Cache-efficient Dynamic Programming Algorithms for Multicores *

Rezaul Alam Chowdhury Vijaya Ramachandran

UTCS Technical Report TR-08-16

April 11, 2008

Abstract

We present cache-efficient chip multiprocessor (CMP) algorithms with good speed-up for some widely used dynamic programming algorithms. We consider three types of caching systems for CMPs: *D-CMP* with a private cache for each core, *S-CMP* with a single cache shared by all cores, and *Multicore*, which has private L_1 caches and a shared L_2 cache. We derive results for three classes of problems: local dependency dynamic programming (LDDP), Gaussian Elimination Paradigm (GEP), and parenthesis problem.

For each class of problems, we develop a generic CMP algorithm with an associated tiling sequence. We then tailor this tiling sequence to each caching model and provide a parallel schedule that results in a cache-efficient parallel execution up to the critical path length of the underlying dynamic programming algorithm.

We present experimental results on an 8-core Opteron for two sequence alignment problems that are important examples of LDDP. Our experimental results show good speed-ups for simple versions of our algorithms.

*Department of Computer Sciences, University of Texas, Austin, TX 78712. Email: {shaikat,vlr}@cs.utexas.edu. This work was supported in part by NSF Grant CCF-0514876 and NSF CISE Research Infrastructure Grant EIA-0303609.

1 Introduction

Chip multiprocessors (CMP) are rapidly becoming the dominant parallel computing platform. A key feature that sets CMPs apart from earlier parallel machines is the organization and cost measure for memory accesses. In a CMP, memory is organized in a hierarchy of caches, and two new features come with this:

- The cost of a memory access is determined by whether the accessed data item is a cache hit or a cache miss. This is in contrast to using the routing cost, or latency and gap parameters, to measure communication cost as in, e.g., [33, 17, 12].
- A cache has a small finite size, hence a data item brought into cache and needed at a later time may not reside there at that later time if it has been evicted to make space for other data items. This is a feature that has not been captured in any of the traditional parallel models and algorithms.

In this paper we consider some important dynamic programming (DP) applications, and we develop good CMP implementations under three types of caching models for CMP (private, shared, and multicore). We elaborate on these three caching models in Section 1.1. For all three CMP models, we develop parallelizations of dynamic programming algorithms that match the best sequential caching complexity while achieving the maximum speed-up available in the algorithm. We obtain results for the following classes of problems:

1. *LDDP (Local Dependency DP) problems,* where each update to an entry in the DP table is determined by contents of neighboring cells. The simplest example of this class is the well-known *LCS* (longest common subsequence) problem. Some widely-used string problems in bioinformatics including *PA* (pairwise sequence alignment with affine gap cost) and *Median* (3-way sequence alignment with affine gap) are examples of LDDP. We present experimental results for PA and Median that show simplified versions of our algorithms to run several times faster than currently used software for these problems even when using the default Linux scheduler for our parallel code.
2. *GEP (Gaussian Elimination Paradigm)*, which solves another wide class of problems including all-pairs shortest paths in weighted graphs (Floyd-Warshall DP), LU decomposition and Gaussian elimination without pivoting, and matrix multiplication. Here we improve the p-processor cache-efficient parallelism from $\Theta(\frac{n^3}{p} + n\log^2 n)$ in the earlier CMP algorithms in [10, 2] to $\Theta\left(\frac{n^3}{p} + n\right)$[1] on an $n \times n$ input, which is the best possible when staying within the GEP framework.
3. *Parenthesis problem*, which includes RNA secondary structure prediction, optimal matrix chain multiplication, construction of optimal binary search trees, and optimal polynomial triangulation.
4. *RNA-SP (RNA secondary structure prediction with simple pseudo-knots).* Combining our results for 3 dimensional LDDP and GEP it is straightforward to obtain results for RNA secondary structure prediction with simple pseudo-knots, using the sequential RNA-SP algorithm in [8]. We do not elaborate on this further.

Given a parallel algorithm, let $T_\infty(n)$ denote the number of parallel steps in the algorithm as a function of the input size n. This is normally referred to as the critical path length of the computation, and represents its intrinsic parallelism. We use the term *cache-efficient parallelism* or *cache-efficient*

[1] In other words, we improve the critical pathlength of the computation from $\Theta\left(n\log^2 n\right)$ to $\Theta(n)$.

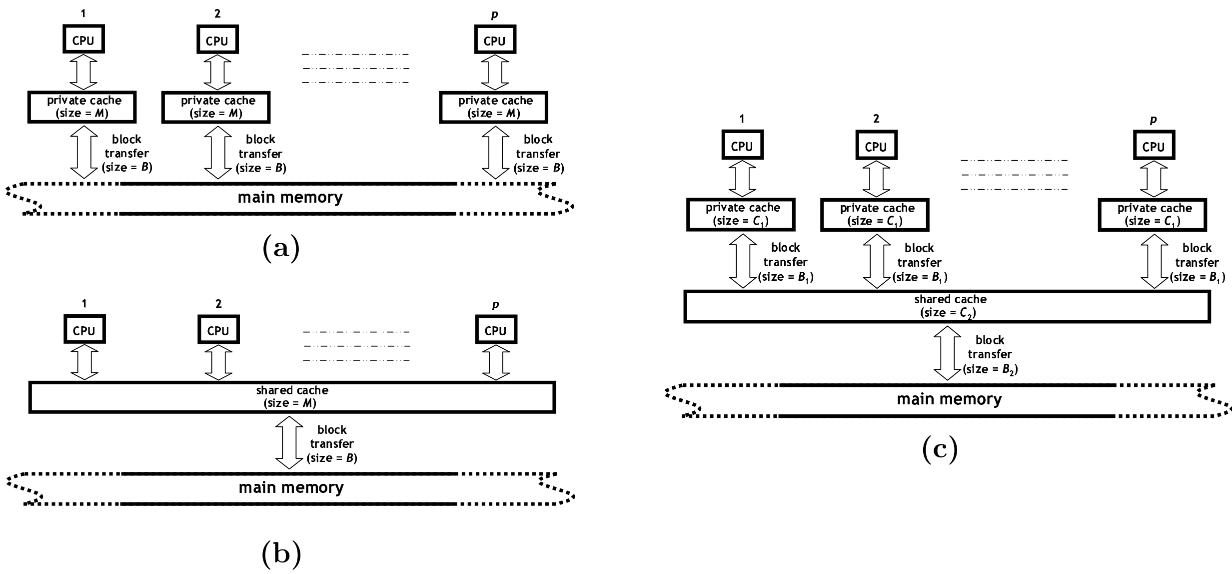

Figure 1: Caching models for CMP: (a) **D-CMP**, (b) **S-CMP**, and (c) **Multicore**.

critical path length $I_\infty(n)$ to denote the number of parallel steps in a parallel algorithm that also matches the sequential work and I/O bound.

For all of the problems mentioned above we present CMP algorithms with $I_\infty(n) = \Theta(n)$ on all three models described in Section 1.1. Since the sequential running time is $\Theta(n^2)$ for LCS and PA, is $\Theta(n^3)$ for Median, GEP and Parenthesis, and $\Theta(n^4)$ for RNA-SP, we obtain a good amount of parallelism with $I_\infty(n) = \Theta(n)$. In fact, the defining DP algorithm for each of the problems, which is the one we use to derive our cache-efficient implementation, has critical path length $\Theta(n)$, hence we achieve the maximum parallelism possible while remaining as cache-efficient as the sequential case.

If we look purely for parallelism, there are NC algorithms known for all of these problems, and there are also work-optimal parallel algorithms that beat our parallel bound I_∞ (see, e.g., [16] and references therein for LDDP and Parenthesis problem). However, none of these results incorporate cache-efficiency in the parallel context, which is crucial for CMPs.

1.1 Caching models for CMP

We consider the following three models for CMPs. In all three models, data is transferred from one memory/caching level to another in a block of a given size.

1. **D-CMP**. Here the CMP is viewed as a p-processor machine, where each processor has its own private cache of size M [15, 10] (i.e., distributed caches). There is also a global shared memory that is arbitrarily large. The performance of a parallel algorithm is measured in terms of the number of parallel steps executed and the total number of block transfers of size B across all caches, assuming the caches are 'ideal' [14].

2. **S-CMP**. This is similar to D-CMP, except that there is a single cache of size $M \geq p \cdot B$ shared by all p processors [3, 10], where B is the block size. There continues to be a global shared-memory that is arbitrarily large, and performance is measured as in D-CMP.

3. **Multicore**. This models the trend in current CMPs: distributed L_1 caches (i.e., private to each core), each of size C_1, and a single shared L_2 cache of size $C_2 \geq p \cdot C_1$ [2]. There continues to be

Problem	I/O	Previous I_∞	Our I_∞	Tiling Parameters of Our Algorithms			
				Sequential	D-CMP	S-CMP	Multicore
LCS & PA	$\mathcal{O}\left(\frac{n^2}{MB}\right)$	n^2 (seq.)	$\mathcal{O}(n)$	$t[d]=2$, $\forall d$	$t[0]=p$, $t[d]=2,\ d>0$	$t[d]=2,\ d<r$, $r=\log(n/p)$, $t[r]=p$	$t[r]=p$, $r=\log(n/C_2)$, $t[d]=2,\ d\neq r$
Median	$\mathcal{O}\left(\frac{n^3}{B\sqrt{M}}\right)$	n^3 (seq.)	$\mathcal{O}(n)$	$t[d]=2$, $\forall d$	$t[0]=\sqrt{p}$, $t[d]=2,\ d>0$	$t[d]=2,\ d<r$, $r=\log(n/\sqrt{p})$, $t[r]=\sqrt{p}$	$t[r]=\sqrt{p}$, $r=\log(n/\sqrt{C_2})$, $t[d]=2,\ d\neq r$
Parenthesis		*					
GEP		$\mathcal{O}\left(n\log^2 n\right)$ [10, 2]					
RNA-SP	$\mathcal{O}\left(\frac{n^4}{B\sqrt{M}}\right)$	n^4 (seq.)	$\mathcal{O}(n)$	–	–	–	–

Table 1: Table of our results. Here the 'I/O' column lists the sequential cache-oblivious bound, and the I_∞ columns list the the number of parallel steps in a work-optimal parallel algorithm whose cache complexity matches the sequential bound. The tiling parameters are explained in the section for each problem. RNA-SP uses 3D-LDDP (similar to Median) and GEP as subroutines, and derives its tiling parameters from both. All results assume that the input is too large to fit into the available cache space, and some of them (i.e., Median, Parenthesis, GEP and RNA-SP) also assume that the cache is *tall* (i.e., $M=\Omega\left(B^2\right)$). * For Parenthesis a cache-efficient parallel algorithm for the IBM Cyclops64 processor is given in [30].

an arbitrarily large global shared memory. The cache complexity is specified by two parameters: the number of block transfers, each of size B_2, into L_2, and the total number of block transfers, each of size B_1, across all L_1 caches. One can consider a more general hierarchy of caches that are successively shared by larger groups of processors, and our results generalize to this model as well. We do not elaborate further on this.

We do not consider cache coherence protocols [20] since the updates that are performed in parallel in the algorithms we present are always on disjoint sets of data, and hence coherence is never invoked. We specify parallelism by parallel ***for*** loops and forking and joining through recursive calls. Our results continue to hold if we replace the ideal cache assumption by LRU.

1.2 On-line Schedules

We view the computation as a dynamic DAG that unfolds as the computation proceeds. Each node in the DAG represents a sequence of instructions being executed, and an edge from node u to node v indicates that computation at v uses a value computed at u. At any point in time, a node can be scheduled on a processor provided its predecessors in the DAG have been evaluated.

