Cache Oblivious Algorithms for Computing the Triplet Distance Between Trees[*]

Gerth Stølting Brodal[1] and Konstantinos Mampentzidis[2]

1   **Department of Computer Science, Aarhus University, Aarhus, Denmark**
    `gerth@cs.au.dk`
2   **Department of Computer Science, Aarhus University, Aarhus, Denmark**
    `kmampent@cs.au.dk`

⎯ Abstract ⎯

We consider the problem of computing the triplet distance between two rooted unordered trees with n labeled leaves. Introduced by Dobson in 1975, the triplet distance is the number of leaf triples that induce different topologies in the two trees. The current theoretically fastest algorithm is an $O(n \log n)$ algorithm by Brodal *et al.* (SODA 2013). Recently Jansson and Rajaby proposed a new algorithm that, while slower in theory, requiring $O(n \log^3 n)$ time, in practice it outperforms the theoretically faster $O(n \log n)$ algorithm. Both algorithms do not scale to external memory.

We present two cache oblivious algorithms that combine the best of both worlds. The first algorithm is for the case when the two input trees are binary trees, and the second is a generalized algorithm for two input trees of arbitrary degree. Analyzed in the RAM model, both algorithms require $O(n \log n)$ time, and in the cache oblivious model $O(\frac{n}{B} \log_2 \frac{n}{M})$ I/Os. Their relative simplicity and the fact that they scale to external memory makes them achieve the best practical performance. We note that these are the first algorithms that scale to external memory, both in theory and in practice, for this problem.

1998 ACM Subject Classification G.2.2 Trees, G.2.1 Combinatorial Algorithms

Keywords and phrases Phylogenetic tree, tree comparison, triplet distance, cache oblivious algorithm

1   Introduction

Trees are data structures that are often used to represent relationships. For example in the field of Biology, a tree can be used to represent evolutionary relationships, with the leaves corresponding to species that exist today, and internal nodes to ancestor species that existed in the past. For a fixed set of n species, different data (e.g., DNA, morphological) or construction methods (e.g., Q* [3], neighbor joining [14]) can lead to trees that look structurally different. An interesting question that arises then is, given two trees T_1 and T_2 over n species, how different are they? An answer to this question could potentially be used to determine whether the difference is statistically significant or not, which in turn could help with evolutionary inferences.

Several distance measures have been proposed in the past to compare two trees that are *unordered*, i.e., trees in which the order of the siblings is not taken into account. A class of them includes distance measures that are based on how often certain features are different in the two trees. Common distance measures of this kind are the Robinson-Foulds distance [13],
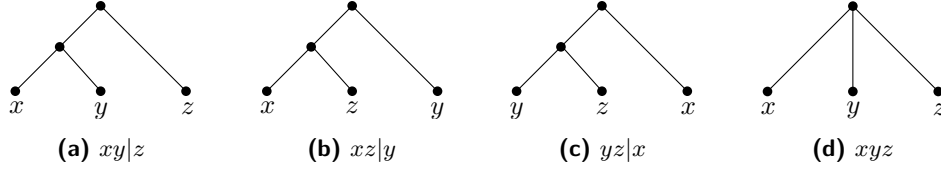
Figure 1 All possible topologies of a triplet with leaves x, y, and z.

the triplet distance [7] for rooted trees and the quartet distance [8] for unrooted trees. The Robinson-Foulds distance counts how many leaf bipartitions are different, where a bipartition in a given tree is generated by removing a single edge from the tree. The triplet distance is only defined for rooted trees, and counts how many leaf triples induce different topologies in the two trees. The counterpart of the triplet distance for unrooted trees, is the quartet distance, which counts how many leaf quadruples induce different topologies in the two trees.

Algorithms exist that can efficiently compute these distance measures. The Robinson-Foulds distance can be optimally computed in $O(n)$ time [6]. The triplet distance can be computed in $O(n \log n)$ time [4]. The quartet distance can be computed in $O(dn \log n)$ time [4], where d is the maximal degree of any node in the two input trees.

The above bounds are in the RAM model. Previous work did not consider any other models, for example external memory models like the I/O model [1] and the cache oblivious model [9]. Typically when hearing about algorithms for external memory models, one might (sometimes incorrectly) think of only algorithms that have to deal with large amounts of data. Hence, any practical improvement that comes from an algorithm that scales to external memory compared to an equivalent that does not, can only be noticed if the inputs are large. However, this is not necessarily the case for cache oblivious algorithms. A cache oblivious algorithm, if built and implemented correctly, can take advantage of the L1, L2, and L3 caches that exist in the vast majority of computers and give a significant performance improvement even for small inputs.

A trivial modification of the algorithm in [6], can give a cache oblivious algorithm for computing the Robinson-Foulds distance that achieves the sorting bound, by requiring $O(\frac{n}{B} \log_{\frac{M}{B}} \frac{n}{M})$ I/Os instead of $O(n)$ I/Os for the standard implementation. For the triplet and quartet distance measures, no such trivial modifications exist.

In this paper we focus on the triplet distance computation and present the first non-trivial algorithms for computing the triplet distance between two rooted trees, that for the first time for this problem, also scale to external memory.

Problem Definition.   For a given rooted unordered tree T where each leaf has a unique label, a *triplet* is defined by a set of three leaf labels x, y, and z and their induced topology in T. The four possible topologies are illustrated in Figure 1. The notation $xy|z$ is used to describe a triplet where the lowest common ancestor of x and y is at a lower depth than the lowest common ancestor of z with either x or y. Note that the triplet $xy|z$ is the same as the triplet $yx|z$ because T is considered to be unordered. Similarly, notation xyz is used to describe a triplet for which every pair of leaves has the same lowest common ancestor. This triplet can only appear if we allow nodes with degree three or larger in T. From here on, when using the word "tree" we imply a "rooted unordered tree".

For two given trees T_1 and T_2 that are built on n identical leaf labels, the *triplet distance* $D(T_1, T_2)$ is the number of triplets the leaves of which induce different topologies in T_1 and T_2. Let $S(T_1, T_2)$ be the number of *shared* triplets in the two trees, i.e., leaf triples with identical

topologies in the two trees. We then have the relationship $D(T_1, T_2) + S(T_1, T_2) = \binom{n}{3}$.

Previous and new results for computing the triplet distance are shown in the table below. Note that the papers [5, 2, 15, 4, 11] do not provide an analysis of the algorithms in the cache oblivious model, so here we provide an upper bound. From here on and unless otherwise stated, any asymptotic bound refers to time.

