n) \right]$$

For completeness, however, we also prove our results in the ideal-cache model, by utilizing the following well-known resource-augmentation result of Sleator and Tarjan [11] (which also applies to other reasonable cache-replacement policies):

LEMMA 2.1. [11] *Any LRU-optimal algorithm in the LRU-cache model with cache parameters B and M is CO-optimal in the ideal-cache model with cache parameters B and $2M$.*

The equivalence in optimality between the ideal-cache and LRU-cache models relies on the cache augmentation, and says nothing about asymptotic equivalence for the same M . However, this is not an issue for a large class of natural problems, which can be solved using *memory-smooth* cache-oblivious algorithms.

DEFINITION 2.3. *A cache-oblivious algorithm A is memory-smooth iff increasing the memory size M by a constant factor does not asymptotically change its execution cost. That is,*

$$\forall_{I \in \mathcal{I}^P} \forall_{c>0} \forall_{B \geq 1} \forall_{M \in \mathcal{M}_P(B)} \left[Q_{M,B}^{\text{CO}}(E(A, I)) = \Theta(Q_{cM,B}^{\text{CO}}(E(A, I))) \right],$$

where the Θ -notation is with respect to the size of instance I .

Finally, we define algorithm optimality in our LoR model:

DEFINITION 2.4. *Let \mathcal{L} be a class of functions. Algorithm A for problem P is asymptotically LoR-optimal with respect to \mathcal{L} iff:*

$$\exists_{c>0, n_0>0} \forall_{n>n_0} \forall_{\ell \in \mathcal{L}} \forall_{A' \in \mathcal{A}_P} \left[W_{\ell}^{\text{LoR}}(P, A, n) \leq c \cdot W_{\ell}^{\text{LoR}}(P, A', n) \right]$$

The notion of an LoR-optimal algorithm with respect to *all* possible functions would be very powerful, as such an algorithm would be asymptotically optimal on any computing device that rewards locality of reference. In this paper we come very close to achieving such optimality, requiring only a natural set of restrictions on the functions in \mathcal{L} .

2.2 B-stable problems To show the equivalence between LoR-optimal and CO-optimal algorithms, we must avoid pathological problems with worst-case behavior that varies dramatically with different instances of the problem for different block sizes.

We say that a problem is *B-stable* if, for any algorithm A that solves P , there is some “worst-case” instance $I_w \in \mathcal{I}_n^P$ that, for every B , has CO cost asymptotically no less than the optimal worst-case cost for that B , over all instances. Formally,

DEFINITION 2.5. *Problem P is B-stable if, for any algorithm $A \in \mathcal{A}_P$ that solves P :*

$$\exists_{c>0, n_0>0} \forall_{n>n_0} \exists_{I_w \in \mathcal{I}_n^P} \forall_{B \geq 1} \forall_{M \in \mathcal{M}_P(B)} \min_{A' \in \mathcal{A}_P} W_{M,B}^{\text{CO}}(P, A', n) \leq c \cdot Q_{M,B}^{\text{CO}}(E(A, I_w))$$

Intuitively, for any algorithm that solves a *B-stable* problem, there must be a single instance that, for all block sizes, has cost no less than the asymptotically worst-case optimal cost.

The following lemma, implied by the definition of CO-optimality (Definition 2.1), states that every algorithm must have an instance on which it performs no better, asymptotically, than the CO-optimal algorithm, for every B :

LEMMA 2.2. *If an asymptotically CO-optimal algorithm A solves a B -stable problem P , then*

$$\exists c > 0, n_0 > 0 \forall n > n_0 \forall A' \in \mathcal{A}_P \exists I_w \in \mathcal{I}_n^P \forall B \geq 1 \forall M \in \mathcal{M}_P(B) \left[W_{M,B}^{\text{CO}}(P, A, n) \leq c \cdot Q_{M,B}^{\text{CO}}(E(A', I_w)) \right]$$

In the full version of the paper we prove the following lemma, which shows existence of non- B -stable problems, for which our main result (Theorem 3.1) does not hold. This justifies our classification and exclusion of these pathological cases.

LEMMA 2.3. *There exists a problem P which is not B -stable and which has a CO-optimal algorithm which is not LoR-optimal.*

2.3 B -smoothed analysis To prove our results, we will work with the *smoothed* version of the analysis. Given the access sequence $E = E(A, I)$ and an integer s chosen uniformly at random in the range $[0, B)$, let $E_{B\text{-smooth}}$ be the sequence derived from E where every element of E is increased by s . Then, the B -smoothed costs are defined as the expected cost on the B -smooth sequence in the respective cache models, i.e., $Q_{M,B}^{\text{CO}}(E) = \mathbb{E}[Q_{M,B}^{\text{CO}}(E_{B\text{-smooth}})]$ and $Q_{M,B}^{\text{LRU}}(E) = \mathbb{E}[Q_{M,B}^{\text{LRU}}(E_{B\text{-smooth}})]$.

LEMMA 2.4. *For any execution sequence E*

$$Q_{M,B}^{\text{SCO}}(E) \leq 2 \cdot Q_{M,B}^{\text{CO}}(E) \leq 4 \cdot Q_{M,B}^{\text{SCO}}(E),$$

and

$$Q_{M,B}^{\text{SLRU}}(E) \leq 2 \cdot Q_{M,B}^{\text{LRU}}(E) \leq 4 \cdot Q_{M,B}^{\text{SLRU}}(E).$$

Proof. Shifting the execution sequence may cause accesses that were in the same block to become in two neighboring blocks, and accesses that were in two neighboring blocks to become in the same block. Thus, the cost may grow or shrink by a factor of two, but not more. \square