A sequential computation will compute in program order, which automatically satisfies the DAG constraints, and this is a depth-first topological sort of the DAG, called a *1DF schedule*. In the parallel context, we need a method to decide how the computation is distributed among the available processors as it unfolds. In *work-stealing* [5], the unfolding DAG is distributed across the processors, and an idle processor 'steals' some work from a random neighbor [5, 15]. Work-stealing has good performance for certain classes of algorithms under D-CMP. A *PDF* schedule assigns to an available processor, the node earliest in the 1DF schedule that is ready to be executed [4, 3], and it often has good performance under S-CMP. The recently proposed *controlled-PDF* scheduler [2] is a refined version of PDF that gives good cache performance for many divide and conquer algorithms under the Multicore caching model.

1.3 Overview of Our Technique

For each of the three classes of problems we consider, we introduce a *tiling sequence* $t[d],\ d\geq 0$, and our CMP algorithm is a parallel recursive algorithm that uses tiling parameter $t[d]$ at level d of the recursion.

For each type of CMP (D-CMP, S-CMP, or Multicore) we specify the tiling parameters and then give a parallel schedule that ensures good performance with respect to both parallelism and cache-efficiency. The known cache-oblivious sequential algorithms for the problems (given in [9] for LCS and GEP, in [8] for LDDP, and in [6] for Parenthesis) can be viewed as special cases of our multicore algorithms, where the tiling parameter $t[d]$ is the same constant for all d (and where no parallelism is specified). Table 1 lists our results together with the tiling parameters used. Details are in the following sections, and the pseudocode for each of our algorithms is included in the appendix.

2 LCS and LDDP

Local Dependence DP (LDDP) includes a very large group of problems solvable by dynamic programming. The key feature of this class is that it applies to some constant number of dimensions $k \geq 2$, and it updates each position in a k-dimensional table by considering the previously computed values immediately adjacent to the current position, and using exactly one of those values to compute the value at the current position. Several LDDP problems are of practical importance, and we present experimental results in Section 5 for two such problems, *pairwise sequence alignment* and *3-way sequence alignment (or median)*, both with affine gaps. Here, for simplicity, we illustrate the LDDP method with arguably the simplest problem in the class, the two-dimensional LCS (or longest common subsequence) problem (see, e.g., [11]). Our results generalize to general k-dimensional LDDP (see [7, 8] for a definition and sequential cache-oblivious algorithm for this problem. For convenience this definition is reproduced in Appendix A.)

Given two sequences $X = x_1 x_2 \cdots x_n$ and $Y = y_1 y_2 \cdots y_n$, (for simplicity, we assume equal-length sequences here), an LCS of X and Y is a sequence of maximum length that is a subsequence of both X and Y. If we define $c[i, j]$, $0 \leq i, j \leq n$, to be the length of an LCS of $x_1 x_2 \cdots x_i$ and $y_1 y_2 \cdots y_j$ then $c[n, n]$ is the LCS length of X and Y, and the $c[i, j]$'s can be computed by dynamic programming using the following recurrence relation (see, e.g., [11]):

$$c[i, j] = \begin{cases} 0 & \text{if } i = 0 \text{ or } j = 0, \\ c[i-1, j-1] + 1 & \text{if } i, j > 0 \wedge x_i = y_j, \\ \max\{\ c[i, j-1],\ c[i-1, j]\ \} & \text{if } i, j > 0 \wedge x_i \neq y_j. \end{cases} \tag{2.1}$$

We denote by $parent(i, j)$, an adjacent cell whose c-value determines $c[i, j]$.

Typically, two types of outputs are expected when evaluating this recurrence: (i) the value of $c[n, n]$, and (ii) the traceback path starting from $c[n, n]$. The *traceback path* from any cell $c[i, j]$ is the path following the chain of *parent* cells through c that ends at some $c[i', j']$ with either $i' = 0$ or $j' = 0$.

Using equation 2.1 LCS can be solved by a simple $\mathcal{O}\left(n^2\right)$ time dynamic programming algorithm, whose DAG has linear critical path length ($T_\infty(n) = \Theta(n)$). Several refinement of this algorithm that optimize for space and cache-efficiency are known (e.g., [21, 9]), and the sequential algorithm in [9] that gives an $\mathcal{O}(n)$-space cache-oblivious algorithm with $\mathcal{O}\left(n^2/(MB)\right)$ cache misses has the best performance across both space and cache-efficiency.

Our tiled LCS algorithm is illustrated in Figure 2. The algorithm has two parts, both of which are recursive. Figure 2(a) illustrates Tiled-Boundary-LCS, which computes the LCS costs for cells corresponding to the output boundary of the LCS table, i.e., the rightmost column and the bottom row, and additionally, for each such cell u, the cell on the input boundary at which the traceback path from u will enter. The algorithm is supplied the subsequences X' and Y' for the current recursive call, together with a positive integer τ, which is the tiling parameter for the current recursive call. Figure 2(a) is shown with $X' = X$ and $Y' = Y$ and with $\tau = 4$. The algorithm proceeds by dividing X and Y into τ equal pieces, thereby tiling the cost table with τ^2 sub-tables. These sub-tables are computed recursively in $2\tau - 1$ parallel steps, where the rth parallel step performs the computation for each sub-table along the

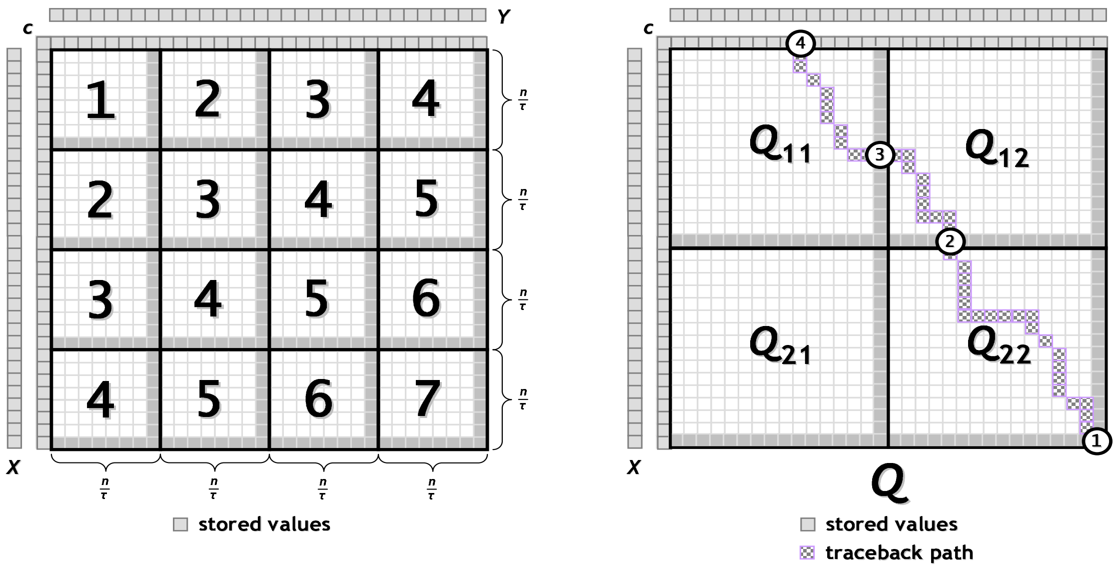

(a) (b)

Figure 2: The recursive tiled LCS algorithm: **(a)** Tiled-Boundary-LCS: Given $\tau = t[\textit{recursion depth}]$, the LCS table is decomposed into τ^2 sub-squares of equal size. In step $r \in [1, 2\tau - 1]$, output boundaries of all sub-squares labelled with r are computed recursively in parallel. Outputs of step r act as inputs for step $r + 1$. For each cell u on an output boundary we also compute the position on the input boundary at which the traceback path from u will enter. **(b)** Tiled-LCS: It first computes the output boundaries of Q_{11}, Q_{12}, Q_{21} and Q_{22} by calling Tiled-Boundary-LCS in sequence. Now given a location labelled 1 on the output boundary of the current table, it determines in $\mathcal{O}(1)$ time the locations at which the traceback path through location 1 intersects the input boundaries of all (at most three) quadrants (e.g., locations 2, 3 and 4 in the figure). Then it recursively calls itself in parallel on all quadrants hit by the traceback path, and extracts the path fragments from each of them.

rth forward diagonal. In the figure all sub-tables computed in step r, for $1 \leq r \leq 7$ are labelled with r. The specification of the algorithm includes the *tiling sequence* $t[d]$, $d \geq 0$, which is a sequence of integers, and at recursion level d the algorithm uses tiling parameter $t[d]$. These parameters will be optimized for the three different caching models we consider. Additionally, we obtain the sequential algorithm in [9] by using $t[d] = 2$ for all d. As in that sequential algorithm, the space requirement of our tiled algorithm is linear in the size of the input (even though a quadratic number of intermediate values are computed).

Figure 2(b) illustrates Tiled-LCS, which computes the trace-back path from a given cell on the output boundary. Tiled-LCS is again a recursive algorithm, and it also calls Tiled-Boundary-LCS. It takes the same inputs as Tiled-Boundary-LCS, together with one additional input u, which a cell on the output boundary of the current table. In figure 2(b), u is the bottom-right position in the table, and is indicated by 1. The algorithm proceeds by calling Tiled-Boundary-LCS on the four sub-tables Q_{11}, Q_{12}, Q_{21} and Q_{22}, which are derived from the current table considering two equal-sized halves of the two input strings. This computation generates for each output boundary cell x in each of these three sub-tables, the entry point on the input boundary of the sub-table of the traceback path from x. After this computation, the algorithm determines the cells 2, 3, and 4 shown in Figure 2(b), as we describe below.

Cell 2 in the figure is the position on the input boundary of Q_{22} where the traceback path from cell 1 meets. This cell could be on the output boundary of either Q_{12} or Q_{21}, and the figure has chosen this to be Q_{12} (as a result, the traceback path from cell 1 does not pass through Q_{21}). Cell 3 is the cell on the input boundary of Q_{12} where the traceback path from cell 2 meets; this cell will be either on the output boundary of Q_{11} or on the input boundary of the overall table (in which case there is no cell 4). The figure has illustrated the case when cell 3 is on the output boundary of Q_{11}, and in this case, cell 4 is the cell on the input boundary of Q_{11} where the traceback path from 3 meets. This cell is guaranteed to be in the input boundary of the overall table. Since the earlier computation of TILED-BOUNDARY-LCS has computed the corresponding traceback cell on the input boundary for each cell on the output boundary, cells 2, 3, and 4 can be determined by at most 3 look-up steps given cell 1, and hence this computation takes constant time and work.

Once the cells 2, 3 and (if needed) 4 of the traceback path are known, TILED-LCS recursively calls itself on the (at most three) subproblems that contain a portion of the traceback path, using the tiling parameter for next level of recursion.

The base case of the recursion occurs when the parent of cell 1 lies on the input boundary of the current table. Each parent cell is computed in such a base case call.

Correctness of this method is straightforward by induction on the input size. A succinct pseudocode for this algorithm is included in Appendix B (see Figure 7).

2.1 Performance Analysis

For performance, we tailor the tiling sequence to the CMP models as well as the sequential case as follows.

(1) Sequential algorithm in [9]. If we use $t[d] = 2$ for all $d \geq 0$, a 1-processor implementation of this multicore algorithm is exactly the sequential cache-oblivious algorithm in [9] and hence its I/O complexity is $\mathcal{O}\left(\frac{n^2}{MB}\right)$, where M and B are the cache and block sizes respectively.