Year	Reference	Time	IOs	Space	Non-Binary Trees
1996	Critchlow *et al.* [5]	$O(n^2)$	$O(n^2)$	$O(n^2)$	no
2011	Bansal *et al.* [2]	$O(n^2)$	$O(n^2)$	$O(n^2)$	yes
2013	Sand *et al.* [15]	$O(n \log^2 n)$	$O(n \log^2 n)$	$O(n)$	no
2013	Brodal *et al.* [4]	$O(n \log n)$	$O(n \log n)$	$O(n \log n)$	yes
2015	Jansson and Rajaby [11]	$O(n \log^3 n)$	$O(n \log^3 n)$	$O(n \log n)$	yes
2017	new	$O(n \log n)$	$O(\frac{n}{B} \log_2 \frac{n}{M})$	$O(n)$	yes

Related Work.   The triplet distance was first suggested as a method of comparing the shapes of trees by Dobson in 1975 [7]. The first non-trivial algorithmic result dates back to 1996, when Critchlow *et al.* [5] proposed an $O(n^2)$ algorithm that however works only for binary trees. Bansal *et al.* [2] introduced an $O(n^2)$ algorithm that works for general (binary and non-binary) trees. Both of these algorithms use $O(n^2)$ space. Sand *et al.* [15] introduced a new $O(n^2)$ algorithm using only $O(n)$ space for the case of binary trees, that they showed how to optimize to reduce the time to $O(n \log^2 n)$. This algorithm was also implemented and shown to be the most efficient in practice. Soon after, Brodal *et al.* [4] managed to extend the $O(n \log^2 n)$ algorithm to work for general trees, and at the same time brought the time down to $O(n \log n)$ but now with the space increased to $O(n \log n)$. The space for binary trees was still $O(n)$. The algorithms from [15] and [4] were implemented and added to the library tqDist [16]. Interestingly, it was shown in [10] that for binary trees the $O(n \log^2 n)$ algorithm had a better practical performance than the $O(n \log n)$ algorithm. Jansson and Rajaby [11, 12] showed that an even slower theoretically algorithm requiring worst case $O(n \log^3 n)$ time and $O(n \log n)$ space could give the best practical performance, both for binary and non-binary trees. A survey of previous results until 2013 can be found in [17].

Contribution.   The common main bottleneck with all previous approaches is that the data structures used rely intensively on $\Omega(n \log n)$ random memory accesses. This means that all algorithms are penalized by cache performance and thus do not scale to external memory. We address this limitation by proposing new algorithms for computing the triplet distance on binary and non-binary trees, that match the previous best $O(n \log n)$ time and $O(n)$ space bounds in the RAM model, but for the first time also scale to external memory. More specifically, in the cache oblivious model, the total number of I/Os required is $O(\frac{n}{B} \log_2 \frac{n}{M})$. The basic idea is to essentially replace the dependency of random access to data structures by scanning contracted versions of the input trees. A careful implementation of the algorithms is shown to achieve the best performance in practice, thus essentially documenting that the theoretical results carry over to practice.

Outline of the Paper.   In Section 2 we provide an overview of previous approaches. In Section 3 we describe the new algorithm for the case where T_1 and T_2 are binary trees. In Section 4 we extend the algorithm to also work for general trees. In Section 5 we provide some implementation details. Section 6 contains our experimental evaluation. Appendix A contains more experimental results. Finally, in Section 7 we provide our concluding remarks.

2   Previous Approaches

A naive approach would enumerate over all $\binom{n}{3}$ sets of three leaf labels and find for each set whether the induced topologies in T_1 and T_2 differ or not, giving an $O(n^3)$ algorithm. This algorithm does not exploit the fact that the triplets are not completely independent. For example, the triplets $xy|z$ and $yx|u$ share the leaves x and y and the fact that the lowest common ancestor of x and y is at a lower depth than the lowest common ancestor of z with either x or y and the lowest common ancestor of u with either x or y. Dependencies like this can be exploited to count the number of shared triplets faster.

Critchlow *et al.* [5] exploit the depth of the leaves' ancestors to achieve the first improvement over the naive approach. Bansal *et al.* [2] exploit the shared leaves between subtrees and reduce the problem to computing the intersection size (number of shared leaves) of all pairs of subtrees, one from T_1 and one from T_2, which can be solved with dynamic programming.

The $O(n^2)$ Algorithm for Binary Trees in [15].   The algorithm for binary trees in [15] is the basis for all subsequent improvements [15, 4, 11], including ours as well, so we will describe it in more detail here. The dependency that was exploited is the same as in [2] but the procedure for counting the shared triplets is completely different. More specifically, each triplet in T_1 and T_2, defined by three leaf labels i, j, and k, is implicitly *anchored* in the lowest common ancestor of i, j, and k. For two nodes u in T_1 and v in T_2, let $s(u)$ and $s(v)$ be the set of triplets that are anchored in u and v respectively. For the number of shared triplets $S(T_1, T_2)$ we then have:

$$S(T_1, T_2) = \sum_{u \in T_1} \sum_{v \in T_2} |s(u) \cap s(v)| \ .$$

For the algorithm to be $O(n^2)$ the value $|s(u) \cap s(v)|$ must be computed in $O(1)$ time. This is achieved by a leaf colouring procedure as follows: Fix an internal node u in T_1 and color the leaves in the left subtree of u *red*, the leaves in the right subtree of u *blue*, let every other leaf have no color and then transfer this coloring to the leaves in T_2, i.e., identically labeled leaves get the same color. The triplets anchored at u are exactly the triplets $xy|z$ where x, y are blue and z is red, or x, y are red and z is blue. To compute $|s(u) \cap s(v)|$ we do as follows: let l and r be the left and right children of v, and let $w_{\texttt{red}}$ and $w_{\texttt{blue}}$ be the number of red and blue leaves in a subtree rooted at a node w in T_2. We then have:

$$|s(u) \cap s(v)| = \binom{l_{\texttt{red}}}{2} r_{\texttt{blue}} + \binom{l_{\texttt{blue}}}{2} r_{\texttt{red}} + \binom{r_{\texttt{red}}}{2} l_{\texttt{blue}} + \binom{r_{\texttt{blue}}}{2} l_{\texttt{red}} \ . \tag{1}$$

Subquadratic Algorithms.   To reduce the time, Sand *et al.* [15] applied the *smaller half trick*, which specifies a depth first order to visit the nodes u of T_1, so that each leaf in T_1 changes color at most $O(\log n)$ times. To count shared triplets efficiently without scanning T_2 completely for each node u in T_1, the tree T_2 is stored in a data structure denoted a *hierarchical decomposition tree* (*HDT*). This HDT of T_2 maintains for the current visited node u in T_1, according to (1) the sum $\sum_{v \in T_2} |s(u) \cap s(v)|$, so that each leaf color change in T_1 can be updated efficiently in T_2. In [15] the HDT is a binary tree of height $O(\log n)$ and every update can be done by a leaf to root path traversal in the HDT, which in total gives $O(n \log^2 n)$ time. In [4] the HDT is generalized to also handle non-binary trees, each query operates the same, and now due to a contraction scheme of the HDT the total time is reduced to $O(n \log n)$. Finally, in [11] as an HDT the so called *heavy-light tree decomposition*

is used. Note that the only difference between all $O(n \operatorname{polylog} n)$ results that are available right now is the type of HDT used.

In terms of external memory efficiency, every $O(n \operatorname{polylog} n)$ algorithm performs $\Theta(n \log n)$ updates to an HDT data structure, which means that for sufficiently large input trees every algorithm requires $\Omega(n \log n)$ I/Os.