3 Memoryless algorithms

We begin with the simplest case, namely, the *memoryless cache model (MCM)* where the internal memory is just a single block, i.e. $\mathcal{M}_P(B) = \{B\}$.² Note that in this case, there is no need to differentiate between LRU-cache model and ideal-cache model because the working sets in both cache models, after accessing $e_i \in E(A, I)$, consist of a single block that contains e_i . Thus, the costs of the algorithm A on instance I in the MCM becomes

$$\begin{aligned} Q_{B,B}^{\text{CO}}(E(A, I)) &= Q_{B,B}^{\text{LRU}}(E(A, I)) \\ &= \sum_{i=1}^{|E(A,I)|} \begin{cases} 1 & \text{if } \lfloor \frac{e_i}{B} \rfloor \neq \lfloor \frac{e_{i-1}}{B} \rfloor \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

This cost rewards spatial locality. Hence, in the LoR model it is natural to define the locality function ℓ to measure the cost of executing the sequence $E(A, I)$ as a function of the *spatial distance* $|e_i - e_{i-1}|$ between accesses:

$$Q_{\ell}^{\text{LoR}}(E(A, I)) = \sum_{i=1}^{|E(A,I)|} \ell(|e_i - e_{i-1}|)$$

Let \mathcal{L} denote the set of all non-negative, non-decreasing, concave functions $\ell : \mathbb{N} \rightarrow \mathbb{R}$. Even though \mathcal{L} encompasses a wide range of (arbitrarily complicated) functions, we will show that any cache-oblivious algorithm is LoR-optimal with respect to \mathcal{L} if and only if it is CO-optimal in the MCM.

²While this might seem overly restrictive, a number of query algorithms on cache-oblivious data structures are CO-optimal, because one is always traversing further into the structure, and never comes back to revisit parts of memory near where it has already gone.

3.1 Main result. Let $\ell_B(d) = \min(1, \frac{d}{B})$. We begin by proving our result for this *specific* locality function $\ell_B \in \mathcal{L}$ and generalize it to any $\ell \in \mathcal{L}$ later.

LEMMA 3.1. For any execution sequence E and block size B , $Q_{B,B}^{\text{SCO}}(E) = Q_{\ell_B}^{\text{LoR}}(E)$.

Proof. Consider a single access $e_i \in E$. Let $d = |e_i - e_{i-1}|$. If $d \geq B$ then $\ell_B(d) = 1$, and the i th term of $Q_{B,B}^{\text{CO}}(E)$ is also 1 because $\lfloor \frac{e_i+s}{B} \rfloor \neq \lfloor \frac{e_{i-1}+s}{B} \rfloor$ for all $0 \leq s < B$. If $d < B$ then $\ell_B(d) = \frac{d}{B}$, and the expected value of the i th term of $Q_{B,B}^{\text{CO}}(E)$ is also $\frac{d}{B}$, because $\lfloor \frac{e_i+s}{B} \rfloor \neq \lfloor \frac{e_{i-1}+s}{B} \rfloor$ for d of the possible shifts $s \in [0, B)$. \square

COROLLARY 3.1. If a cache-oblivious algorithm A for problem P is asymptotically LoR-optimal with respect to \mathcal{L} , then it is asymptotically CO-optimal in the MCM.

Proof. Since A is LoR-optimal, then it is within a constant factor of optimal for all locality functions, including $\ell_B(d) = \min(1, \frac{d}{B})$. The corollary follows from Lemmas 3.1 and 2.4. \square

We now show that *any* locality function $\ell \in \mathcal{L}$ can be represented as a linear combination of N functions $\ell_B(d)$, for N various values of B .

For every locality function $\ell \in \mathcal{L}$ there exist non-negative constants $\alpha_1, \alpha_2 \dots \alpha_N$ and $\beta_1, \beta_2 \dots \beta_N$ such that $\ell(d) = \sum_{i=1}^N \alpha_i \ell_{\beta_i}(d)$ for integers d in $[1..N]$.

Proof. [Proof (Sketch)] Let $\gamma_i = \begin{cases} 2\ell(i) - \ell(i+1) & \text{if } i = 1 \\ 2\ell(i) - \ell(i+1) - \ell(i-1) & \text{if } 2 \leq i \leq N-1 \\ \ell(i) - \ell(i-1) & \text{if } i = N \end{cases}$ and let $\alpha_i = i\gamma_i$ and $\beta_i = i$.

Then for any integer $x \in [1..N]$: $\sum_{i=1}^N \alpha_i \ell_{\beta_i}(x) = \sum_{i=1}^N \min(\alpha_i, \frac{\alpha_i x}{\beta_i}) = \sum_{i=1}^N \min(i\gamma_i, \gamma_i x) = \sum_{i=1}^{x-1} i\gamma_i + \gamma_x x + \sum_{i=x+1}^N \gamma_i x$ and this expression telescopes to $\ell(x)$ (the full derivation can be found in the full version of the paper).

Since function ℓ is non-negative and concave, all values of α_i and β_i are non-negative. \square

For every locality function $\ell \in \mathcal{L}$ there exists a sets of n non-negative constants $\alpha_1, \alpha_2 \dots \alpha_n$, and $\beta_1, \beta_2 \dots \beta_n$ such that, for any execution sequence E , $\sum_{i=1}^n \alpha_i Q_{\ell_{\beta_i}}^{\text{LoR}}(E) = Q_{\ell}^{\text{LoR}}(E)$.

THEOREM 3.1. Let \mathcal{L} be a set of all non-negative, non-decreasing, concave functions $l : \mathbb{N} \rightarrow \mathbb{R}$. Any cache-oblivious algorithm A that solves a B -stable problem P is LoR-optimal with respect to \mathcal{L} if and only if it is CO-optimal in the memoryless cache model.