(2) p-processor D-CMP. Since this model has a private cache for each core, we use a tiling sequence that gives the largest amount of locality to each processor, so that the private caches can be used most effectively. For this we use $t[0] = p$ and $t[d] = 2$ for $d \geq 1$. Note that this results in a cache-oblivious D-CMP algorithm.

The parallel schedule consists of assigning the ith column of sub-tables (see figure 2) to the ith processor at the top level recursion ($d = 0$). Further levels of recursion are executed entirely on the processor that was assigned the subproblem at level 0. On an input of size n and with p processors, the parallel running time for TILED-BOUNDARY-LCS is $T_B(n,p) = \mathcal{O}\left((\frac{n}{p})^2 \cdot (2p-1) + n\right) = \mathcal{O}\left(n^2/p + n\right)$ since there are $2p - 1$ parallel steps, each executing the sequential algorithm on an input of size n/p. There are p^2 subproblems of size n/p, each executing the sequential cache-oblivious algorithm, hence the number of cache misses is $\mathcal{O}\left(\left(\frac{(n/p)^2}{MB} + \frac{n/p}{B} + 1\right) \cdot p^2\right) = \mathcal{O}\left(\frac{n^2}{MB}\right)$ under the natural assumptions that $n \geq pM$ (i.e., the input does not fit within the caches), and block-size $B \leq n/p$ (which will hold since $B \leq M$).

For TILED-LCS, we schedule the (up to) 3 recursive subproblems with $p/3$ processors each, hence the parallel running time $T(n,p)$ is given by $T(n,p) = \mathcal{O}\left(n^2/p + n\right) + T(n/2, p/3)$, which remains $\mathcal{O}\left(n^2/p + n\right)$. By a similar analysis the cache complexity also remains the same as for TILED-BOUNDARY-LCS. Hence we obtain $I_\infty(n) = \Theta(n)$.

(3) p-processor S-CMP. Here all processors share a single cache, hence we choose a tiling schedule that causes the processors to work on parallel tasks that are close to one another in the sequential order. For this, we use $t[d] = 2$ for $d < \log(n/p)$ and $t[\log_2(n/p)] = p$. This is again a cache-oblivious strategy.

The parallel schedule consists of a Brent-type processor allocation (see, e.g., [22]) at recursion

G(c, n, f, Σ_G)

(Input $c[1\ldots n, 1\ldots n]$ is an $n\times n$ matrix, $f(\cdot,\cdot,\cdot,\cdot)$ is an arbitrary problem-specific function, and Σ_G is a problem-specific set of triplets such that $c[i,j] \leftarrow f(c[i,j], c[i,k], c[k,j], c[k,k])$ is executed in line 4 if $\langle i,j,k\rangle \in \Sigma_G$.)

1. ***for*** $k \leftarrow 1$ ***to*** n ***do***
2. ***for*** $i \leftarrow 1$ ***to*** n ***do***
3. ***for*** $j \leftarrow 1$ ***to*** n ***do***
4. ***if*** $\langle i,j,k\rangle \in \Sigma_G$ ***then*** $c[i,j] \leftarrow f(c[i,j], c[i,k], c[k,j], c[k,k])$

Figure 3: GEP: Triply nested ***for*** loops typifying code fragment with structural similarity to the computation in Gaussian elimination without pivoting.

level $\log n/p$. At this level of recursion the input has two substrings of length p and hence we are applying parallelism at the finest level of granularity. Since the shared cache can be expected to have size $M \geq p$ this computation will have the same cache-complexity as the sequential case. There are $\Theta\left(\frac{n^2}{p^2}\right)$ subproblems executed, each with parallel time $\Theta(p)$, hence the parallel running time is $\mathcal{O}\left(p \cdot (n^2/p^2) + n\right)$ and hence remains $\mathcal{O}\left(n^2/p + n\right)$. The analysis of Tiled-LCS is similar to that for D-CMP, and again we have cache-efficient parallelism I_∞ to be $\Theta(n)$.

(4) Multicore. Here we need to adapt to the the private L_1 caches of size C_1 and the shared L_2 cache of size C_2. Hence we use a strategy that combines our approach for S-CMP and D-CMP. Let $r = \log \frac{n}{C_2}$. We use $t[r] = p$ and $t[d] = 2$ for all $d \neq r$.

As in D-CMP, at recursion level r we assign the ith column of sub-tables to the ith processor, $1 \leq i \leq p$. There are $\frac{n^2}{C_2^2}$ subproblems, each with two input strings of length C_2. Each such subproblem is solved similar to the D-CMP case, hence the parallel time is $\mathcal{O}\left(\frac{n^2}{C_2^2} \cdot (2p-1) \cdot (C_2/p)^2 + n\right) = \mathcal{O}\left(n^2/p + n\right)$. The C_2 cache complexity is clearly the sequential complexity $\mathcal{O}\left(n^2/(C_2 B)\right)$ since all parallelism is exposed at a problem size when the entire input fits into the L_2 cache. Since each subproblem whose input size is C_2 is solved with private caches using the same method as in D-CMP, the L_1 cache complexity is $\mathcal{O}\left(\frac{n_2^2}{C_2^2} \cdot \frac{C_2^2}{C_1 B}\right)$, which is $\mathcal{O}\left(n^2/(C_1 B)\right)$. Hence this give a Multicore LCS algorithm with cache-efficient parallelism $I_\infty(n) = \Theta(n)$.

3 GEP (Gaussian Elimination Paradigm) Problems

Let $c[1\ldots n, 1\ldots n]$ be an $n \times n$ matrix with entries chosen from an arbitrary set $\mathcal{S}$, and let $f : \mathcal{S}\times\mathcal{S}\times\mathcal{S}\times\mathcal{S} \rightarrow \mathcal{S}$ be an arbitrary function. By GEP (or the *Gaussian Elimination Paradigm*) introduced by the authors in [9], we refer to the computation in Figure 3. Here the algorithm G modifies c by applying a given set of updates of the form $c[i,j] \leftarrow f(c[i,j], c[i,k], c[k,j], c[k,k])$, where $i,j,k \in [1,n]$. We use the notation $\langle i,j,k\rangle$ $(1 \leq i,j,k \leq n)$ denotes an update of the form $c[i,j] \leftarrow f(c[i,j], c[i,k], c[k,j], c[k,k])$, and we let Σ_G denote the set of such updates that the algorithm needs to perform.

As noted in [9] many practical problems can be solved using the GEP construct, including all-pairs shortest paths, LU decomposition and Gaussian elimination without pivoting, and matrix multiplication.

An $\mathcal{O}\left(\frac{n^3}{B\sqrt{M}}\right)$ I/O cache-oblivious sequential algorithm for solving some important special cases of GEP including all problems mentioned above was presented by the authors in [9]. Later in [10], this implementation was named I-GEP, extended to C-GEP which solves all instances of GEP within the same performance bounds, and was also parallelized to run in $\mathcal{O}\left(\frac{n^3}{p} + n\log^2 n\right)$ time on p processors with schedulers to match its sequential cache complexity on D-CMP and S-CMP. Recently in [2] we presented a scheduler to run I-GEP and C-GEP on multicores within the same performance bounds. In this section we improve the parallel algorithm presented in [10] to run in $\mathcal{O}\left(\frac{n^3}{p} + n\right)$ parallel time on D-CMP, S-CMP and multicores while matching its sequential cache complexity. In other words, the new algorithm achieves $I_\infty(n) = T_\infty(n) = \Theta(n)$.

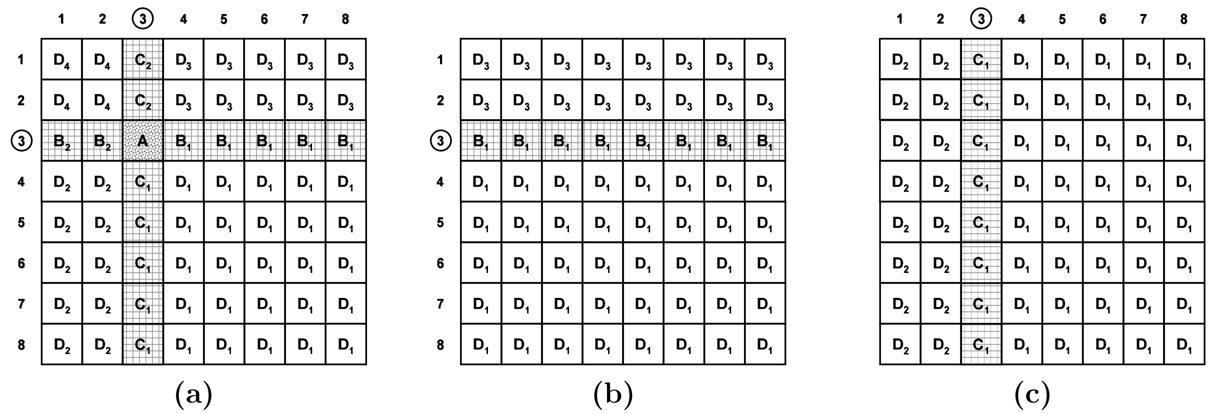

Figure 4: Execution of superstep $k = 3$ of I-GEP functions A (Figure 4(a)), B_1 (Figure 4(b)) and C_1 (Figure 4(c)). Each cell contains the name of the function used to update the entries in the corresponding submatrix. Submatrices corresponding to dotted cells (if any) are updated first. In the next step all submatrices corresponding to cells with grids are updated in parallel. In the last step submatrices corresponding to white cells are updated simultaneously. Pseudocode for these functions are included in Figure 8 of Appendix B.

3.1 Improved CMP Algorithm

We present an improved parallel implementation of I-GEP which is an extension of the parallel implementation we presented in [10]. We modify each I-GEP function (A, B_i, C_i, D_j, where $i \in [1, 2]$ and $j \in [1, 4]$) introduced in [10] so that one can control the way the input matrices are subdivided at each level of recursion by specifying a vector t of tiling parameters. At recursion depth $d \geq 0$, an $n \times n$ input matrix is subdivided into $r \times r$ submatrices of size $\frac{n}{r} \times \frac{n}{r}$ each, where $r = t[d]$. We show that by appropriately setting the tiling parameters we can reduce the parallel time complexity of the algorithm from $\mathcal{O}\left(\frac{n^3}{p} + n \log^2 n\right)$ [10] to $\mathcal{O}\left(\frac{n^3}{p} + n\right)$ for p processors. Pseudocode for each of these functions is given in Figure 8 of Appendix B.

Each function accepts a (recursion) depth parameter $d \geq 0$, and four (not necessarily distinct) equal size square submatrices X, U, V and W of the input matrix c. We assume that W has a diagonal aligned to the $(1, 1)$ to (n, n) diagonal of c, and for each $c[i, j] \in X$ and $c[k, k] \in W$, the entries $c[i, k]$ and $c[k, j]$ can be found in U and V, respectively. Each function updates the entries in X using appropriate entries from U, V and W. The functions differ in the amount of overlap X, U and V have among them. Function A assumes that all three matrices completely overlap, while D_l ($l \in [1, 4]$) expects completely non-overlapping matrices. In the intermediate cases, function B_l assumes that only X and V overlap, while C_l assumes overlap only between X and U, where $l \in [1, 2]$. Intuitively, the less the overlap among the input matrices the more flexibility the function has in ordering its recursive calls, and thus leading to better parallelism.