3 The New Algorithm for Binary Trees

In this section, we provide a cache oblivious algorithm that for two binary trees T_1 and T_2, built on the same leaf label set of size n, computes $D(T_1, T_2)$ using $O(n \log n)$ time and $O(n)$ space in the RAM model, and $O(\frac{n}{B} \log_2 \frac{n}{M})$ I/Os in the cache oblivious model.
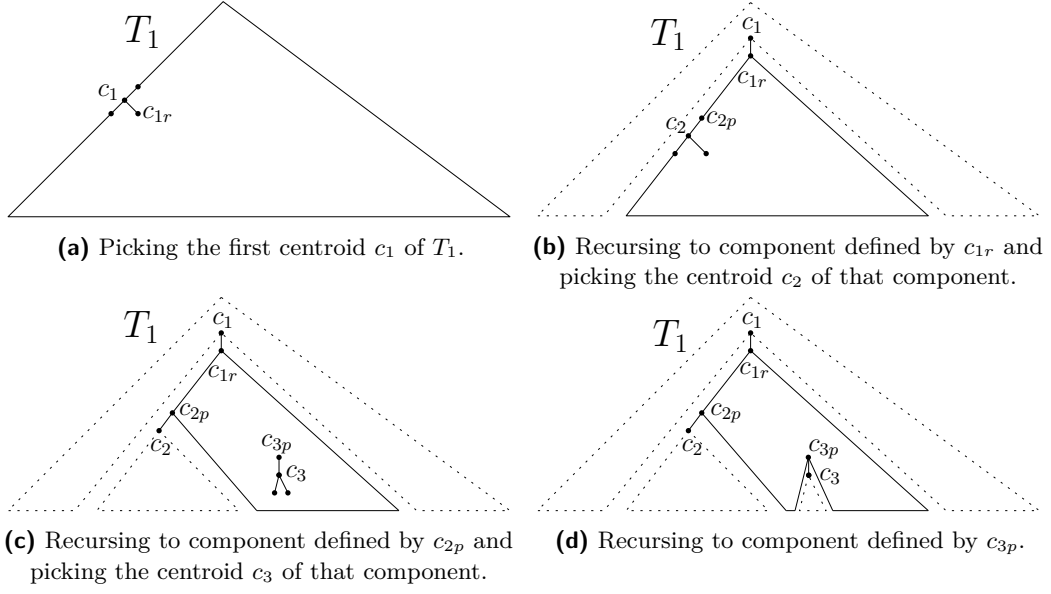
Overview. We use the $O(n^2)$ algorithm from Section 2 as a basis. The main difference between this algorithm and our new algorithm is in the order that we visit the nodes of T_1, and how we process T_2 when we count. We propose a new order of visiting the nodes of T_1, which is found by applying a hierarchical decomposition on T_1. Every component in this decomposition corresponds to a connected part of T_1 and a contracted version of T_2. In simple terms, if Λ is the set of leaves in a component of T_1, the contracted version of T_2 is a binary tree on Λ that preserves the topologies induced by Λ in T_2 and has size $O(|\Lambda|)$. To count shared triplets, every component of T_1 has a representative node u that we use to scan the corresponding contracted version of T_2 in order to find $\sum_{v \in T_2} |s(u) \cap s(v)|$. Unlike previous algorithms, we do not store T_2 in a data structure. We process T_2 by contracting and counting, both of which can be done by scanning. At the same time, even though we apply a hierarchical decomposition on T_1, the only reason why we do so, is so we can find the order in which to visit the nodes of T_1. This means that we do not need to store T_1 in a data structure either. Hence, we completely remove the need of data structures (and thereby random memory accesses), and scanning becomes the basic primitive in the algorithm. To make our algorithm I/O efficient, all that remains to be done is to use a proper layout to store the contracted trees in memory, so that every time we scan a tree of size s we spend $O(s/B)$ I/Os.

3.1 Modified Centroid Decomposition

For a given binary tree T let $|T|$ denote the number of nodes in T (internal nodes and leaves). For a node u in T let l and r be the left and right children of u, and p the parent of u. Removing u from T partitions T into three (possibly empty) *connected components* T_l, T_r, and T_p containing l, r, and p, respectively. A *centroid* is a node u in T such that $\max\{|T_l|, |T_r|, |T_p|\} \le |T|/2$. A centroid always exists and can be found by starting from the root of T and iteratively visiting the child with a largest subtree, eventually we will reach a centroid. Finding the size of every subtree and identifying u takes $O(|T|)$ time in the RAM model. By recursively finding centroids in each of the three components, we in the end get a ternary tree of centroids, which is called the *centroid decomposition* of T, denoted $CD(T)$. We can generate a level of $CD(T)$ in $O(|T|)$ time, given the decomposition of T into components by the previous level. Since $CD(T)$ has at most $1 + \log_2(|T|)$ levels, the total time required to build $CD(T)$ is $O(|T| \log |T|)$, thus we get Lemma 1.

▶ **Lemma 1.** *For any given binary tree T with n leaves, there exists an algorithm that builds $CD(T)$ using $O(n \log n)$ time and $O(n)$ space in the RAM model.*

A component in a centroid decomposition $CD(T)$, might have several edges crossing its boundaries (connecting nodes inside and outside the component). An example of creating a

(a) Picking the first centroid c_1 of T_1.

(b) Recursing to component defined by c_{1r} and picking the centroid c_2 of that component.

(c) Recursing to component defined by c_{2p} and picking the centroid c_3 of that component.

(d) Recursing to component defined by c_{3p}.

Figure 2 Generating a component in $CD(T_1)$ that has two edges from below. The black polygon is the component.

component that has two edges from below can be found in Figure 2. It is trivial to see that by following the same pattern of generating components as depicted in Figure 2d, $CD(T)$ can have a component with an arbitrary number of edges from below. The below *modified centroid decomposition*, denoted $MCD(T)$, generates components with at most two edges crossing the boundary, one going towards the root and one down to exactly one subtree.

An $MCD(T)$ is built as follows: The first component is defined by T, just like in $CD(T)$. To find recursively the rest of the components, if a component C has no edge from below, we select the centroid c of C as a splitting node, just like when building $CD(T)$. Otherwise, let (x, y) be the edge that crosses the boundary from below, where x is in C, y is a child of x and y is not in C, and let c be the centroid of C (possibly $x = c$). As a splitting node choose the lowest common ancestor of x and c (possibly x or c). By induction every component has at most one edge from below and one edge from above. A useful property of $MCD(T)$ is captured by the following lemma:

▶ **Lemma 2.** *For any given binary tree T, we have $h(MCD(T)) \leq 2 + 2\log_2 |T|$, where $h(MCD(T))$ denotes the height of $MCD(T)$.*