Proof. The first direction follows from Corollary 3.1. To prove the other direction, consider the CO-optimal algorithm $A_{\text{CO-Opt}}$ and some algorithm A that solves P . Since P is B -stable, by Definition 2.5,

$$\exists_{c,n_0} \forall_{n > n_0} \exists_{I_w \in \mathcal{I}_n^P} \forall_{B \geq 1} \left[W_B(P, A_{\text{CO-Opt}}, n) \leq c \cdot Q_{B,B}^{\text{CO}}(E(A, I_w)) \right]$$

Using the definition of the worst-case cost W_B , we get

$$\exists_{c,n_0} \forall_{n > n_0} \exists_{I_w \in \mathcal{I}_n^P} \forall_{B \geq 1} \left[\max_{I \in \mathcal{I}_n^P} Q_{B,B}^{\text{CO}}(E(A_{\text{CO-Opt}}, I)) \leq c \cdot Q_{B,B}^{\text{CO}}(E(A, I_w)) \right]$$

By Lemma 2.4, we get

$$\exists_{c,n_0} \forall_{n > n_0} \exists_{I_w \in \mathcal{I}_n^P} \forall_{B \geq 1} \left[\max_{I \in \mathcal{I}_n^P} Q_{B,B}^{\text{SCO}}(E(A_{\text{CO-Opt}}, I)) \leq 4c \cdot Q_{B,B}^{\text{SCO}}(E(A, I_w)) \right]$$

and since, by Lemma 3.1, for any B the smoothed CO cost is equivalent to the LoR cost with the corresponding ℓ_B function,

$$\exists_{c,n_0} \forall_{n > n_0} \exists_{I_w \in \mathcal{I}_n^P} \forall_{B \geq 1} \left[\max_{I \in \mathcal{I}_n^P} Q_{\ell_B}^{\text{LoR}}(E(A_{\text{CO-Opt}}, I)) \leq 4c \cdot Q_{\ell_B}^{\text{LoR}}(E(A, I_w)) \right]$$

This inequality holds for all B and thus all linear combinations of various B . For any locality function ℓ in the set of valid locality functions, \mathcal{L} , consider $\alpha_1^\ell, \alpha_2^\ell, \dots, \alpha_n^\ell$ and $\beta_1^\ell, \beta_2^\ell, \dots, \beta_n^\ell$ given by Lemma 3.1. We use the β 's as the B values and the α 's as the coefficients in the linear combination to get

$$\exists_{c', n_0} \forall_{n > n_0} \exists_{I_w \in \mathcal{I}_n^P} \forall_{\ell \in \mathcal{L}} \sum_{k=1}^n \max_{I \in \mathcal{I}_n^P} \left(\alpha_k^\ell Q_{\ell_{\beta_k^\ell}}^{\text{LoR}}(E(\text{ACO-Opt}, I)) \right) \leq c' \cdot \sum_{k=1}^n \left(\alpha_k^\ell Q_{\ell_{\beta_k^\ell}}^{\text{LoR}}(E(A, I_w)) \right)$$

I_w is a single instance of P , therefore it cannot have a greater total cost than the single instance that maximizes the cost

$$\exists_{c', n_0} \forall_{n > n_0} \forall_{\ell \in \mathcal{L}} \sum_{k=1}^n \max_{I \in \mathcal{I}_n^P} \left(\alpha_k^\ell Q_{\ell_{\beta_k^\ell}}^{\text{LoR}}(E(\text{ACO-Opt}, I)) \right) \leq c' \cdot \max_{I \in \mathcal{I}_n^P} \sum_{k=1}^n \left(\alpha_k^\ell Q_{\ell_{\beta_k^\ell}}^{\text{LoR}}(E(A, I)) \right)$$

Moving the max outside the summation can only decrease the overall cost of the left side of the inequality, thus

$$\exists_{c', n_0} \forall_{n > n_0} \forall_{\ell \in \mathcal{L}} \max_{I \in \mathcal{I}_n^P} \sum_{k=1}^n \left(\alpha_k^\ell Q_{\ell_{\beta_k^\ell}}^{\text{LoR}}(E(\text{ACO-Opt}, I)) \right) \leq c' \cdot \max_{I \in \mathcal{I}_n^P} \sum_{k=1}^n \left(\alpha_k^\ell Q_{\ell_{\beta_k^\ell}}^{\text{LoR}}(E(A, I)) \right)$$

Using Corollary 3.1, we get

$$\exists_{c', n_0} \forall_{n > n_0} \forall_{\ell \in \mathcal{L}} \max_{I \in \mathcal{I}_n^P} (Q_\ell^{\text{LoR}}(E(\text{ACO-Opt}, I))) \leq c' \cdot \max_{I \in \mathcal{I}_n^P} (Q_\ell^{\text{LoR}}(E(A, I)))$$

We are considering an arbitrary algorithm A that solves P , so this applies to all $A \in \mathcal{A}_P$. By the definition of worst-case LoR cost,

$$\exists_{c', n_0} \forall_{n > n_0} \forall_{\ell \in \mathcal{L}} \forall_{A \in \mathcal{A}_P} \left[W_\ell(P, \text{ACO-Opt}, n) \leq c' \cdot W_\ell(P, A, n) \right]$$

Thus, by Definition 2.4, ACO-Opt is asymptotically LoR-optimal. \square

4 General models for algorithms with memory

In the previous section, we considered execution sequences that did not utilize more than one block of memory, and thus locality to other than the previously accessed memory location was irrelevant. Now we generalize the model and apply it to execution sequences without this restriction. This requires that we consider the size and contents of internal memory when computing the expected cost of an access.

4.1 General LoR model To capture the concept of the working set for algorithms that use internal memory, we define bidimensional locality functions that compute LoR cost based on two dimensions: distance and time. This bidimensional locality function, $\ell(d, \delta)$ represents the cost of a jump from a *source* element, s , to a *target* element, t , where d and δ are the *spatial distance* and *temporal distance*, respectively, between s and t . This captures the concept of the working set by using “time” to determine if the source element is in memory or not. If the source is temporally close (was accessed recently) and spatially close to the target, t , the resulting locality cost of the jump is small.