The initial call is to function A with $d = 0$ and $X = U = V = W = c$. If X is a 1×1 matrix, then X is updated directly by setting $X \leftarrow f(\ X,\ U,\ V,\ W\ )$. Otherwise each matrix is subdivided into $r \times r$ submatrices of size $\frac{n}{r} \times \frac{n}{r}$ each, where $r = t[d] = t[0]$. The submatrix of X at the i-th position from the top and the j-th position from the left is denoted by $X_{i,j}$. Then the function executes r supersteps. Superstep $k \in [1, r]$ of function A consists of 3 steps (see Figure 4(a)). In step 1, submatrix $X_{k,k}$ is updated recursively by function A. In step 2, the remaining submatrices $X_{i,k}$ and $X_{k,j}$ ($i \neq k, j \neq k$)

are updated in parallel using the entries in $X_{k,k}$ by appropriate calls to functions B_1, B_2, C_1 and C_2 (see Figure 4(a) for details). In step 3, the remaining submatrices of X are updated in parallel using entries computed in step 2 by calling appropriate D_l ($l \in [1,4]$) functions (see Figure 4(a)). Superstep $k \in [1,r]$ of functions B_l and C_l ($l \in [1,2]$) consists of 2 steps. In case of function B_l (C_l), step 1 updates all $X_{i,k}$ ($X_{k,j}$, resp.) by parallel calls to B_l (C_l, resp.). In step 2, the remaining submatrices of X are updated by parallel calls to appropriate D_l ($l \in [1,4]$) functions (e.g., see Figures 4(b) and 4(c)). Each superstep of function D_l ($l \in [1,4]$) has only one step which updates all submatrices of X in parallel by calling D_l recursively.

We now analyze the performance of the algorithm under different configurations below.

(1) Sequential. We use $t[d] = 2$ for all $d \geq 0$, and execute the entire algorithm on a single processor. We have already analyzed this case in [9, 10], and found the cache complexity of the algorithm to be $\mathcal{O}\left(n^3/(B\sqrt{M}) + n^3/M + n^2/B + 1\right)$ which reduces to $\mathcal{O}\left(n^3/(B\sqrt{M})\right)$ provided the cache is tall (i.e., $M = \Omega\left(B^2\right)$) and the input matrix is too large to fit into the cache (i.e., $n^2 > M$). It is straight-forward to see that the algorithm runs in $\mathcal{O}\left(n^3\right)$ time.

(2) D-CMP. We set $t[d] = \sqrt{p}$ for $d = 0$, and $t[d] = 2$ for $d > 0$. Hence at level 0 of recursion, function A will execute $\sqrt{p}$ supersteps, and in each superstep it will make a total of p parallel recursive functions calls with submatrices of size $(n/\sqrt{p}) \times (n/\sqrt{p})$ each. Each such function will be executed sequentially in $\mathcal{O}\left(\left(n/\sqrt{p}\right)^3\right)$ time on a single processor. Thus the parallel running time of the algorithm is $\sqrt{p} \cdot \mathcal{O}\left(\left(n/\sqrt{p}\right)^3\right) = \mathcal{O}\left(n^3/p\right)$. Now let $Q_{\mathrm{A}}(n)$, $Q_{\mathrm{BC}}(n)$ and $Q_{\mathrm{D}}(n)$ denote the cache complexity of functions A, $\mathrm{B}_1/\mathrm{B}_2/\mathrm{C}_1/\mathrm{C}_2$ and $\mathrm{D}_1/\mathrm{D}_2/\mathrm{D}_3/\mathrm{D}_4$, respectively, on an input of size n. Assuming that a submatrix of size $(n/\sqrt{p}) \times (n/\sqrt{p})$ is too large to fit into the cache, and that the cache is tall, the total number of cache misses is $\sqrt{p} \cdot Q_{\mathrm{A}}(n/2) + \sqrt{p} \cdot (2\sqrt{p} - 1) \cdot Q_{\mathrm{BC}}(n/2) + \sqrt{p} \cdot (p - 2\sqrt{p}) \cdot Q_{\mathrm{D}}(n/2) = \mathcal{O}\left(n^3/(B\sqrt{M}) + n^3/M + \sqrt{p} \cdot n^2/B + p\sqrt{p}\right) = \mathcal{O}\left(n^3/(B\sqrt{M})\right)$.

(3) S-CMP. We set $t[d] = \sqrt{p}$ if $d \geq \log_2 \frac{n}{\sqrt{p}}$, and $t[d] = 2$ otherwise. We do not make any parallel function calls until we reach a level $d \geq \log_2 \frac{n}{\sqrt{p}}$. Hence, we can use the same recurrence relations as in the sequential case and get the same cache complexity provided $p \leq M$. Now $T_{\mathrm{A}}(n)$, $T_{\mathrm{BC}}(n)$ and $T_{\mathrm{D}}(n)$ denote the parallel running times of functions A, $\mathrm{B}_1/\mathrm{B}_2/\mathrm{C}_1/\mathrm{C}_2$ and $\mathrm{D}_1/\mathrm{D}_2/\mathrm{D}_3/\mathrm{D}_4$, respectively on an input of size n. Then $T_{\mathrm{A}}(n) = T_{\mathrm{BC}}(n) = T_{\mathrm{D}}(n) = \mathcal{O}\left(\sqrt{p}\right)$ if $n \leq \sqrt{p}$. Otherwise,

$$T_{\mathrm{D}}(n) = 8 \cdot T_{\mathrm{D}}\left(n/2\right) + \mathcal{O}\left(1\right)$$

$$T_{\mathrm{BC}}(n) = 4 \cdot T_{\mathrm{BC}}\left(n/2\right) + 4 \cdot T_{\mathrm{D}}\left(n/2\right) + \mathcal{O}\left(1\right)$$

$$T_{\mathrm{A}}(n) = 2 \cdot T_{\mathrm{A}}\left(n/2\right) + 4 \cdot T_{\mathrm{BC}}\left(n/2\right) + 2 \cdot T_{\mathrm{D}}\left(n/2\right) + \mathcal{O}\left(1\right)$$

Solving the recurrences and assuming that $p \leq n^2$, we obtain $T_{\mathrm{A}}(n) = T_{\mathrm{BC}}(n) = T_{\mathrm{D}}(n) = \mathcal{O}\left(n^3/p\right)$.

(4) Multicore. We set $t[d] = \sqrt{p}$ if $d = \log_2\left(n/\sqrt{C_2}\right)$, and $t[d] = 2$ otherwise. All parallel calls are made only at level $d = \log_2\left(n/\sqrt{C_2}\right)$. Observe that whenever we reach a subproblem of input size $n_1 = \sqrt{C_2/p}$, it is executed entirely on a single processor, and there are $\left(n/n_1\right)^3$ such subproblems. Hence, the number of L_1 cache-misses is $\left(n/n_1\right)^3 \cdot \mathcal{O}\left({n_1}^3/(B_1\sqrt{C_1}) + {n_1}^3/C_1 + {n_1}^2/B_1 + 1\right)$, which reduces to $\mathcal{O}\left(n^3/(B_1\sqrt{C_1})\right)$ under the assumption that $C_2 > p \cdot C_1$ ($> p \cdot B_1$), and that the L_1 cache is tall (i.e., $C_1 = \Omega\left({B_1}^2\right)$). The number of L_2 cache-misses can be computed using the same recurrence relations as in the sequential case and in **S-CMP**, and thus the L_2 cache complexity is $\mathcal{O}\left(n^3/(B_2\sqrt{C_2})\right)$ provided the L_2 cache is tall (i.e., $C_2 = \Omega\left({B_2}^2\right)$), and the input matrix is too large to fit into that cache

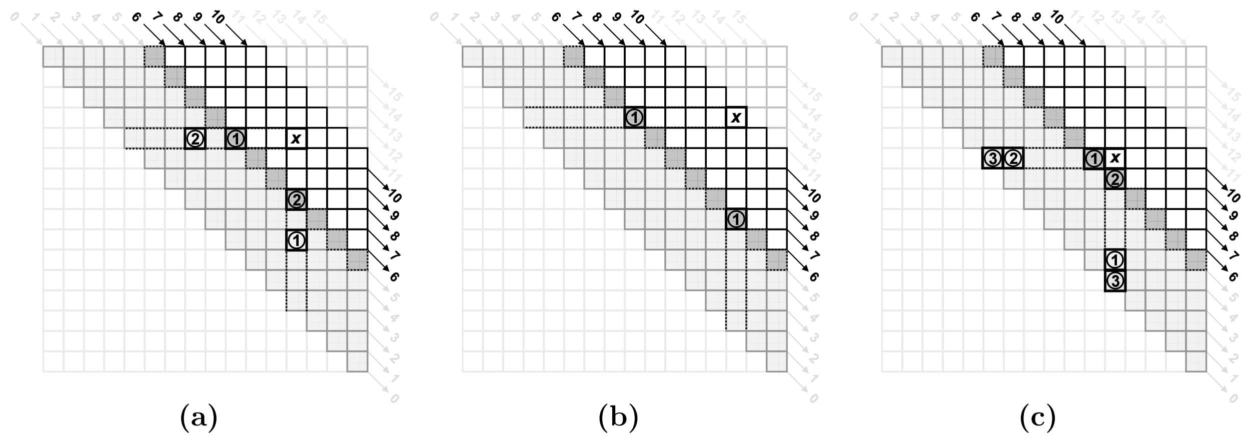

Figure 5: Execution of superstep $k = 6$ of function E for solving the parenthesis problem. Figure 5(a) marks the submatrices used in steps 1 and 2 to update a submatrix x on diagonal $k' \in [k+1, 2k-3]$. Figure 5(b) shows that if x is on diagonal $2k-2$ then it is updated only in step 1, and as shown in Figure 5(c) any submatrix on diagonal k is updated in all three steps (submatrices used in step $r \in [1,3]$ are labelled with r).

(i.e., $n^2 > C_2$). The parallel running time can be computed using the same recurrence relations as in **S-CMP**, but with a different base condition. We use $T_{\mathrm{A}}(n) = T_{\mathrm{BC}}(n) = T_{\mathrm{D}}(n) = \mathcal{O}\left(\sqrt{p} \cdot \left(\sqrt{C_2/p}\right)^3\right)$ for $n \le \sqrt{C_2}$. Assuming $n^2 > C_2$, we obtain $T_{\mathrm{A}}(n) = T_{\mathrm{BC}}(n) = T_{\mathrm{D}}(n) = \mathcal{O}\left(n^3/p\right)$.

Observe that for all three models we obtain $I_\infty(n) = \Theta(n)$. For D-CMP and S-CMP models, our algorithm is cache-oblvious.

4 The Parenthesis Problem

We consider the *parenthesis problem* [16] which is defined by the following recurrence relation:

$$c[i,j] = \begin{cases} x_j & \text{if } 0 \le i = j-1 < n, \\ \min_{i<k<j} \left\{ \begin{matrix} (c[i,k] + c[k,j]) \\ +w(i,k,j) \end{matrix} \right\} & \text{if } 0 \le i < j-1 < n; \end{cases} \tag{4.2}$$

where x_j's are assumed to be given for $j \in [1,n]$. We also assume that $w(\cdot,\cdot,\cdot)$ is a function that can be computed in-core without incurring any cache misses.

The class of problems defined by the recurrence relation above includes RNA secondary structure prediction, optimal matrix chain multiplication, construction of optimal binary search trees, and optimal polygon triangulation. A variant of this recurrence which does not include the $w(i,k,j)$ term and is defined as the *simple dynamic program*, was considered in [6], where an $\mathcal{O}\left(n^3/(B\sqrt{M})\right)$ I/O sequential cache-oblivious algorithm based on Valiant's context-free language recognition algorithm [32] was given for solving the recurrence. A parallel algorithm for solving the parenthesis problem which runs in $\mathcal{O}\left(n^{\frac{3}{4}} \log n\right)$ time and performs optimal $\mathcal{O}\left(n^3\right)$ work was given in [16], but the algorithm is not cache-efficient. A cache-efficient multicore algorithm for the IBM Cyclops64 processor was given in [30].