Proof. In $MCD(T)$ if a component C does not have an edge from below then the centroid of C is used as a splitting node, thus generating three components C_l, C_r, and C_p such that $|C_l| \leq \frac{|C|}{2}$, $|C_r| \leq \frac{|C|}{2}$, and $|C_p| \leq \frac{|C|}{2}$. Otherwise, C has one edge (x, y) from below, with x being the node that is part of C. Let c be a centroid of C. We have to consider the following two cases: if c happens to be the lowest common ancestor of c and x, then our algorithm will split C according to the actual centroid, so we will have that $|C_l| \leq \frac{|C|}{2}$, $|C_r| \leq \frac{|C|}{2}$, and $|C_p| \leq \frac{|C|}{2}$. Otherwise, the splitting node will produce components C_l, C_r, and C_p, where C_l contains x and C_r contains c, i.e., we have $|C_l| + |C_p| \leq \frac{|C|}{2}$ and $|C_r| \geq \frac{|C|}{2}$. From the first inequality, we have that $|C_l| \leq \frac{|C|}{2}$ and $|C_p| \leq \frac{|C|}{2}$. Notice that C_r is going to be a component corresponding to a complete subtree of T, so it will have no edges from below. This means that in the next recursion level when working with C_r the actual centroid of C_r

will be chosen as a splitting node, thus in the following recursion level the three components produced from C_r will be such that their sizes are at most half the size of C. From the analysis given so far, it becomes clear that when we have a component of size $|C|$ with one edge from below, then we will need at most 2 levels in $MCD(T)$ before producing components all of which will have a guaranteed size of at most $\frac{|C|}{2}$. Hence, the statement follows. ◀

Since every level of $MCD(T)$ can be constructed in O($|T|$) time and we have $|T| = 2n - 1$, we obtain the following:

▶ **Theorem 3.** *For any given binary tree T with n leaves, there exists an algorithm that constructs $MCD(T)$ using* O($n \log n$) *time and* O(n) *space in the RAM model.*

3.2 The Main Algorithm

There is a preprocessing step and a counting (of shared triplets between T_1 and T_2) step.

In the preprocessing step, first we apply a depth first traversal on T_1 to make T_1 *left-heavy*, by swapping children so that for every node u in T_1 the left subtree is larger than the right subtree. Second, we change the leaf labels of T_1, which can also be done by a depth first traversal of T_1, so that the leaves are numbered 1 to n from left to right. Both steps take O(n) time in the RAM model. The second step is performed to simplify the process of transferring the leaf colors between T_1 and T_2. The coloring of a subtree in T_1 will correspond to assigning the same color to a contiguous range of leaf labels. Determining the color of a leaf in T_2 will then require one `if-statement` to find in what range (red or blue) its label belongs to. Finally, we build $MCD(T_1)$ according to the description after Lemma 1.

In the counting step, we visit the nodes of T_1, given by the depth first traversal of the ternary tree $MCD(T_1)$, where the children of every node u in $MCD(T_1)$ are visited from left to right. For every such node u we compute $\sum_{v \in T_2} |s(u) \cap s(v)|$. We achieve this by processing T_2 in two phases, the *contraction* phase and the *counting* phase.

Contraction Phase of T_2. Let $L(T_2)$ denote the set of leaves in T_2 and $\Lambda \subseteq L(T_2)$. In the contraction phase, T_2 is compressed into a binary tree of size O($|\Lambda|$) whose leaf set is Λ. The contraction is done in a way so that all the topologies induced by Λ in T_2 are preserved in the compressed binary tree. This is achieved by the following three sequential steps:

- Prune all leaves of T_2 that are not in Λ,
- Repeatedly prune all internal nodes of T_2 with no children, and
- Repeatedly contract unary internal nodes, i.e., nodes having exactly one child.

Let u be a node of $MCD(T_1)$ and C_u the corresponding component of T_1. For every such node u we have a contracted version of T_2, from now on referred to as $T_2(u)$, where $L(T_2(u)) = L(C_u)$. The goal is to augment $T_2(u)$ with counters (see counting phase below), so that we can find $\sum_{v \in T_2} |s(u) \cap s(v)|$ by scanning $T_2(u)$ instead of T_2. One can imagine $MCD(T_1)$ as being a tree where each node u is augmented with $T_2(u)$. To generate all contractions of T_2 for level i of $MCD(T_1)$, which correspond to a set of disjoint connected components in T_1, we can reuse the contractions of T_2 at level $i - 1$ in $MCD(T_1)$. This means that we can generate the contractions of level i in O(n) time, thus we can generate all contractions of T_2 in O($n \log n$) time. Note that by explicitly storing all contractions, we will also need to use O($n \log n$) space. For our problem, because we traverse $MCD(T_1)$ in a depth first manner, we only need to store the contractions corresponding to the stack of nodes of $MCD(T_1)$ that we have to remember during the traversal of $MCD(T_1)$. Since the

components at every second level of $MCD(T_1)$ have at most half the size of the components two levels above, Lemma 4 states that the size of this stack is always O(n).

▶ **Lemma 4.** *Let T_1 and T_2 be two binary trees with n leaves and $u_1, u_2, \ldots, u_k$ a root to leaf path of $MCD(T_1)$. For the sizes of the corresponding contracted versions $T_2(u_1)$, $T_2(u_2)$, ..., $T_2(u_k)$ we have that $\sum_{i=1}^{k} |T_2(u_i)| = $ O(n).*

Proof. For the root u_1 we have $T_2(u_1) = T_2$, thus $|T_2(u_1)| \leq 2n$. From the proof of Lemma 2 we have that for every component of size x, we need at most two levels in $MCD(T_1)$ before producing components all of which have a size of at most $\frac{x}{2}$. This means that $\sum_{i=1}^{k} |T_2(u_i)| \leq 2n + 2n + \frac{2n}{2} + \frac{2n}{2} + \frac{2n}{4} + \frac{2n}{4} + \cdots + \frac{2n}{2^i} + \frac{2n}{2^i} + \cdots = 2 \sum_{j=0}^{\infty} \frac{2n}{2^j} \leq 8n = $ O(n). ◀

Counting Phase of T_2**.** In the counting phase, we find the value of $\sum_{v \in T_2} |s(u) \cap s(v)|$ by scanning $T_2(u)$ instead of T_2. This makes the total time of the algorithm in the RAM model O($n \log n$), with the space being O(n) because of Lemma 4. We consider the following two cases:
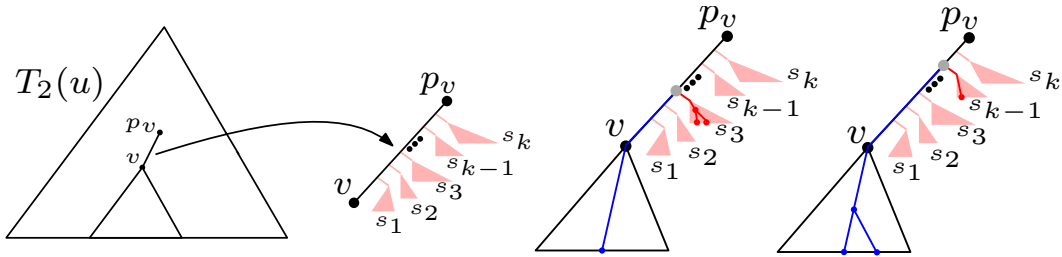
- C_u *has no edges from below.*
    In this case C_u corresponds to a complete subtree of T_1. We act exactly like in the O(n^2) algorithm (Section 2) but now instead of scanning T_2 we scan $T_2(u)$.
- C_u *has one edge from below.*
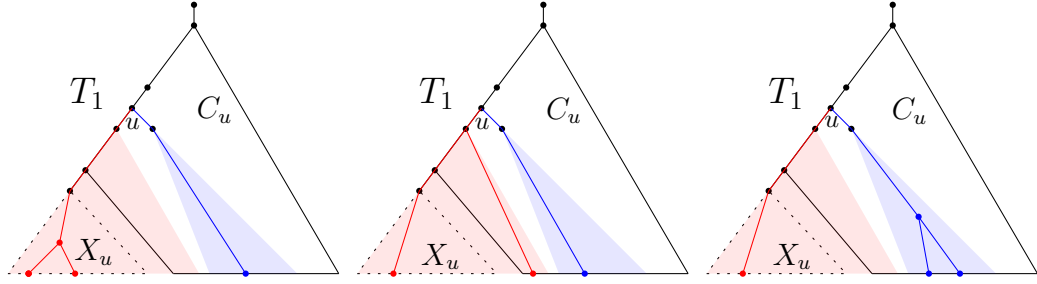
Figure 3 $MCD(T_1)$: Triplets (red and blue) that can be anchored in u with the leaves not being in the component C_u.