Details of the bidimensional locality functions. Let \mathcal{L}_B denote the set of *bidimensional locality functions* that we consider. The functions in this set have the following properties. An element of this set is a function of the form $\ell(d, \delta) = \max(f(d), g(\delta))$, where $f(d)$ is non-negative, non-decreasing, and concave, while $g(\delta)$ is a 0-1 threshold function, i.e.,

$$g(\delta) = g_x(\delta) = \begin{cases} 0 & \text{if } 0 \leq \delta \leq x \\ 1 & \text{otherwise} \end{cases}$$

for some value x . For any $k \leq i$, the bidimensional locality cost of a jump from source element e_k to target e_i in the sequence E is $\ell(|e_i - e_k|, t(E, i) - t(E, k))$, where $t(E, i)$ is the *time* of the i -th access. For simplicity of notation, we define $\delta_t(E, k, i)$ to be the temporal distance between the i -th access (e_i) and k -th access (e_k), i.e., $\delta_t(E, k, i) = t(E, i) - t(E, k)$. Intuitively, we can think of $\delta_t(E, k, i)$ as the time from access e_k to “present” when accessing e_i . In addition, we require that the functions cannot be more “sensitive” to temporal locality than spatial

locality, i.e., for any locality function $\ell(d, \delta) = \max(f(d), g(\delta))$, we have that $\forall_x [f(x) \geq g(x)]$. This corresponds to the *tall cache assumption* $M \geq B^2$, which is typically used in the analysis of cache-oblivious algorithms [6, 10]. Therefore, we restrict the machine parameters, M and B , to all values $B \geq 1$ and $\mathcal{M}_P(B) = \{M : M \geq B^2\}$. A more in-depth discussion of the tall cache assumption and how it relates to the LoR model can be found in the full version of the paper.

We form our definition of time based on the amount of change that occurs to the working set. For example, if an access causes a block of B elements to be evicted, we say that time increases by 1. Thus, time depends on the locality function, so we define the time of the i -th access of E , for the given locality function ℓ to be $t_\ell(E, i) = \sum_{k=1}^{i-1} (Q_\ell^{\text{LoR}}(E, k))$. That is, the time of access $e_i \in E$ is simply the sum of costs of all accesses prior to e_i in sequence E . We note that the time after the last access of E is the total LoR cost (i.e., $Q_\ell^{\text{LoR}}(E) = t_\ell(E, |E| + 1)$).

Unlike the memoryless LoR cost, we cannot simply compute the cost of access e_i using the distance from the previous access, e_{i-1} , since any of the prior accesses may be in the working set when accessing e_i . Furthermore, since we no longer consider only non-decreasing execution sequences, when accessing e_i , there may be accesses to both the *left* and *right* that could be in the same block as e_i . Therefore, computing $Q_\ell^{\text{LoR}}(E, i)$ using the locality function from a single source is insufficient to capture the idea of the working set, and a detailed example showing why this is the case is included in the full version of the paper. We define the general LoR cost of access $e_i \in E$ as $Q_\ell^{\text{LoR}}(E, i) =$

$$\max \left(\left(\begin{array}{c} \text{left side} \\ \min_{\substack{\forall L < i \text{ s.t.} \\ e_L \leq e_i}} \ell(e_i - e_L, \delta_t(E, i, L)) \\ + \\ \min_{\substack{\forall R < i \text{ s.t.} \\ e_R \geq e_i}} \ell(e_R - e_i, \delta_t(E, i, R)) \\ \text{right side} \end{array} \right) - 1, 0 \right).$$

Intuitively, the LoR cost of access $e_i \in E$ is computed from the minimum cost jumps from both the left side and right side of e_i . We note that this generalizes the LoR cost definition of the memoryless setting (Section 3), as the locality function from source e_R will always evaluate to 1 for non-decreasing accesses.

This formulation has the added benefit that it lets us easily visualize an execution sequence in a graphical representation, illustrated in Figure 2. We consider a series of accesses in execution sequence E as points in a 2-dimensional plane. The point representing access e_i is plotted with the x and y coordinates corresponding to the spatial position, e_i , and the temporal position, $t(E, i)$, respectively. The cost of access e_i is simply computed from the LoR cost with sources e_L and e_R (the previous access with the minimum locality function cost to the left and right, respectively). We can visually determine which previous accesses correspond to e_L and e_R : if a previous access is outside the gray region (i.e., $\delta > \frac{M}{B}$ or $d > B$), the cost is 1. Otherwise, it is simply $\frac{d}{B}$.

4.2 Equivalence to cache-oblivious cost As with the memoryless LoR model, we first prove our result for a specific bidimensional locality function, $\ell_{M,B}(d, \delta) = \max\left(\min\left(1, \frac{d}{B}\right), \min\left(1, \left\lfloor \frac{\delta}{M/B} \right\rfloor\right)\right)$, and later generalize our result to all bidimensional locality function in \mathcal{L}_B .

For any cache-oblivious algorithm A , $B \geq 1$, $M \geq B^2$, execution sequence E generated by A , and $\ell_{M,B}(d, \delta) = \max\left(\min\left(1, \frac{d}{B}\right), \min\left(1, \left\lfloor \frac{\delta}{M/B} \right\rfloor\right)\right)$: $Q_{\ell_{M,B}}^{\text{LoR}}(E) = Q_{M,B}^{\text{SLRU}}(E)$ and, consequently, $Q_{M,B}^{\text{SCO}}(E) \leq Q_{\ell_{M,B}}^{\text{LoR}}(E) \leq 2 \cdot Q_{\frac{M}{2}, B}^{\text{SCO}}(E)$.