As in [6], instead of recurrence 4.2 we will use the following slightly augmented version of 4.2 which will considerably simplify the recursive subdivision process in our cache-oblivious algorithm.

$$c[i,j] = \begin{cases} \infty & \text{if } 0 \le i = j \le n, \\ x_j & \text{if } 0 \le i = j-1 < n, \\ \min_{i \le k \le j} \left\{ \begin{matrix} (c[i,k] + c[k,j]) \\ +w(i,k,j) \end{matrix} \right\} & \text{if } 0 \le i < j-1 < n; \end{cases} \quad (4.3)$$

where $w(i,k,j)$ is defined to be ∞ when $k = i$ or $k = j$. It is straight-forward to see that recurrences 4.2 and 4.3 are equivalent, i.e., they compute the same values for any given $c[i,j]$, $0 \le i < j-1 < n$.

4.1 CMP Algorithm

We assume that the input to the algorithm is an $n \times n$ matrix c that has all $c[i,j]$ with $0 \le i \le j \le i+1 \le n$ initialized as in recurrence 4.3, and the remaining entries (i.e., all $c[i,j]$ with $0 \le i < j-1 < n$) initialized to ∞. The algorithm works by recursively subdividing the input matrix, and assumes the existence of a global vector t of tiling parameters which for each level of recursion specifies how the input matrix is subdivided. Pseudocode for the algorithm is given in Figure 9 of Appendix B. Here we give an informal description of the algorithm.

The algorithm (i.e., function E in Figure 9 of Appendix B) splits the input matrix X (which is initially set to c) into $r \times r$ submatrices of size $\frac{n}{r} \times \frac{n}{r}$ each, where $r = t[d]$ and d (≥ 0) is the level (i.e., depth) of recursion. The submatrix of X at the i-th position from the top and the j-th position from the left is denoted by $X_{i,j}$. The k-th diagonal of X includes the submatrices $X_{i,i+k}$, where $k \in [0, n-1]$ and $i \in [1, n-k]$ (see Figure 5). Observe that the submatrices on diagonal 0 of X are smaller instances of the original parenthesis problem defined by X and all of them are independent, and hence are solved recursively by $\frac{n}{r}$ parallel calls to E. The algorithm then executes $r-1$ supersteps, and maintains the following invariant.

Invariant 4.1. *At the start of superstep $k \in [1, r-1]$,*

(i) all entries of c in the submatrices on diagonals 0 to $k-1$ of X have their values finalized, and
(ii) all remaining cells $c[i,j]$ are updated with $(c[i,l] + c[l,j]) + w(i,l,j)$, where both $c[i,l]$ and $c[l,j]$ belong to the submatrices on diagonals 0 to $k-2$ of X.

Superstep k ($1 \le k < r$) consists of up to 3 parallel steps. The first two steps are executed only for $k > 1$. In step 1 each submatrix $X_{i,j}$ on diagonals k to $\min\{2k-2, r-1\}$ is updated using data from the submatrix $X_{i,r+k-i}$ on diagonal $k-1$ that lies to the left of $X_{i,j}$, and the corresponding submatrix $X_{i-k+1,j}$ below $X_{i,j}$. All such $X_{i,j}$'s are updated in parallel using another recursive function H (see Figure 9 of Appendix B for pseudocode) which is implemented in exactly the same way as the I-GEP function D_3. Similarly, step 2 updates each submatrix $X_{i,j}$ on diagonals k to $\min\{2k-3, r-1\}$ using data from the submatrix $X_{r+k-j,j}$ on diagonal $k-1$ that lies below $X_{i,j}$, and the corresponding submatrix $X_{i,j-k+1}$ to the left of $X_{i,j}$. All such updates are also performed in parallel by calling H. Observe that in step 1, for $X_{i,j}$'s on diagonal $2k-2$ both input submatrices lie on diagonal $k-1$. Hence, submatrices on diagonal $2k-2$ need not be updated in step 2. In step 3 each submatrix $X_{i,j}$ on diagonal k is updated using the two submatrices on diagonal 0 that lie to the left of $X_{i,j}$ and below it, i.e., $X_{i,r-i+1}$ and $X_{r-j+1,j}$, respectively. All these updates are performed in parallel by calling function F. The implementation of this function is similar to that of function E (see Figure 9 of Appendix B for pseudocode). Since invariant 4.1 was true before the start of superstep k, it is easy to see that after step 3 of this superstep all entries of c inside submatrices on diagonal k will have their values finalized.

Now we analyze the performance of the algorithm under different machine configurations.

(1) Sequential. We use $t[d] = 2$ for all $d \ge 0$, and execute the entire algorithm on a single processor. In this case the algorithm reduces to the sequential algorithm given in [6]. It is straight-forward to see that the running time of each of the three functions (i.e., E, F and H) is $\mathcal{O}\left(n^3\right)$. Let the sequential cache-complexity of the three functions on an input of size n be $Q_E(n)$, $Q_F(n)$ and $Q_H(n)$, respectively. Then $Q_E(n) = Q_F(n) = Q_H(n) = \mathcal{O}\left(n + n/B\right)$ for $n^2 \leq \gamma M$, where γ is a suitable constant. Otherwise,

$$Q_H(n) = 8 \cdot Q_H\left(n/2\right), \quad Q_F(n) = 4 \cdot Q_F\left(n/2\right) + 4 \cdot Q_H\left(n/2\right)$$

$$\text{and} \quad Q_E(n) = 2 \cdot Q_E\left(n/2\right) + Q_F\left(n/2\right)$$

Solving the recurrences we obtain, $Q_E(n) = Q_F(n) = Q_H(n) = \mathcal{O}\left(n^3/(B\sqrt{M}) + n^3/M + n^2/B + 1\right)$ which reduces to $\mathcal{O}\left(n^3/(B\sqrt{M})\right)$ provided the cache is tall and the input matrix is too large to fit into the cache.

(2) D-CMP. We use the same tiling parameters as in the **D-CMP** case of I-GEP in Section 3.1, and the analyses of performance bounds are also similar. The parallel running time of the algorithm turns out to be $\mathcal{O}\left(n^3/p\right)$, and the cache complexity $\mathcal{O}\left(n^3/(B\sqrt{M}) + n^3/M + \sqrt{p} \cdot n^2/B + p\sqrt{p}\right)$, which reduces to $\mathcal{O}\left(n^3/(B\sqrt{M})\right)$ provided the input is too large to fit into the cache (i.e., $n^2 > p \cdot M$) and the cache is tall (i.e., $M = \Omega\left(B^2\right)$).

(3) S-CMP. We use the same tiling parameters as in the **S-CMP** case of I-GEP in Section 3.1, and compute the cache complexity similarly. Let $T_{\mathrm{E}}(n)$, $T_{\mathrm{F}}(n)$ and $T_{\mathrm{H}}(n)$ be the parallel running times of E, F and H, respectively, on input size n. Then $T_E(n) = T_F(n) = T_H(n) = \mathcal{O}\left(\sqrt{p}\right)$ if $n \leq \sqrt{p}$. Otherwise,

$$T_H(n) = 8 \cdot T_H\left(n/2\right) + \mathcal{O}\left(1\right)$$

$$T_F(n) = 4 \cdot T_F\left(n/2\right) + 4 \cdot T_H\left(n/2\right) + \mathcal{O}\left(1\right)$$

$$T_E(n) = 2 \cdot T_E\left(n/2\right) + T_F\left(n/2\right) + \mathcal{O}\left(1\right)$$

Solving the recurrences and assuming that $p \leq n^2$, we obtain $T_E(n) = T_F(n) = T_H(n) = \mathcal{O}\left(n^3/p\right)$. The cache complexity remains the same as in the sequential case.

(4) Multicore. The analysis is exactly the same (using the same tiling sequence) as in the multicore case of I-GEP in Section 3, and we get $T_E(n) = T_F(n) = T_H(n) = \mathcal{O}\left(n^3/p\right)$. The cache complexities for both the L_1 and L_2 caches match its sequential cache complexity.

Observe that similar to the results for I-GEP in Section 3, this algorithm acheives $I_\infty(n) = \mathcal{O}\left(n\right)$ for all three CMP models, and is cache-oblivious for D-CMP and S-CMP.

5 Experimental Results

We ran experiments for PA *(pairwise global sequence alignment with affine gap penalty)* and for Median *(median of 3 sequences, again with affine gap penalty)*. Definitions for PA and Median problems as 2- and 3-dimensional LDDP respectively can be found in Appendix A. Our sequential cache-oblivious algorithms and experimental results for them are given in [8]. We ran our experiments on an 8-core 2.2 GHz AMD Opteron 850, with cache sizes 64KB and 1 MB (8-way) for L1 and L2 caches, 32 GB RAM.

Since we were dealing with only $p = 8$ cores, we used a simple CMP algorithm with tiling parameter $t[d] = 2$ for all d, which is the same tiling sequence used by the sequential algorithm in [9]. This gives rise to parallelism $T_\infty(n) = \mathcal{O}\left(n^{\log_2 3}\right)$ since it satisfies the recurrence $T_\infty(n) = 3 \cdot T_\infty(n/2) + \mathcal{O}\left(1\right)$, with $T_\infty(1) = \mathcal{O}\left(1\right)$. Using techniques similar to those we have described in section 2 for LCS on D-CMP, S-CMP and Multicore, it is straightforward to obtain schedules that achieve $I_\infty(n) = \mathcal{O}\left(n^{\log_2 3}\right)$ on all three models. However, since our experiments were actual runs on an existing 8-core machine and not

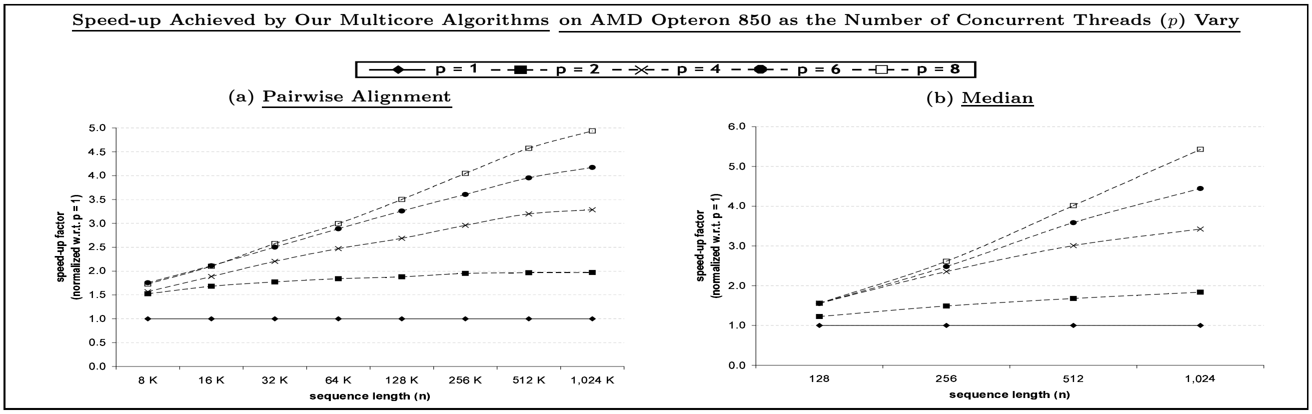

Figure 6: Performance of algorithms for PA and Median on randomly generated sequences over $\{ A, T, G, C \}$.