Figure 4 Contracted subtrees on edges in $T_2(u)$ and shared triplets rooted at contracted nodes.

In this case C_u does not correspond to a complete subtree of T_1, since the edge from below C_u, points to a subtree X_u, that is located outside of C_u (see Figure 3). Note that because in the preprocessing step T_1 was made to be left-heavy, X_u is always rooted at a node on the leftmost path from u. The leaves in X_u are important because they can be used to form triplets that are anchored in u. Acting in the exact same manner as in the previous case is not sufficient because we need to count these triplets as well.

To address this problem, every edge (p_v, v) in $T_2(u)$ between a node v and its parent p_v, is augmented with some counters about the leaves from X_u that were contracted away in T_2. If v is the root of $T_2(u)$, we add an extra edge to store this information. For every such edge (p_v, v), let $s_1, s_2, \ldots, s_k$ be the contracted subtrees rooted at the edge (see Figure 4). Every such subtree contains either leaves with no color or leaves from X_u that have the color red (the color cannot be blue because T_1 was made to be left-heavy). For every node v in $T_2(u)$ the counters that we have are the following:

- $v_{\tt red}$: total number of red leaves in the subtree of v (including those coming from X_u).
- $v_{\tt blue}$: total number of blue leaves in the subtree of v.
- v_{ts}: total number of red leaves in $s_1, s_2, \ldots, s_k$.
- v_{ps}: total number of pairs of red leaves in $s_1, s_2, \ldots, s_k$ such that each pair comes from the same contracted subtree, i.e., $\sum_{i=1}^{k} \binom{r_i}{2}$ where r_i is the number of red leaves in s_i.

The number of shared triplets that are anchored in a non-contracted node v of $T_2(v)$ can be found like in the $O(n^2)$ algorithm using the counters $v_{\tt red}$ and $v_{\tt blue}$ in (1). As for the number of shared triplets that are anchored in a contracted node on edge (p_v, v), this value is exactly $\binom{v.{\tt blue}}{2} \cdot v_{ts} + v_{\tt blue} \cdot v_{ps}$.

## 3.3   Scaling to External Memory.

We now describe how to make the algorithm scale to external memory. The tree T_1 is stored in an array of size $2n - 1$ by following a preorder layout, i.e., if a node w of T_1 is stored in position p, the left child of w is stored in position $p + 1$ and if x is the size of the left subtree of w, the right child of w is stored in position $p + x + 1$. The components of T_1 are connected parts of T_1, so they can be identified in T_1 without having to make a unique copy for each one of them. For T_2 and its contractions, we use the proof of Lemma 4 to initialize a large enough array that can fit T_2 and every contraction of T_2 that we need to remember while traversing $MCD(T_1)$. This array is used as a stack that we use to push and pop the contractions of T_2. The tree T_2 and its contractions are stored in memory following a post order layout, i.e., if a node w is stored in position p and y is the size of the right subtree of w, the left child of w is stored in position $p - y - 1$ and the right child of w is stored in position $p - 1$.

In the preprocessing step, T_1 can be made left-heavy with two depth first traversals. The first traversal computes for every node u in T_1 the size of the subtree rooted at u. The second traversal starts from the root of T_1, recursively visits the children by first visiting a largest child, and prints all nodes visited along the way to an output array. This output array will at the end of the traversal contain the left-heavy version of T_1 in a preorder layout. From the following Lemma 5 we have that both the first and second depth first traversal of T_1 require $O(n/B)$ I/Os in the cache oblivious model, i.e., making T_1 left-heavy requires $O(n/B)$ I/Os in the cache oblivious model.

In Lemma 5 we consider the I/Os required to apply a depth first traversal on a binary tree T that is stored in memory following a local layout, i.e., the nodes of every subtree of T are stored consecutively in memory and every node has $O(1)$ occurrences in memory. From here on, when we refer to an edge (u, v), we imply that u is the parent of v in T. During a depth first traversal of T, an edge (u, v) is either *processed* to discover v or to backtrack from v to u. In any case, w.l.o.g. we assume that when an edge is processed, both u and v are visited, i.e., both u and v are accessed in memory.

▶ **Lemma 5.** *Let T be a binary tree with n leaves that is stored in an array following a local layout, i.e., the nodes of every subtree of T are stored consecutively in memory and every*

node has $O(1)$ *occurrences in memory. Any depth first traversal that starts from the root of T, and in which for every internal node u in T the children of u are discovered in any order, requires* $O(n/B)$ *I/Os in the cache oblivious model.*
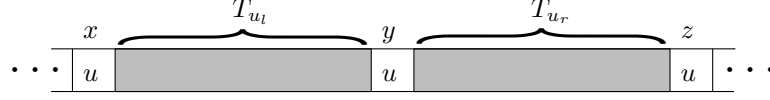
■ **Figure 5** Position of a node u in memory with respect to the two children subtrees of u.

Proof. For a node u in T, let T_u denote the set of nodes in the subtree defined by u. From here on, T_u will be referred to as a *subtree of T*. Let u_l and u_r be the two children of u. In Figure 5 we illustrate the three possibilities for the position of u in memory with respect to T_{u_l} and T_{u_r}. W.l.o.g. and to simplify the presentation of the proof, in our analysis we assume that u is stored in all these three possible positions, denoted x, y, and z. This assumption is w.l.o.g. because in any local layout one or more of these positions is used, thus the number of I/Os is upper bounded by the number of I/Os incurred if we follow our assumption. This placement of u in memory implies that when u is visited in a depth first traversal of T, all the three copies of u are accessed in memory. Note that according to the definition of a local layout, T_{u_l} and T_{u_r} can be interchanged in Figure 5. In the following, the aim is to bound the number of I/Os implied.

Define a node u in T to be *B-light* if $3|T_u| \leq B - 2$, otherwise the node is said to be *B-heavy*. Observe that the children of a B-light node are all B-light. We consider the following disjoint sets of nodes from T:

S_1: Every B-light node
S_2: Every B-heavy node with only B-light children
S_3: Every B-heavy node with two B-heavy children
S_4: Every B-heavy node with one B-heavy child and one B-light child.