Proof. To prove that $Q_{\ell_{M,B}}^{\text{LoR}}(E) = Q_{M,B}^{\text{SLRU}}(E)$, we consider the cost of performing access $e_i \in E$. Assume that, when accessing e_i , e_L is the nearest element to the *left* of e_i ($e_L \leq e_i$) that is in the working set, i.e., $L < i$, $e_L \in \mathcal{W}_{M,B}^{\text{LRU}}(E, i-1)$, and $e_i - e_L$ is minimized. If there is no such element to the left of e_i in the working set, then we say that $e_L = -\infty$. Similarly, assume that e_R is the nearest element in $\mathcal{W}_{M,B}^{\text{LRU}}(E, i-1)$ to the *right* of e_i ($e_i \leq e_R$), and if there is no such element, then $e_R = +\infty$. By this definition, for any access e_i , we can simply consider the spatial components, because, if no element is within temporal distance $\frac{M}{B}$, the spatial distance is ∞ and the $\ell_{M,B}$ cost is 1.

We consider three possible cases for the spatial distance of e_L and e_R from access e_i :

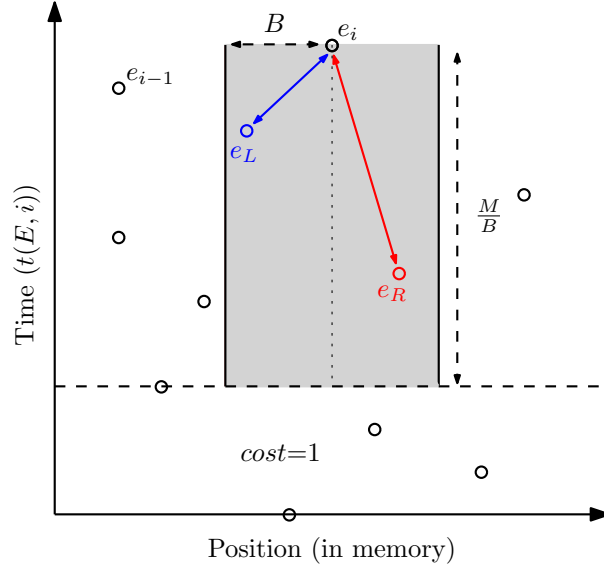


Figure 2: Graphical visualization of accesses in the general LoR model with locality function $\ell_{M,B}$. At time of access e_i , all prior accesses within the past time $\frac{M}{B}$ (above the dashed horizontal line) have $g(\delta) = 0$. The locality cost to jump from any element outside the gray region has the maximum cost of 1. In this example, there is both an e_L and e_R with cost < 1 .

Case 1: $(e_i - e_L) \geq B$ AND $(e_R - e_i) \geq B$

There is no element in the working set within distance B of e_i , then, for all alignment shifts, $0 \leq s < B$, we know that $e_i \notin \mathcal{W}_{M,B}^{\text{LRU}}(E, i-1)$. Thus,

$$\begin{aligned} Q_{M,B}^{\text{SLRU}}(E, i) &= \mathbb{E}[Q_{M,B}^{\text{CO}}(E_{B\text{-smooth}}, i)] \\ &= \frac{1}{B} \sum_{s=0}^{B-1} 1 \\ &= 1 \end{aligned}$$

and the LoR model cost is

$$\begin{aligned} Q_{\ell_{M,B}}^{\text{LoR}}(E, i) &= \max(\ell_{M,B}(e_i - e_L, \delta_t(E, i, L)) \\ &\quad + \ell_{M,B}(e_R - e_i, \delta_t(E, i, R)) - 1, 0) \\ &= \max(1 + 1 - 1, 0) = 1 \end{aligned}$$

We note that this includes the cases where e_L and/or e_R do not exist, since, we set $e_L = -\infty$ and/or $e_R = \infty$, respectively, in such cases.

Case 2: $(e_i - e_L) < B$ OR $(e_R - e_i) < B$

Only one side (left or right) is within distance B of e_i . W.l.o.g, assume that $(e_R - e_i) < B$ and $(e_i - e_L) \geq B$. Since $(e_i - e_L) \geq B$, for all shifts $0 \leq s < B-1$, we know that $e_L \notin \mathcal{W}_{M,B}^{\text{LRU}}(E, i-1)$. Thus, the smoothed LRU cost is simply

$$\begin{aligned} Q_{M,B}^{\text{SLRU}}(E, i) &= \frac{1}{B} \sum_{s=0}^{B-1} \begin{cases} 1 & \text{if } \lfloor \frac{e_R+s}{B} \rfloor \neq \lfloor \frac{e_i+s}{B} \rfloor \\ 0 & \text{otherwise} \end{cases} \\ &= \frac{e_R - e_i}{B} \end{aligned}$$

The LoR cost is

$$\begin{aligned}
Q_{\ell_M, B}^{\text{LoR}}(E, i) &= \max \left(\min \left(1, \frac{e_i - e_L}{B} \right) \right. \\
&\quad \left. + \min \left(1, \frac{e_R - e_i}{B} \right) - 1, 0 \right) \\
&= \max \left(1 + \frac{(e_R - e_i)}{B} - 1, 0 \right) \\
&= \frac{(e_R - e_i)}{B}
\end{aligned}$$

A symmetric argument holds in the case where $(e_i - e_L) < B$ and $(e_R - e_i) \geq B$.