Seq. Pair	Lengths ($\times 10^6$)	Cost ($\times 10^3$)	FASTA [27]	Multicore PA 1 core	2 cores	4 cores	6 cores	8 cores
human/baboon	1.80|1.51	689	$21h$ $43m$ (5.87)	$17h$ $41m$ (4.79)	$8h$ $58m$ (2.43)	$5h$ $26m$ (1.47)	$4h$ $21m$ (1.17)	$3h$ $42m$ (1.00)
human/chimp	1.80|1.32	585	$20h$ $15m$ (6.34)	$14h$ $28m$ (4.53)	$7h$ $20m$ (2.30)	$4h$ $28m$ (1.40)	$3h$ $32m$ (1.11)	$3h$ $12m$ (1.00)
baboon/chimp	1.51|1.32	574	$17h$ $57m$ (6.22)	$13h$ $0m$ (4.51)	$6h$ $36m$ (2.29)	$4h$ $3m$ (1.40)	$3h$ $16m$ (1.13)	$2h$ $52m$ (1.00)
human/rat	1.80|1.50	1,143	$27h$ $55m$ (6.15)	$21h$ $47m$ (4.80)	$11h$ $2m$ (2.43)	$6h$ $37m$ (1.46)	$5h$ $11m$ (1.14)	$4h$ $32m$ (1.00)
rat/mouse	1.50|1.49	822	$18h$ $39m$ (5.31)	$15h$ $42m$ (4.47)	$7h$ $58m$ (2.27)	$4h$ $55m$ (1.40)	$3h$ $54m$ (1.11)	$3h$ $31m$ (1.00)

Table 2: Performance of pairwise alignment algorithms on 8-core AMD Opteron 850 on CFTR DNA sequences [31]. Parameters used: gap open cost = 2, gap extension cost = 1, mismatch cost = 1. Each number outside parentheses in columns 4–9 is the time for a single run, and the ratio of that running time to the corresponding running time for Multicore PA with 8 cores is given within parentheses.

simulations, we had no control over the scheduler, so our parallel code was run with the default Linux thread scheduler.

Our algorithms were implemented in C++ (compiled with *g++* 3.3.4) while some software packages we used for comparison were written in C (compiled with *gcc* 3.3.4). Optimization parameter -O3 was used in all cases.

For PA we compared our code (PA-CO) with FASTA *(fasta2)* [27], and for median we compared our code (Median-CO) with Knudsen [25], and with ukk.alloc and ukk.checkp, both from Powell et al. [29]. These are all well-known software, and FASTA especially is very widely used. None of the code we used for comparison were designed for parallelism, but in all cases even the 1-core version of our parallel code ran faster than the all compared code on random triples from the data set.

Experimental Performance. Our experimental results for random input strings in Figure 6 show that both PA-CO and Median-CO achieve good speed-up as the number of processors increases. For example, with 8 processors PA-CO achieves a speed-up factor of about 5 when $n = 1024$ K, and MED-CO achieves speed-up 5.5 when $n = 1024$.

Tables 2 and 3 present a sample of runs on real data, both on our code and the other software. As seen from the tables, the 1-core runs of our code are faster than the software we compare against. The speed-up for these real data samples as we increase the number of cores is similar to our results for random input strings in Figure 6.

No	Lengths	Cost	Knudsen [24]	ukk.alloc [28]	ukk.checkp [28]	Multicore Median 1 core	2 cores	4 cores	6 cores	8 cores
1	367|387|388	299	1,061 (6.93)	396 (2.59)	548 (3.58)	431 (2.82)	274 (1.79)	181 (1.18)	158 (1.03)	153 (1.00)
2	378|388|403	324	1,136 (7.01)	–	707 (4.36)	460 (2.84)	298 (1.84)	190 (1.17)	173 (1.07)	162 (1.00)
3	342|367|389	339	936 (6.37)	–	795 (5.41)	388 (2.64)	256 (1.74)	166 (1.13)	152 (1.03)	147 (1.00)
4	342|370|474	432	1,154 (7.08)	–	1,595 (9.79)	464 (2.85)	281 (1.72)	195 (1.20)	175 (1.07)	163 (1.00)
5	370|388|447	336	–	–	768 (4.74)	494 (3.05)	307 (1.90)	204 (1.25)	170 (1.05)	162 (1.00)

Table 3: Performance on an 8-core AMD Opteron 850 of Median algorithms on triples of random sequences from 16S bacterial rDNA sequences from the Pseudanabaena group [13]. Parameters used: gap open cost = 2, gap extension cost = 1, mismatch cost = 1. A '–' in a column denotes that the corresponding algorithm could not be run due to high space overhead. Each number outside parentheses in columns 4–11 is the time in seconds for a single run, and the ratio of that running time to the corresponding running time for Multicore Median with 8 cores is given within parentheses.

References

[1] S. Altschul and B. Erickson. Optimal sequence alignment using affine gap costs. *Bulletin of Mathematical Biology*, 48:603–616, 1986.

[2] G. Blelloch, R. Chowdhury, P. Gibbons, V. Ramachandran, S. Chen, and M. Kozuch. Provably good multicore cache performance for divide-and-conquer algorithms. In *Proc. ACM-SIAM SODA*, pages 501–510, 2008.

[3] G. Blelloch and P. Gibbons. Effectively sharing a cache among threads. In *Proc. ACM SPAA*, pages 235–244, 2004.

[4] G. Blelloch, P. Gibbons, and Y. Matias. Provably efficient scheduling for languages with fine-grained parallelism. *JACM*, 46, 1999.

[5] R. Blumofe and C. Leiserson. Scheduling multithreaded computations by work stealing. *JACM*, 46:720–748, 1999.

[6] C. Cherng and R. Ladner. Cache efficient simple dynamic programming. In *Proc. Intl Conf Analysis of Algorithms*, pages 49–58, 2005.

[7] R. Chowdhury. *Cache-efficient Algorithms and Data Structures: Theory and Experimental Evaluation*. PhD thesis, Department of Computer Sciences, The University opf Texas at Austin, 2007.

[8] R. Chowdhury, H. Le, and V. Ramachandran. Efficient cache-oblivious string algorithms for Bioinformatics. Technical Report TR-07-03, Dept of Comp Sci, UT-Austin, 2007, submitted for journal publication.

[9] R. Chowdhury and V. Ramachandran. Cache-oblivious dynamic programming. In *Proc. ACM-SIAM SODA*, pages 591–600, 2006.

[10] R. Chowdhury and V. Ramachandran. The cache-oblivious gaussian elimination paradigm: Theoretical framework, parallelization and experimental evaluation. In *Proc. ACM SPAA*, pages 71–80, 2007.

[11] T. Cormen, C. Leiserson, R. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, second edition, 2001.

[12] D. Culler, R. Karp, D. Patterson, A. Sahay, K. Schauser, S. E., R. Subramonian, and T. von Eicken. Logp: Toward a realistic model of parallel computation. In *Proc. 4th SIGPLAN Symp. Principles Practices of Parallel Programming*, pages 1–12, 1993.

[13] T. DeSantis, I. Dubosarskiy, S. Murray, and G. Andersen. Comprehensive aligned sequence construction for automated design of effective probes (cascade-p) using 16S rDNA. *Bioinformatics*, 19:1461–1468, 2003. url: http://greengenes.llnl.gov/16S/.

[14] M. Frigo, C. Leiserson, H. Prokop, and S. Ramachandran. Cache-oblivious algorithms. In *Proc. IEEE FOCS*, pages 285–297, 1999.

[15] M. Frigo and V. Strumpen. The cache complexity of multithreaded cache oblivious algorithms. In *Proc ACM SPAA*, pages 271–280, 2006.

[16] Z. Galil and K. Park. Parallel algorithms for dynamic programming recurrences with more than $o(1)$ dependency. *JPDC*, 21:213–222, 1994.

[17] P. Gibbons, Y. Matias, and V. Ramachandran. Can shared-memory model serve as a bridging model for parallel computation? In *Proc. ACM SPAA*, pages 72–83, 1997.
[18] O. Gotoh. An improved algorithm for matching biological sequences. *Journal of Molecular Biology*, 162:705–708, 1982.
[19] J. Grice, R. Hughey, and D. Speck. Reduced space sequence alignment. *Computer Applications in the Biosciences*, 13(1):45–53, 1997.
[20] J. L. Hennessy and D. A. Patterson. *Computer Architecture: A Quantitative Approach.* Morgan Kaufmann, third edition, 2002.
[21] D. Hirschberg. A linear space algorithm for computing maximal common subsequences. *CACM*, 18(6):341–343, 1975.
[22] R. Karp and V. Ramachandran. Parallel algorithms for shared memory machines. In *Handbook of Theor Comp Sci*, pages 869–941. Elsevier, 1990.
[23] J. Kleinberg and E. Tardos. *Algorithm Design.* Addison-Wesely, 2005.
[24] B. Knudsen. Multiple parsimony alignment with "affalign". Software package `multalign.tar`.
[25] B. Knudsen. Optimal multiple parsimony alignment with affine gap cost using a phylogenetic tree. In *Proc. Workshop Algs in Bioinf.*, pages 433–446, 2003.
[26] E. Myers and W. Miller. Optimal alignments in linear space. *Computer Applications in the Biosciences*, 4(1):11–17, 1988.
[27] W. Pearson and D. Lipman. Improved tools for biological sequence comparison. In *Proc. Natl Acad. Sciences*, volume 85, pages 2444–2448, 1988.
[28] D. Powell. Software package `align3str_checkp.tar.gz`.
[29] D. Powell, L. Allison, and T. Dix. Fast, optimal alignment of three sequences using linear gap cost. *Journal of Theoretical Biology*, 207(3):325–336, 2000.
[30] G. Tan, N. Sun, and G. R. Gao. A parallel dynamic programming algorithm on a multi-core architecture. In *ACM SPAA*, pages 135–144, 2007.
[31] J. Thomas et al. Comparative analyses of multi-species sequences from targeted genomic regions. *Nature*, 424:788–793, 2003.
[32] L. Valiant. General context-free recognition in less than cubic time. *JCSS*, 10:308–315, 1975.
[33] L. Valiant. A bridging model for parallel computation. *CACM*, 33(8):103–111, 1990.

APPENDIX

A Dynamic Programs with Local Dependencies (LDDP)

For convenience, we describe the definition of LDDP here, which is also available in [7, 8]. Given $d \geq 1$ sequences $S_i = s_{i,1}s_{i,2}\ldots s_{i,n}$, $1 \leq i \leq d$, and functions $h(\cdot)$ and $f(\cdot,\cdot,\cdot)$, we consider dynamic programs that compute entries of a d-dimensional matrix $c[0:n, 0:n, \ldots, 0:n]$ as follows, where $\mathbf{i} = i_1, i_2, \ldots, i_d$ and $\mathbf{S_i}$ is the tuple $\langle\ s_{1,i_1},\ s_{2,i_2}, \ldots,\ s_{d,i_d}\ \rangle$ containing the i_j-th symbol of S_j in j-th position.

$$c[\ \mathbf{i}\ ] = \begin{cases} h\left(\ \langle\ \mathbf{i}\ \rangle\ \right) & \text{if } \exists\ i_j = 0, \\ f\left(\langle\ \mathbf{i}\ \rangle,\ \mathbf{S_i},\ c\begin{bmatrix} i_1 - 1 : i_1, \\ i_2 - 1 : i_2, \\ \ldots, \\ i_d - 1 : i_d \end{bmatrix} \setminus c[\ \mathbf{i}\ ]\right) & \text{otherwise.} \end{cases} \tag{A.4}$$

Function f can be arbitrary except that it is allowed to use exactly one cell from its third argument to compute the final value of $c[i_1, i_2, \ldots, i_d]$ (though it can consider all cells), and we call that specific cell the *parent cell* of $c[i_1, i_2, \ldots, i_d]$. We also assume that f does not access any memory locations in addition to those passed to it as inputs except possibly some constant size local variables.