For a B-light node u in T, let w be the first B-heavy node we reach in the path from u to the root of T. An I/O incurred by visiting the node u in T is charged to w. This node w can be either in S_2 or S_4. Let w' be the child of w such that $T_{w'}$ contains u. Since $3|T_{w'}| \leq B - 2$, at most 1 I/O is sufficient to visit all nodes in $T_{w'}$. We say that $T_{w'}$ is a subtree that is B-light. In Figure 6a we have an example of a tree, where the gray subtrees denote B-light subtrees.

We now argue that $|S_2| = O(n/B)$ and $|S_3| = O(n/B)$. Let T' be the binary tree created by pruning every B-light node from T and their incident edges, and subsequently contracting nodes with in-degree of 1 and out-degree of 1. An example for T and the corresponding tree T' can be found in Figures 6a and 6b. Let $l_1, l_2, \ldots, l_k$ be the leaves of T' and $T_{l_1}, \ldots, T_{l_k}$ the corresponding subtrees in T. Since all these subtrees are disjoint and for every $1 \leq i \leq k$ we have $|T_{l_i}| > \frac{B-2}{3}$, for the total number of leaves x in T' we have $x \leq 3|T|/(B - 2)$. Hence, we have $|S_2| = x = O(n/B)$. By construction T' is a binary tree, thus we have that $|T'| \leq 2x \leq 6|T|/(B - 2) = O(n/B)$. Since the nodes in S_3 correspond to internal nodes in T', we have $|S_3| = O(n/B)$.

We now argue that the total number of I/Os incurred by the nodes in S_4 is $O(n/B)$, thus proving the statement. Let (u, v) be an edge in T'. This edge corresponds to a unique path, denoted $P_{(u,v)}$ in T that contains every B-heavy node, except u and v, that is in the path from u to v. For example the edge (a, h) in Figure 6b corresponds to $P_{(a,h)} = b \rightarrow c \rightarrow d \rightarrow e \rightarrow f \rightarrow g$. Let $C_{(u,v)}$ contain all B-light and B-heavy nodes, except u and v, rooted at the path from u to v in T. By the local layout followed to store T in memory, the nodes in $C_{(u,v)}$ are stored in two segments of memory (e.g., see

(a)

(b)

(c)
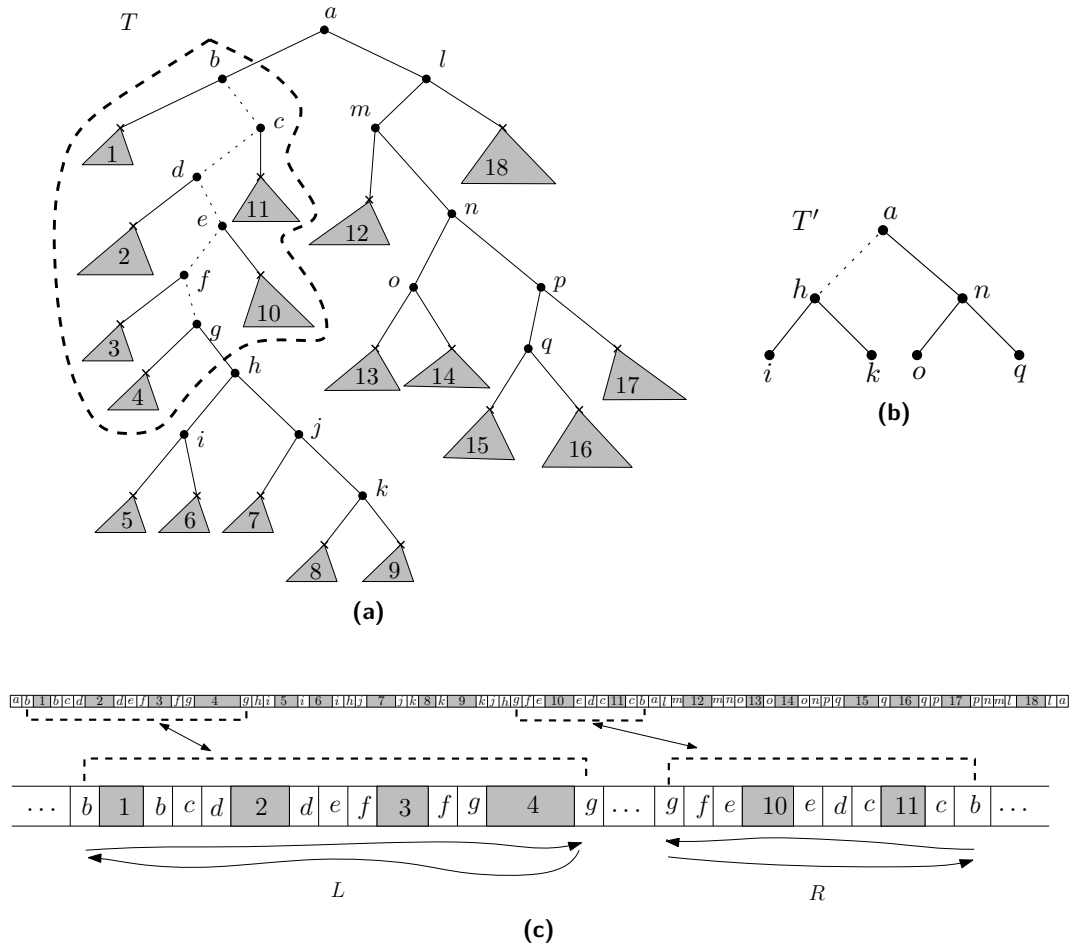
Figure 6 (a) A tree T. The gray subtrees are B-light subtrees and every node not in a B-light subtree is a B-heavy node. (b) The corresponding tree T' according to the proof of Lemma 5. (c) How T is stored in memory, the two segments of memory (in dashed lines) that correspond to the edge (a, h) in T' and how the nodes in $P_{(a,h)}$ are visited (defined by the one directional lines) during a depth first traversal of T.

Figure 6c). Let L be the left segment and R the right segment. During a depth first traversal of T, visiting all nodes in $P_{(u,v)}$ corresponds to visiting L from left to right and then from right to left, and visiting R from right to left and then from left to right. Since each of the B-light subtrees in L and R use at most $B - 2$ positions in memory, by accessing all three copies of a node w in $P_{(u,v)}$ every time w is visited in a depth first traversal of T, we guarantee that the corresponding B-light subtree rooted at w is in cache, i.e., it can be accessed in memory for free. Hence, the total number of I/Os that are sufficient to pay for traversing all nodes in $C_{(u,v)}$ is $4 + \lceil 3|C_{(u,v)}|/B \rceil$, where the $+4$ comes from the 4 I/Os we need to pay (in the worst case) to visit the first and last node of L and R. In total, the total number of I/Os we need to spend for all paths of T that correspond to edges of T' is $\sum_{(u,v) \in T'} (4 + \lceil 3|C_{(u,v)}|/B \rceil) = \mathrm{O}(n/B)$. Together with the fact that for every node of T that corresponds to a node of T' we only spend $\mathrm{O}(1)$ I/Os and there are $\mathrm{O}(n/B)$ such nodes, the statement follows.   ◀

Changing the labels of T_1 can be done in $\mathrm{O}(\frac{n}{B} \log_2 \frac{n}{M})$ I/Os with a cache oblivious sorting routine, e.g., with merge sort. Overall, the preprocessing step requires $\mathrm{O}(\frac{n}{B} \log_2 \frac{n}{M})$ I/Os.