Case 3: $(e_i - e_L) < B$ AND $(e_R - e_i) < B$

Both e_L and e_R are within distance B of e_i , so the smoothed LRU cost depends on the number of alignment shifts, s , for which e_i is not in the same block as either e_L or e_R , i.e.,

$$Q_{M, B}^{\text{SLRU}}(E, i) = \frac{1}{B} \sum_{s=0}^{B-1} \begin{cases} 1 & \text{if } \lfloor \frac{e_i+s}{B} \rfloor \neq \lfloor \frac{e_R+s}{B} \rfloor \\ & \text{and } \lfloor \frac{e_i+s}{B} \rfloor \neq \lfloor \frac{e_L+s}{B} \rfloor \\ 0 & \text{otherwise} \end{cases}$$

For simplicity, assume that at alignment shift $s = 0$, e_i is in the last location of the block of size B . Thus, the shifts from $s = 0$ to $s = (B - 1)$ define a $2B$ range around e_i (i.e., $[e_i - B, e_i + B]$). We define $p(e_L)$ and $p(e_R)$ to be the indexes of e_L and e_R in this $2B$ range, respectively. For all $0 \leq s \leq p(e_L)$, e_i is in the same block as e_L . Similarly, for all $(p(e_R) - B) \leq s < B$, e_i is in the same block as e_R . Thus, the cost is simply the number of shifts, s , where the entire block of size B containing e_i is strictly between $p(e_L)$ and $p(e_R)$, i.e.,

$$\begin{aligned}
Q_{M, B}^{\text{SLRU}}(E, i) &= \frac{1}{B} \sum_{s=p(e_L)}^{p(e_R)-B} 1 \\
&= \frac{p(e_R) - B - p(e_L)}{B}
\end{aligned}$$

and, since the cost cannot be negative, this becomes

$$= \max \left(\frac{p(e_R) - p(e_L)}{B} - 1, 0 \right)$$

We know that $p(e_R) = B + e_R - e_i$ and $p(e_L) = B - (e_i - e_L)$, thus

$$\begin{aligned}
&= \max \left(\frac{B + e_R - e_i}{B} \right. \\
&\quad \left. - \frac{B - (e_i - e_L)}{B} - 1, 0 \right) \\
&= \max \left(\frac{e_R - e_i}{B} + \frac{e_i - e_L}{B} - 1, 0 \right)
\end{aligned}$$

Since both e_L and e_R are within distance B of e_i , this is equal to LoR cost, i.e.,

$$\begin{aligned}
Q_{\ell_M, B}^{\text{LoR}}(E, i) &= \max \left(\min \left(1, \frac{e_i - e_L}{B} \right) \right. \\
&\quad \left. + \min \left(1, \frac{e_R - e_i}{B} \right) - 1, 0 \right) \\
&= \max \left(\frac{e_i - e_L}{B} + \frac{e_R - e_i}{B} - 1, 0 \right)
\end{aligned}$$

Thus, for any access access, $e_i \in E$,

$$Q_{\ell_{M,B}}^{\text{LoR}}(E, i) = Q_{M,B}^{\text{SLRU}}(E, i)$$

where

$$\ell_{M,B}(d, \delta) = \max\left(\min\left(1, \frac{d}{B}\right), \min\left(1, \left\lfloor \frac{\delta}{M/B} \right\rfloor\right)\right).$$

Since they are equivalent for any access, $e_i \in E$, then for any execution sequence E ,

$$Q_{\ell_{M,B}}^{\text{LoR}}(E) = Q_{M,B}^{\text{SLRU}}(E)$$

Since the cache-oblivious cost is computed assuming ideal cache replacement, and LRU cache replacement with twice the memory is 2-competitive with ideal cache, we have

$$\begin{aligned} Q_{M,B}^{\text{SCO}}(E) &\leq Q_{M,B}^{\text{SLRU}}(E) \\ &\leq Q_{\ell_{M,B}}^{\text{LoR}}(E) \\ &\leq Q_{M,B}^{\text{SLRU}}(E) \\ &\leq 2 \cdot Q_{\frac{M}{2},B}^{\text{SCO}}(E) \end{aligned}$$

□

We can also prove similar asymptotic equivalence result between the LoR and the ideal-cache models for the same M , if we consider memory-smooth algorithms:

LEMMA 4.1. *For any memory-smooth cache-oblivious algorithm A , $B \geq 1$, $M \geq B^2$, execution sequence E generated by A , and $\ell_{M,B}(d, \delta) = \max\left(\min\left(1, \frac{d}{B}\right), \min\left(1, \left\lfloor \frac{\delta}{M/B} \right\rfloor\right)\right)$:*

$$Q_{\ell_{M,B}}^{\text{LoR}}(E) = \Theta(Q_{M,B}^{\text{SCO}}(E)).$$

Proof. If A is memory-smooth, then $Q_{M,B}^{\text{SCO}}(E) = \Theta(Q_{\frac{M}{2},B}^{\text{SCO}}(E))$, and, by Theorem 4.2, $Q_{\ell_{M,B}}^{\text{LoR}}(E) = \Theta(Q_{M,B}^{\text{SCO}}(E))$.
□

4.3 Main result We now extend our result to any bidimensional locality function $\ell \in \mathcal{L}_B$. A cache-oblivious algorithm A for a B -stable problem P is LRU-optimal if and only if it is LoR-optimal with respect to \mathcal{L}_B , where \mathcal{L}_B is a set of all functions of the form $\ell(d, \delta) = \max(f(d), g(\delta))$, where $g(\delta)$ is a 0-1 threshold function, $f(x) \geq g(x)$ for all $x \geq 0$, and f is a non-negative, non-decreasing, concave function.

Proof. If algorithm A_{LoR} is LoR-optimal for all bidimensional locality functions, then it is optimal for locality functions $\ell_{M,B}$, for any M and B . By Theorem 4.2, it follows that A_{LoR} is LRU-optimal for any M and B .