Typically, two types of outputs are expected when evaluating this recurrence: (i) the final value of $c[n, n, \ldots, n]$, and (ii) the traceback path starting from $c[n, n, \ldots, n]$. The *traceback path* from any cell

$c[i_1, i_2, \ldots, i_d]$ is the path following the chain of *parent* cells through c that ends at some $c[i'_1, i'_2, \ldots, i'_d]$ with $\exists\ i'_j = 0$.

Each cell of c can have multiple fields and in that case f must compute a value for each field, though as before, it is allowed to use exactly one field from its third argument to compute the final value of any field in $c[i_1, i_2, \ldots, i_d]$. The definition of traceback path extends naturally.

The standard DP evaluation of recurrence A.4 runs in $\mathcal{O}\left(n^d\right)$ time, $\mathcal{O}\left(n^d\right)$ space and $\mathcal{O}\left(n^d/B\right)$ cache-misses (assuming a *tall cache*, i.e., $M = \Omega\left(B^2\right)$). In [8] we present a sequential cache-oblivious algorithm with the same time bounds, but using only $\mathcal{O}\left(n^{d-1}\right)$ space, and incurring only $\mathcal{O}\left(\frac{n^d}{BM^{\frac{1}{d-1}}}\right)$ (when $M = \Omega\left(B^{d-1}\right)$). The CMP algorithm for LDDP is a fairly straightforward extension of the LCS algorithm.

LDDP includes many important problems, especially in bioinformatics, including *longest common subsequence (LCS)*, *global pairwise sequence alignment (PA)* and *Median* (both with affine gap costs), and *RNA secondary structure prediction with simple pseudoknots (RNA-SP).* It also includes local alignment, generalized global alignment with intermittent similarities, multiple sequence alignment under several scoring functions such as 'sum-of-pairs' objective function and RNA secondary structure prediction with simple pseudoknots using energy functions based on adjacent base pairs.

A.1 Pairwise Global Sequence Alignment with Affine Gap Penalty. Sequence alignment plays a central role in biological sequence comparison, and can reveal important relationships among organisms. Given two strings $X = x_1x_2 \ldots x_m$ and $Y = y_1y_2 \ldots y_n$ over a finite alphabet Σ, an *alignment* of X and Y is a matching M of sets $\{1, 2, \ldots, m\}$ and $\{1, 2, \ldots, n\}$ such that if $(i, j), (i', j') \in M$ and $i < i'$ hold then $j < j'$ must also hold [23]. The i-th letter of X or Y is said to be in a *gap* if it does not appear in any pair in M. Given a *gap penalty* g and a mismatch cost $s(a, b)$ for each pair $a, b \in \Sigma$, the *basic (global) pairwise sequence alignment problem* asks for a matching M_{opt} for which $(m+n-|M_{opt}|) \times g + \sum_{(a,b)\in M_{opt}} s(a, b)$ is minimized [23].

For simplicity we will assume $m = n$. The formulation of the basic sequence alignment problem favors a large number of small gaps while real biological processes favor the opposite. The alignment can be made more realistic by using an *affine gap penalty* [18, 1] which has two parameters: a *gap introduction cost* g_i and a *gap extension cost* g_e. A run of k gaps incurs a total cost of $g_i + g_e \times k$.

In [18] Gotoh presented an $\mathcal{O}\left(n^2\right)$ time and $\mathcal{O}\left(n^2\right)$ space DP algorithm for solving the global pairwise alignment problem with affine gap costs. The algorithm incurs $\mathcal{O}\left(\frac{n^2}{B}\right)$ cache misses. The space complexity of the algorithm can be reduced to $\mathcal{O}(n)$ using Hirschberg's space-reduction technique [26] or the diagonal checkpointing technique described in [19]. However, the time and cache complexities remain unchanged. Gotoh's algorithm solves the following DP recurrences.

$$D(i,j) = \begin{cases} G(0,j) + g_e & \text{if } i = 0 \wedge j > 0 \\ \min\left\{\begin{matrix} D(i-1,j), \\ G(i-1,j) + g_i \end{matrix}\right\} + g_e & \text{if } i > 0 \wedge j > 0. \end{cases}$$

$$I(i,j) = \begin{cases} G(i,0) + g_e & \text{if } i > 0 \wedge j = 0 \\ \min\left\{\begin{matrix} I(i,j-1), \\ G(i,j-1) + g_i \end{matrix}\right\} + g_e & \text{if } i > 0 \wedge j > 0. \end{cases}$$

$$G(i,j) = \begin{cases} 0 & \text{if } i = 0 \wedge j = 0 \\ g_i + g_e \times j & \text{if } i = 0 \wedge j > 0 \\ g_i + g_e \times i & \text{if } i > 0 \wedge j = 0 \\ \min \left\{ \begin{array}{c} D(i,j), I(i,j), \\ G(i-1,j-1) \\ +s(x_i, y_j) \end{array} \right\} & \text{if } i > 0 \wedge j > 0. \end{cases}$$

The optimal alignment cost is $\min\{G(n,n), D(n,n), I(n,n)\}$ and an optimal alignment can be traced back from the smallest of $G(n,n)$, $D(n,n)$ and $I(n,n)$.

Connection to 2-dimensional LDDP. The three recurrences above can be viewed as a single recurrence evaluating a single matrix $c[0:n, 0:n]$ with three fields: D, I and G. When $i = 0$ or $j = 0$ each field of $c[i,j]$ depends only on the indices i and j, and constants g_i and g_e, and hence each such entry can be computed using a function similar to h in the general recurrence A.4 in Section A. When both i and j are positive, $c[i,j]$ depends on x_i, y_j, the entries in $c[i-1:i, j-1:j] \setminus c[i,j]$ (i.e., in $D(i-1:i, j-1:j) \setminus D(i,j)$, $I(i-1:i, j-1:j) \setminus I(i,j)$ and $G(i-1:i, j-1:j) \setminus G(i,j)$,), and constants g_i and g_e. Hence, in this case $c[i,j]$ can be computed using a function similar to function f in recurrence A.4. Thus the recurrences used by Gotoh's algorithm completely match the general recurrence A.4 for $d = 2$.

A.2 Median of Three Sequences. The Median problem is the problem of obtaining an optimal alignment of three sequences using an affine gap penalty. The median sequence under the optimal alignment is also computed. Knudsen [25] presented a dynamic program to find multiple alignment of N sequences, each of length n in $\mathcal{O}\left(16.81^N n^N\right)$ time and $\mathcal{O}\left(7.442^N n^N\right)$ space. For the median problem, this gives an $\mathcal{O}\left(n^3\right)$ time and space algorithm that incurs $\mathcal{O}\left(\frac{n^3}{B}\right)$ cache-misses. An Ukkonen-based algorithm is presented in [29], which performs well especially for sequences whose (3-way) edit distance δ is small. On average, it requires $\mathcal{O}\left(n + \delta^3\right)$ time and space [29].

Knudsen's Algorithm [25] for three sequences (say, $X = x_1 x_2 \ldots x_n$, $Y = y_1 y_2 \ldots y_n$ and $Z = z_1 z_2 \ldots, z_n$) is a dynamic program over a three-dimensional matrix K. Each entry $K(i,j,k)$ is composed of 23 fields. Each field corresponds to an indel configuration q, which describes how the last characters x_i, y_j and z_k are matched. A residue configuration defines how the next three characters of the sequences will be matched. Each configuration is a vector $e = (e_1, e_2, e_3, e_4)$, where $e_i \in \{0,1\}$, $1 \le i \le 4$. The entry e_i, $1 \le i \le 3$ indicates if the aligned character of sequence i is a gap or a residue, while e_4 corresponds to the aligned character of the median sequence. There are 10 residue configurations out of 16 possible ones. The recursive step calculates the value of the next entry by applying residue configurations to each indel configuration. We define $\nu(e,q) = q'$ if applying the residue configuration e to the indel configuration q gives the indel configuration q'. The recurrence relation used by Knudsen's algorithm is:

$$K(i,j,k)_q = \begin{cases} 0 & \text{if } i = j = k = 0 \wedge q = q_o \\ \infty & \text{if } i = j = k = 0 \wedge q \neq q_o \\ \min\limits_{e,q' : q = \nu(e,q')} \left\{ \begin{array}{c} K(i',j',k')_{q'} \\ +G_{e,q} \\ +M_{(i',j',k') \to (i,j,k)} \end{array} \right\} & \text{otherwise.} \end{cases}$$

where q_o is the configuration where all characters match, $i' = i - e_1$, $j' = j - e_2$ and $k' = k - e_3$, $M_{(i',j',k') = (i,j,k)}$ is the matching cost between characters of the sequences, and $G_{e,q}$ is the cost of introducing or extending the gap.

The M and G matrices can be pre-computed. Therefore, Knudsen's algorithm runs in $\mathcal{O}\left(n^3\right)$ time and space with $\mathcal{O}\left(\frac{n^3}{B}\right)$ cache-misses.

Connection to 3-Dimensional LDDP. In order to make the recurrence used by Knudsen's algorithm match the general recurrence A.4 for $d = 3$ given in Section A, we shift all characters of X, Y and Z one position to the right, introduce a dummy character in front of each of those three sequences, and obtain the following recurrence.

$$c[i,j,k]_q = \begin{cases} \infty & \text{if } i = 0 \vee j = 0 \vee k = 0 \\ 0 & \text{if } i = j = k = 1 \wedge q = q_o \\ \infty & \text{if } i = j = k = 1 \wedge q \neq q_o \\ \min\limits_{e,q' : q = \nu(e,q')} \left\{ \begin{array}{c} c[i',j',k']_{q'} \\ +G_{e,q} \\ +M_{(i',j',k') \to (i,j,k)} \end{array} \right\} & \text{otherwise.} \end{cases}$$

It is easy to see that $K(i,j,k)_q = c[i+1, j+1, k+1]_q$ for $0 \le i,j,k \le n$ and any q. If $i = 0$ or $j = 0$ or $k = 0$ then $c[i,j,k]_q$ can be evaluated using a function $h(\ \langle\ i,\ j,\ k\ \rangle\ ) = \infty$ as in the general recurrence A.4. Otherwise the value of $c[i,j,k]_q$ depends on the values of i, j, and k, values in some constant size arrays (G and M), and on the cells to its left, back and top. Hence, in this case, $c[i,j,k]_q$ can be evaluated using a function similar to f in recurrence A.4 for $d = 3$. Therefore, the above recurrence matches the 3-dimensional version of the general recurrence A.4.