When building $MCD(T_1)$, by scanning the leftmost path that starts from the root of a component C_u, we can find the splitting node of C_u in $1 + \lceil |C_u|/B \rceil$ I/Os. In $T_2(u)$ we spend $1 + \Theta(\lceil |T_u|/B \rceil)$ I/Os for the contraction and counting phase. Since $|T_2(u)| = \Theta(|C_u|)$, overall for a given $(C_u, T_2(u))$ pair the algorithm requires $2 + \Theta(\lceil |C_u|/B \rceil)$ I/Os. However, after $\mathrm{O}(\log_2 \frac{n}{M})$ levels in $MCD(T_1)$, any $(C_u, T_2(u))$ pair will fit in a cache of size M. All such pairs together incur $\mathrm{O}(n/B)$ I/Os. By using a stack to store the contractions of T_2, the remaining pairs incur $\mathrm{O}(\frac{n}{B} \log_2 \frac{n}{M})$ I/Os. Overall, the algorithm requires $\mathrm{O}(\frac{n}{B} \log_2 \frac{n}{M})$ I/Os in the cache oblivious model.

4   The New Algorithm for General Trees

Unlike a binary tree, a general tree can have internal nodes with an arbitrary number of children. By anchoring the triplets of T_1 and T_2 in edges instead of nodes, we show that with only four colors we can count all the shared triplets between the two trees. We start by describing a new $\mathrm{O}(n^2)$ algorithm for general trees. We then show how we can use the same ideas presented in the previous section to extend the $\mathrm{O}(n^2)$ algorithm and reduce the time to $\mathrm{O}(n \log n)$.

4.1   Quadratic Algorithm

For a given tree T, let t be a triplet with leaves i, j, and k that is either a resolved triplet $ij|k$ or an unresolved triplet ijk, where i is to the left of j and for the triplet ijk, k is also to the right of j. Let w be the lowest common ancestor of i and j and (w, c) the edge from w to the child c whose subtree contains j. We anchor t in edge (w, c). Let $s'(w, c)$ be the set containing all triplets anchored in edge (w, c). For the number of shared triplets $S(T_1, T_2)$ we have:

$$S(T_1, T_2) = \sum_{(u,c) \in T_1} \sum_{(v,c') \in T_2} |s'(u, c) \cap s'(v, c')| \ .$$

For the efficient computation of $S(T_1, T_2)$ we use the following coloring procedure: Fix a node u in T_1 and a child c. Color the leaves of every child subtree of u to the left of c red, the leaves of the child subtree defined by c blue, the leaves of every child subtree to the right of c green and give the color black to every other leaf of T_1. We then transfer this coloring to the leaves of T_2. For the resolved triplet $ij|k$, i corresponds to the red color, j corresponds to the

blue color and k corresponds to the black color. For the unresolved triplet ijk, i corresponds to the red color, j corresponds to the blue color and k corresponds to the green color.

Suppose that the node v in T_2 has k children. We are going to compute all shared triplets that are anchored in the k children edges of v in $O(k)$ time. This will give an $O(n^2)$ total running time, because for every edge in T_1 we spend $O(n)$ time in T_2 and there are $O(n)$ edges in T_1. In v we have the following counters:

- $v_{\mathtt{red}}$: total number of red leaves in the subtree of v.
- $v_{\mathtt{blue}}$: total number of blue leaves in the subtree of v.
- $v_{\mathtt{green}}$: total number of green leaves in the subtree of v.
- $\overline{v}_{\mathtt{black}}$: total number of black leaves not in the subtree of v.

While scanning the k children edges of v from left to right, for the child c' that is the m-th child of v, we also maintain the following:

- $a_{\mathtt{red}}$: total number red leaves from the first $m-1$ children subtrees.
- $a_{\mathtt{blue}}$: total number blue leaves from the first $m-1$ children subtrees.
- $a_{\mathtt{green}}$: total number of green leaves from the first $m-1$ children subtrees.
- $p_{\mathtt{red,green}}$: total number of pairs of leaves from the first $m-1$ children subtrees, where one is red, the other is green, and they both come from different subtrees.
- $p_{\mathtt{red,blue}}$: total number of pairs of leaves from the first $m-1$ children subtrees, where one is red, the other is blue, and they both come from different subtrees.
- $p_{\mathtt{blue,green}}$: total number of pairs of leaves from the first $m-1$ children subtrees, where one is blue, the other is green, and they both come from different subtrees.
- $t_{\mathtt{red,blue,green}}$: total number of leaf triples from the first $m-1$ children subtrees, where one is red, one is blue and one is green, and all three leaves come from different subtrees.

Before scanning the children edges of v, every variable is initialized to 0. Then for the child c' every variable is updated in $O(1)$ time as follows:

- $a_{\mathtt{red}} = a_{\mathtt{red}} + c'_{\mathtt{red}}$
- $a_{\mathtt{blue}} = a_{\mathtt{blue}} + c'_{\mathtt{blue}}$
- $a_{\mathtt{green}} = a_{\mathtt{green}} + c'_{\mathtt{green}}$
- $p_{\mathtt{red,green}} = p_{\mathtt{red,green}} + a_{\mathtt{green}} \cdot c'_{\mathtt{red}} + a_{\mathtt{red}} \cdot c'_{\mathtt{green}}$
- $p_{\mathtt{red,blue}} = p_{\mathtt{red,blue}} + a_{\mathtt{blue}} \cdot c'_{\mathtt{red}} + a_{\mathtt{red}} \cdot c'_{\mathtt{blue}}$
- $p_{\mathtt{blue,green}} = p_{\mathtt{blue,green}} + a_{\mathtt{green}} \cdot c'_{\mathtt{blue}} + a_{\mathtt{blue}} \cdot c'_{\mathtt{green}}$
- $t_{\mathtt{red,blue,green}} = t_{\mathtt{red,blue,green}} + p_{\mathtt{red,green}} \cdot c'_{\mathtt{blue}} + p_{\mathtt{red,blue}} \cdot c'_{\mathtt{green}} + p_{\mathtt{blue,green}} \cdot c'_{\mathtt{red}}$

After finishing scanning the k children edges of v, we can compute the shared triplets that are anchored in every child edge of v as follows: for the total number of shared resolved triplets, denoted $\mathtt{tot_{res}}$, we have that $\mathtt{tot_{res}} = p_{\mathtt{red,blue}} \cdot \overline{v}_{\mathtt{black}}$ and for the total number of shared unresolved triplets, denoted $\mathtt{tot_{unres}}$, we have that $\mathtt{tot_{unres}} = t_{\mathtt{red,blue,green}}$. We are now ready to describe the $O(n \log n)$ algorithm.

4.2   Subquadratic Algorithm

Similarly to the case of binary trees in Section 3, there is a preprocessing step and a counting step. The counting step is divided into two phases, the contraction and counting phase of T_2.