To prove that LRU-optimal algorithms are also LoR-optimal, we consider problem P and algorithm A_{LRU} that solves P with optimal LRU cost, i.e.,

$$\exists_{c,n_0} \forall_{n > n_0} \forall_{B \geq 1} \forall_{M \geq B^2} \forall_{A \in \mathcal{A}_P} \quad W_{M,B}^{\text{LRU}}(P, A_{\text{LRU}}, n) \leq c \cdot W_{M,B}^{\text{LRU}}(P, A, n)$$

And by the definition of the worst-case cost W ,

$$\exists_{c,n_0} \forall_{n > n_0} \forall_{B \geq 1} \forall_{M \geq B^2} \forall_{A \in \mathcal{A}_P} \quad \max_{I \in \mathcal{I}_n^P} (Q_{M,B}^{\text{LRU}}(E(A_{\text{LRU}}, I))) \leq c \cdot \max_{I \in \mathcal{I}_n^P} (Q_{M,B}^{\text{LRU}}(E(A, I)))$$

Since P is B -stable, there is some instance $I_w \in \mathcal{I}_n^P$ for each A such that

$$\exists_{c,n_0} \forall_{n > n_0} \forall_{B \geq 1} \forall_{M \geq B^2} \forall_{A \in \mathcal{A}_P} \quad \max_{I \in \mathcal{I}_n^P} (Q_{M,B}^{\text{LRU}}(E(A_{\text{LRU}}, I))) \leq c \cdot Q_{M,B}^{\text{LRU}}(E(A, I_w))$$

and by Lemma 2.4,

$$\exists_{c,n_0} \forall_{n>n_0} \forall_{B \geq 1} \forall_{M \geq B^2} \forall_{A \in \mathcal{A}_P} \frac{1}{2} \max_{I \in \mathcal{I}_n^P} (Q_{M,B}^{\text{SLRU}}(E(A_{\text{LRU}}, I))) \leq 2c \cdot Q_{M,B}^{\text{SLRU}}(E(A, I_w))$$

therefore, by Lemma 4.2

$$\exists_{c,n_0} \forall_{n>n_0} \forall_{B \geq 1} \forall_{M \geq B^2} \forall_{A \in \mathcal{A}_P} \max_{I \in \mathcal{I}_n^P} (Q_{\ell_{M,B}}^{\text{LoR}}(E(A_{\text{LRU}}, I))) \leq 4c \cdot Q_{\ell_{M,B}}^{\text{LoR}}(E(A, I_w))$$

Since this inequality holds for all $\ell_{M,B}$ functions, we define a series of such functions that we use to represent any bidimensional locality function. Recall that $\ell_{M,B}$ functions are of the form

$$\ell_{M,B}(d, \delta) = \max \left(\min \left(1, \frac{d}{B} \right), \min \left(1, \left\lfloor \frac{\delta}{M/B} \right\rfloor \right) \right)$$

where $B \geq 1$ and $M \geq B^2$. Consider an arbitrary bidimensional locality function $\ell(d, \delta) = \max(f(d), g(\delta))$. By Lemma 3.1, we can represent the $f(d)$ component by a linear combination of n memoryless ℓ_B functions (and therefore using the spatial component of $\ell_{M,B}$ functions). By our definition of bidimensional locality functions, $g(\delta) = \lfloor \frac{\delta}{x} \rfloor$, for some integer x . Thus, we simply set the temporal component of every one of our $\ell_{M,B}$ functions to be $g(\delta)$. For a given bidimensional locality function ℓ , we define ℓ_k^ℓ to be the k -th such $\ell_{M,B}$ function that we use to represent it, i.e., $\ell_k^\ell = \max(\alpha_k^\ell \ell_{\beta_k^\ell}, g(\delta))$.³ Thus, we have

$$\exists_{c',n_0} \forall_{n>n_0} \forall_{A \in \mathcal{A}_P} \forall_{\ell \in \mathcal{L}_B} \sum_{k=1}^n \max_{I \in \mathcal{I}_n^P} (Q_{\ell_k^\ell}^{\text{LoR}}(E(A_{\text{LRU}}, I))) \leq c' \sum_{k=1}^n (Q_{\ell_k^\ell}^{\text{LoR}}(E(A, I_w)))$$

Instance I_w cannot result in greater cost than the instance that maximizes the total cost, so

$$\exists_{c',n_0} \forall_{n>n_0} \forall_{A \in \mathcal{A}_P} \forall_{\ell \in \mathcal{L}_B} \sum_{k=1}^n \max_{I \in \mathcal{I}_n^P} (Q_{\ell_k^\ell}^{\text{LoR}}(E(A_{\text{LRU}}, I))) \leq c' \max_{I \in \mathcal{I}_n^P} \sum_{k=1}^n (Q_{\ell_k^\ell}^{\text{LoR}}(E(A, I)))$$

Moving the max outside of the summation can only decrease the cost of the left hand side of the inequality, thus

$$\exists_{c',n_0} \forall_{n>n_0} \forall_{A \in \mathcal{A}_P} \forall_{\ell \in \mathcal{L}_B} \max_{I \in \mathcal{I}_n^P} \sum_{k=1}^n (Q_{\ell_k^\ell}^{\text{LoR}}(E(A_{\text{LRU}}, I))) \leq c' \max_{I \in \mathcal{I}_n^P} \sum_{k=1}^n (Q_{\ell_k^\ell}^{\text{LoR}}(E(A, I)))$$

The proof of Corollary 3.1 applies, giving us

$$\exists_{c'',n_0} \forall_{n>n_0} \forall_{A \in \mathcal{A}_P} \forall_{\ell \in \mathcal{L}_B} \max_{I \in \mathcal{I}_n^P} (Q_{\ell}^{\text{LoR}}(E(A_{\text{LRU}}, I))) \leq c'' \cdot \max_{I \in \mathcal{I}_n^P} (Q_{\ell}^{\text{LoR}}(E(A, I)))$$

Using our definition of the worst-case LoR cost,

$$\exists_{c'',n_0} \forall_{n>n_0} \forall_{A \in \mathcal{A}_P} \forall_{\ell \in \mathcal{L}_B} W_{\ell}(P, A_{\text{LRU}}, n) \leq c'' \cdot W_{\ell}(P, A, n)$$

Therefore, any LRU-optimal algorithm is also LoR-optimal. \square

A memory-smooth cache-oblivious algorithm A for a B -stable problem P is CO-optimal if and only if A is LoR-optimal with respect to \mathcal{L}_B , where \mathcal{L}_B is a set of all functions of the form $\ell(d, \delta) = \max(f(d), g(\delta))$, where $g(\delta)$ is a 0-1 threshold function, $f(x) \geq g(x)$ for all $x \geq 0$, and f is a non-negative, non-decreasing, concave function.