B Pseudocode for CMP Algorithms

Tiled-Boundary-LCS(X', Y', d)
Input. X' and Y' are substrings of LCS inputs X and Y, respectively, with $|X'| = |Y'|$. We let $Q[1..s, 1..s]$ be the (virtual) subsquare of cost table c for pairs of positions falling in X' and Y' (where $s = |X'|$). The c values at the input boundary of Q are available (the *input boundary* consists of the portions of the row and column of c just above and to the left of Q).
The third parameter d is a non-negative integer that gives the recursion level for this call, hence this call with use *tiling parameter* $t[d]$ (see text for more details).
Output. The c values at the *output boundary* of Q (i.e., at $Q[i,s]$ and $Q[s,j]$, $1 \le i,j \le s$), and for each such position, the entry point of its traceback path on input boundary of Q. Space usage is $O(s)$.

if $s = 1$ ***then*** directly compute the output values for $Q[1,1]$ in constant time ***else***
  Let $\tau := t[d]$ (we assume for convenience that τ divides s); we divide X' and Y' into τ equal-sized
    pieces X'_i and Y'_j, hence we tile Q with τ^2 subsquares $Q_{i,j}$ where, for $1 \le i,j \le \tau$,
    $Q_{i,j} = Q[(i-1)(s/\tau)+1\ ..\ is/\tau,\ (j-1)(s/\tau)+1\ ..\ js/\tau]$
  for $1 \le k \le 2\tau - 1$ ***do***
    pfor all pairs (i,j) with $i + j = k + 1$ ***do*** Tiled-Boundary-LCS$(X'_i, Y'_j, d+1)$

Tiled-LCS(X', Y', u, d)
Input/Output. X', Y', d are same as above; u is a position on output boundary of Q from which a traceback path is required.
Output. The traceback path from u to the input boundary of Q, together with the LCS cost at u. The space usage is $O(s)$.

if $s = 1$ ***then*** directly compute value of $Q[1,1]$, and its parent in the c-table in constant time ***else***
  Let $\tau := t[d]$; divide each of X', Y' into two equal halves s.t. $X' =< X'_1, X'_2 >$, $Y' =< Y'_1, Y'_2 >$,
    and let $Q_{i,j}[1..s/2,\ 1..s/2]$ be the (virtual) Q table for X'_i and Y'_j, $1 \le i,j \le 2$
1. Call Tiled-Boundary-LCS with (X'_1, Y'_1, d), then (X'_1, Y'_2, d) and (X'_2, Y'_1, d), then (X'_2, Y'_2, d)
2. Using the entry points of the traceback paths in the $Q_{i,j}$, identify the set Σ of those triples X'_i, Y'_j, v such that
    the traceback path from u enters $Q_{i,j}$ at v (by starting at $Q_{2,2}$ and working backward) in $O(1)$ time
3. ***pfor*** each of the (at most 3) $S \in \Sigma$ ***do*** Tiled-LCS$(S, d+1)$

Figure 7: CMP algorithms for LCS. The initial call is Tiled-LCS(X, Y, (n,n), 0), where X and Y are the input sequences, both of length n.

A(d, X, U, V, W)

(Each of X, U, V and W points to the same $m \times m$ square submatrix of the $n \times n$ input matrix c. We assume for simplicity that r divides m (if $m > 1$), where $r = t[d]$. The initial call to A is $A(0, c, c, c, c)$.)

1. ***if*** $T_{XUV} \cap \Sigma_G = \emptyset$ ***then return*** $\{T_{XUV} = \{$ *updates on* X *using* $(i,k) \in U$ *and* $(k,j) \in V$ $\}$, *and* Σ_G *is the set of updates performed by iterative GEP*$\}$

2. ***if*** X is a 1×1 matrix ***then*** $X \leftarrow f(X, U, V, W)$

else

3. $r \leftarrow t[d]$

Split each of X, U, V and W into $r \times r$ square submatrices of size $\frac{n}{r} \times \frac{n}{r}$ each. The submatrix of X (or U or V or W) at the i-th position from the top and the j-th position from the left is denoted by X_{ij} (or U_{ij} or V_{ij} or W_{ij}, resp.).

4. ***for*** $k \leftarrow 1$ ***to*** r ***do***

5. $A(\, d+1,\ X_{kk},\ U_{kk},\ V_{kk},\ W_{kk}\,)$

6. ***parallel*** : $B_{1+|j<k|}(\, d+1,\ X_{kj},\ U_{kk},\ V_{kj},\ W_{kk}\,),\ j \in [1,r],\ j \neq k$

$C_{1+|i<k|}(\, d+1,\ X_{ik},\ U_{ik},\ V_{kk},\ W_{kk}\,),\ i \in [1,r],\ i \neq k$

7. ***parallel*** : $D_{1+2|i<k|+|j<k|}(\, d+1,\ X_{ij},\ U_{ik},\ V_{kj},\ W_{kk}\,),\ i,j \in [1,r],\ i \neq k,\ j \neq k$

B_l(d, X, U, V, W) $\{\, l \in \{1,2\}\, \}$

($X \equiv V \equiv c[i_1..i_2, j_1..j_2]$ and $U \equiv W \equiv c[i_1..i_2, k_1..k_2]$, where $i_2 - i_1 = j_2 - j_1 = k_2 - k_1$, $[i_1, i_2] = [k_1, k_2]$ and $[j_1, j_2] \cap [k_1, k_2] = \emptyset$.)

1–3. As in steps 1–3 of A

4. ***for*** $k \leftarrow 1$ ***to*** r ***do***

5. ***parallel*** : $B_l(\, d+1,\ X_{kj},\ U_{kk},\ V_{kj},\ W_{kk}\,)$

$j \in [1,r]$

6. ***parallel*** : $D_{l+2|i<k|}(d+1,\ X_{ij},\ U_{ik},\ V_{kj},\ W_{kk})$

$i,j \in [1,r],\ i \neq k$

C_l(d, X, U, V, W) $\{\, l \in \{1,2\}\, \}$

($X \equiv U \equiv c[i_1..i_2, j_1..j_2]$ and $V \equiv W \equiv c[k_1..k_2, j_1..j_2]$, where $i_2 - i_1 = j_2 - j_1 = k_2 - k_1$, $[j_1, j_2] = [k_1, k_2]$ and $[i_1, i_2] \cap [k_1, k_2] = \emptyset$.)

1–3. As in steps 1–3 of A

4. ***for*** $k \leftarrow 1$ ***to*** r ***do***

5. ***parallel*** : $C_l(\, d+1,\ X_{ik},\ U_{ik},\ V_{kk},\ W_{kk}\,)$

$i \in [1,r]$

6. ***parallel*** : $D_{2l-|j>k|}(d+1,\ X_{ij},\ U_{ik},\ V_{kj},\ W_{kk})$

$i,j \in [1,r],\ j \neq k$

D_l(d, X, U, V, W) $\{\, l \in \{1,2,3,4\}\, \}$

($X \equiv c[i_1..i_2, j_1..j_2]$, $U \equiv c[i_1..i_2, k_1..k_2]$, $V \equiv c[k_1..k_2, j_1..j_2]$ and $W \equiv c[k_1..k_2, k_1..k_2]$, where $i_2 - i_1 = j_2 - j_1 = k_2 - k_1$, $[i_1, i_2] \cap [k_1, k_2] = \emptyset$, and $[j_1, j_2] \cap [k_1, k_2] = \emptyset$.)

1–3. As in steps 1–3 of A

4. ***for*** $k \leftarrow 1$ ***to*** r ***do***

5. ***parallel*** : $D_l(\, d+1,\ X_{ij},\ U_{ik},\ V_{kj},\ W_{kk}\,),\ i,j \in [1,r]$

Figure 8: CMP Implementation of I-GEP. Initial call is A(0, c, c, c, c) on an $n \times n$ matrix c.

E(d, X)

(X points to a $m \times m$ square submatrix of the $n \times n$ input matrix c such that the top-left to bottom-right diagonal of X lies on the top-left to bottom-right diagonal of c. We assume for simplicity that r divides m (if $m > 1$), where $r = t[d]$. The initial call to A is E$(0, c)$.)

1. ***if*** X is not a 1×1 matrix ***then***
2. $\quad r \leftarrow t[d]$
   $\quad$ Split X into $r \times r$ square submatrices of size $\frac{n}{r} \times \frac{n}{r}$ each. The submatrix of X at the i-th position from the top and the j-th position from the left is denoted by $X_{i,j}$.
3. $\quad$ ***parallel*** : E($d+1$, $X_{i,j}$), $i, j \in [1, r]$, $i + j = r + 1$
4. $\quad$ ***for*** $k \leftarrow 1$ ***to*** $r - 1$ ***do***
5. $\quad\quad$ ***parallel*** : H($d+1$, $X_{i,j}$, $X_{i,r+k-i}$, $X_{i-k+1,j}$), $i, j \in [1, r]$, $r + k < i + j < r + 2k$
6. $\quad\quad$ ***parallel*** : H($d+1$, $X_{i,j}$, $X_{i,j-k+1}$, $X_{r+k-j,j}$), $i, j \in [1, r]$, $r + k < i + j < r + 2k - 1$
7. $\quad\quad$ ***parallel*** : F($d+1$, $X_{i,j}$, $X_{i,r-i+1}$, $X_{r-j+1,j}$), $i, j \in [1, r]$, $i + j = r + k + 1$

F(d, X, U, V)

(X, U and V are $m \times m$ disjoint submatrices of the $n \times n$ input matrix c such that X lies entirely above the top-left to bottom-right diagonal of c, and the top-left to bottom-right diagonals of U and V lie on that of c. The submatrix U lies to the left of submatrix X, and both lie on the same rows of c. The submatrix V lies below submatrix X, and both lie on the same columns of c. We assume for simplicity that r divides m (if $m > 1$), where $r = t[d]$.)

1. ***if*** X is a 1×1 matrix ***then***
2. $\quad X \leftarrow \min\{ X, (U + V) + W \}$, where $W = w(i, k, j)$ assuming $X \equiv c[i, j]$, $U \equiv c[i, k]$ and $V \equiv c[k, j]$

   else
3. $\quad r \leftarrow t[d]$
   $\quad$ Split each of X, U and V into $r \times r$ square submatrices of size $\frac{n}{r} \times \frac{n}{r}$ each. The submatrix of X (or U or V) at the i-th position from the top and the j-th position from the left is denoted by $X_{i,j}$ (or $U_{i,j}$ or $V_{i,j}$, resp.).
   $\quad$ For $i, j \in [1, r]$, we define the following two mappings.

$$U'_{i,l} = \begin{cases} X_{i,l} & \text{if } l > 0, \\ U_{i,l+r} & \text{otherwise.} \end{cases} \qquad V'_{l,j} = \begin{cases} X_{l,j} & \text{if } l > 0, \\ V_{l+r,j} & \text{otherwise.} \end{cases}$$

4. $\quad$ ***for*** $k \leftarrow 1$ ***to*** r ***do***
5. $\quad\quad$ ***parallel*** : H($d+1$, $X_{i,j}$, $U'_{i,k-i}$, $V'_{i-k+1,j}$), $i, j \in [1, r]$, $k < i + j < 2k$
6. $\quad\quad$ ***parallel*** : H($d+1$, $X_{i,j}$, $U'_{i,j-k+1}$, $V'_{k-j,j}$), $i, j \in [1, r]$, $k < i + j < 2k - 1$
7. $\quad\quad$ ***parallel*** : F($d+1$, $X_{i,j}$, $U'_{i,1-i}$, $V'_{1-j,j}$), $i, j \in [1, r]$, $i + j = k + 1$

H(d, X, U, V)

(X, U and V are $m \times m$ disjoint submatrices of the $n \times n$ input matrix c each of which lies entirely above the top-left to bottom-right diagonal of c. The submatrix U lies to the left of submatrix X, and both lie on the same rows of c. The submatrix V lies below submatrix X, and both lie on the same columns of c. We assume for simplicity that r divides m (if $m > 1$), where $r = t[d]$.)

1–3. As in steps 1–3 of F

4. $\quad$ ***for*** $k \leftarrow 1$ ***to*** r ***do***
5. $\quad\quad$ ***parallel*** : H($d+1$, $X_{i,j}$, $U_{i,k}$, $V_{k,j}$), $i, j \in [1, r]$

Figure 9: CMP algorithm for the parenthesis problem defined by recurrence 4.3. Initial call is E(0, c) for an $n \times n$ input matrix c.