Proof. Since the cache-oblivious model assumes ideal cache replacement, for any execution sequence E , $Q_{M,B}^{\text{SCO}}(E) \leq Q_{M,B}^{\text{SLRU}}(E) \leq 2 \cdot Q_{\frac{M}{2},B}^{\text{SCO}}(E)$. Since algorithm A is memory-smooth, for any execution sequence E generated by A , $Q_{\frac{M}{2},B}^{\text{SCO}}(E) = \Theta(Q_{M,B}^{\text{SCO}}(E))$. Therefore, $Q_{M,B}^{\text{SLRU}}(E) = \Theta(Q_{M,B}^{\text{SCO}}(E))$. Since the LRU cost and CO cost are asymptotically equivalent for every execution sequence generated by A , then A is asymptotically LRU-optimal if and only if it is asymptotically CO-optimal and, by Theorem 4.3, A is LoR-optimal if and only if it is CO-optimal. \square

³We note that, because we are limited to $M \geq B^2$ for our $\ell_{M,B}$ functions, we can only construct functions where $f(k) \geq g(k)$, for all $k > 0$. However, our definition of bidimensional locality functions includes this restriction, as it corresponds to the tall cache assumption (discussed in Section 4.1).

5 Conclusion

Despite the increasing complexity of modern hardware architectures, the goal of many design and optimization principles remain the same: maximize locality of reference. Even many of the optimization techniques used by modern compilers, such as branch prediction or loop unrolling [12], can be seen as methods of increasing spatial and/or temporal locality. As we demonstrated in this work, cache-oblivious algorithms do just that, suggesting that the performance benefits of such algorithms extend beyond what was originally envisioned.

That is to say, although we have introduced a new way to model computation via locality functions, we are not advocating algorithm design and analysis using locality functions. Instead, though our transformations, we have shown that creating the best possible algorithms in the existing cache-oblivious models is the right way to design algorithms not just for a multi-level cache, but for any locality-of-reference-rewarding system. One can thus conclude that the cache-oblivious model is better than we thought it was.

References

- [1] Alok Aggarwal, Bowen Alpern, Ashok K. Chandra, and Marc Snir. A model for hierarchical memory. In Alfred V. Aho, editor, *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 305–314. ACM, 1987. URL: <https://doi.org/10.1145/28395.28428>, doi:10.1145/28395.28428.
- [2] Alok Aggarwal, Ashok K. Chandra, and Marc Snir. Hierarchical memory with block transfer. In *28th Annual Symposium on Foundations of Computer Science, Los Angeles, California, USA, 27-29 October 1987*, pages 204–216. IEEE Computer Society, 1987. URL: <https://doi.org/10.1109/SFCS.1987.31>, doi:10.1109/SFCS.1987.31.
- [3] Alok Aggarwal and Jeffrey Scott Vitter. The input/output complexity of sorting and related problems. *Commun. ACM*, 31(9):1116–1127, 1988. URL: <http://doi.acm.org/10.1145/48529.48535>, doi:10.1145/48529.48535.
- [4] Bowen Alpern, Larry Carter, and Ephraim Feig. Uniform memory hierarchies. In *31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume II*, pages 600–608. IEEE Computer Society, 1990. URL: <https://doi.org/10.1109/SFCS.1990.89581>, doi:10.1109/SFCS.1990.89581.
- [5] Michael A. Bender, Gerth Stølting Brodal, Rolf Fagerberg, Riko Jacob, and Elias Vicari. Optimal sparse matrix dense vector multiplication in the i/o-model. *Theory Comput. Syst.*, 47(4):934–962, 2010. URL: <https://doi.org/10.1007/s00224-010-9285-4>, doi:10.1007/s00224-010-9285-4.
- [6] Gerth Stølting Brodal and Rolf Fagerberg. On the limits of cache-obliviousness. In *STOC*, pages 307–315. ACM, 2003.
- [7] Peter J. Denning. The locality principle. In *Communication Networks And Computer Systems: A Tribute to Professor Erol Gelenbe*, pages 43–67. 2006.
- [8] Matteo Frigo, Charles E. Leiserson, Harald Prokop, and Sridhar Ramachandran. Cache-oblivious algorithms. *ACM Trans. Algorithms*, 8(1):4:1–4:22, 2012.
- [9] Tomasz Jurkiewicz and Kurt Mehlhorn. The cost of address translation. In *Proceedings of the Meeting on Algorithm Engineering and Experiments (ALENEX)*, pages 148–162, 2013.
- [10] Francesco Silvestri. On the limits of cache-oblivious rational permutations. *Theor. Comput. Sci.*, 402(2-3):221–233, 2008.
- [11] Daniel Dominic Sleator and Robert Endre Tarjan. Amortized efficiency of list update and paging rules. *Commun. ACM*, 28(2):202–208, 1985. URL: <http://doi.acm.org/10.1145/2786.2793>, doi:10.1145/2786.2793.
- [12] William Stallings. *Computer Organization and Architecture - Designing for Performance (7. ed.)*. Pearson / Prentice Hall, 2006.
- [13] Leslie G. Valiant. A bridging model for multi-core computing. *J. Comput. Syst. Sci.*, 77(1):154–166, 2011. URL: <https://doi.org/10.1016/j.jcss.2010.06.012>, doi:10.1016/j.jcss.2010.06.012.