In the preprocessing step of the algorithm, we start by transforming T_1 into a binary tree, denoted $b(T_1)$. Let w be a node of T_1 that has exactly k children, where $k > 2$. The k edges
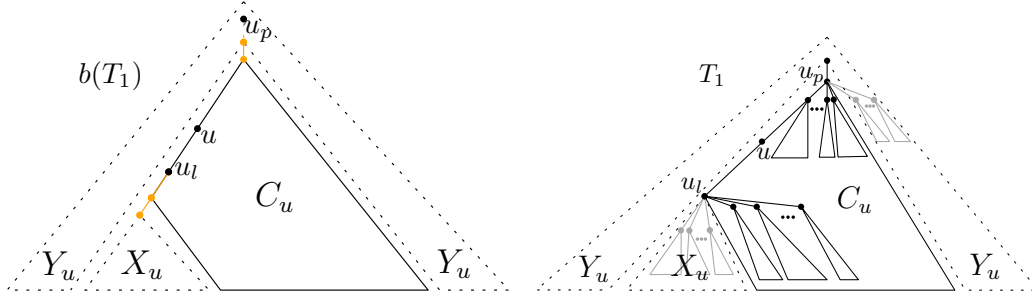
Figure 7 How a component in $b(T_1)$ translates to a component in T_1.

that connect w to its children in T_1 are replaced in $b(T_1)$ by a so called *orange binary tree*. The root of this binary tree is w and the leaves are the k children of w in T_1. Every internal node (except the root) and edge is colored orange, hence the given name. We assume that node w and its k children in T_1, in $b(T_1)$ have the color *black*. This binary tree is built in a way so that every orange node is on the leftmost path that starts from w, and its leftmost leaf stores the heaviest child of w in T_1 (i.e., the child whose subtree is the largest among all other children subtrees of w, when transformed in $b(T_1)$), thus making $b(T_1)$ left-heavy. The order in which the other children of w in T_1 are stored in the remaining leaves does not matter, however for the notation below to be mathematically correct, we assume that after constructing $b(T_1)$, the left to right order of the children of w in T_1 is implicitly updated, so that it matches the left to right order in which they appear in the leaves of the orange binary tree below w in $b(T_1)$.

Let u be a node in $b(T_1)$ and c its right child. By construction, c must be a black node. If u is orange, then let $u_{\tt root}$ be the root of the orange binary tree that u is part of. If u is black, then let $u_{\tt root} = u$. Again by construction, $u_{\tt root}$ must be the parent of c in T_1. For the edge (u, c) in $b(T_1)$, we define $s''(u, c)$ to be the set of triplets that are anchored in edge (u_{root}, c) of T_1, i.e., $s''(u, c) = s'(u_{root}, c)$. Note that for an edge (u', c') in $b(T_1)$ connecting u' with its left child c', we have $s''(u', c') = 0$.

For the number of shared triplets we then have:

$$S(T_1, T_2) = \sum_{(u,c) \in b(T_1)} \sum_{(v,c') \in T_2} |s''(u, c) \cap s'(v, c')| \,.$$

We can capture all triplets in T_1 by coloring $b(T_1)$ instead of T_1. For the nodes u and c where c is the right child of u, the leaves of $b(T_1)$ are colored according to edge (u, c) as follows: the leaves in the left subtree of u are colored red, the leaves in the right right subtree of u are colored blue. If u is an orange node, then the black leaves in the remaining subtrees of the orange binary tree that u is part of are colored green. All other leaves of $b(T_1)$ maintain their color black.

The reason behind transforming T_1 into the binary tree $b(T_1)$, is because now we can use exactly the same core ideas described in Section 3. The tree $b(T_1)$ is a binary tree, so we apply the same preprocessing step, except we do not need to make it left-heavy because by construction it already is. However, we change the labels of the leaves in $b(T_1)$ and T_2, so that the leaves in $b(T_1)$ are numbered 1 to n from left to right. The order in which we visit the nodes of $b(T_1)$ is determined by a depth first traversal of $MCD(b(T_1))$, where the children of every node u in $MCD(b(T_1))$ are visited from left to right.

In Figure 7 we see that a component C_u of $b(T_1)$ structurally looks like a component of T_1 in the binary algorithm of Section 3. However, the edges crossing the boundary can
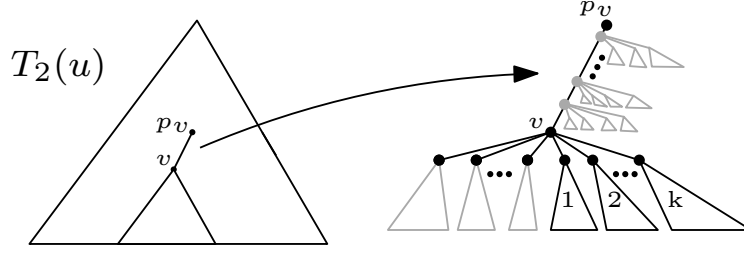
Figure 8 $T_2(u)$: Contracted children subtrees rooted at node v and contracted subtrees rooted at contracted nodes (gray color) on the edge (p_v, v).

now be orange edges as well, which in T_1 translates to more than one consecutive subtrees.

Like in the case of binary input trees, while traversing $MCD(b(T_1))$ we process T_2 in two phases, the contraction phase and the counting phase. The only difference after this point between the algorithm for binary trees and the algorithm for general trees, is in the counters that we have to maintain in the contracted versions of T_2. Otherwise, the same analysis from Section 3 holds.

Contraction Phase of T_2**.** The contraction of T_2 with respect to a set of leaves $\Lambda \subseteq L(T_2)$, happens in the exact same way as described in Section 3, i.e., we start by pruning all leaves of T_2 that are not in Λ, then we prune all internal nodes of T_2 with no children, and finally, we contract the nodes that have exactly one child.

Let u be a node of $MCD(b(T_1))$ and C_u the corresponding component of $b(T_1)$. For every such node u we have a contracted version of T_2, denoted $T_2(u)$, where $L(T_2(u)) = L(C_u)$. Like in the binary algorithm of Section 3, the goal is to augment $T_2(u)$ with counters, so that we can find $\sum_{(v,c') \in T_2} |s''(u,c) \cap s'(v,c')|$ by scanning $T_2(u)$ instead of T_2.

Because of the location where the triplets are anchored, every leaf that was contracted when constructing $T_2(u)$ must have a color and be stored in some way. The color of each such leaf depends on the type of the corresponding component that we have in $b(T_1)$ and the splitting node that is used for that component. For example, in Figure 7 the contracted leaves from X_u will have the red color because $b(T_1)$ is left-heavy. The contracted leaves from the children subtrees of u_p in T_1 can either have the color green or black. If u in $b(T_1)$ happens to be orange and part of the orange binary tree that u_p is the root of, then the color must be green, otherwise black. Finally, every leaf that is not in the subtree defined by u_p, and thus is in Y_u, must have the color black. The way we store this information is described in the counting phase below.

Counting Phase of T_2**.** In Figure 8 we illustrate how a node v in $T_2(u)$ can look like. The contracted subtrees are illustrated with the dark gray color. Every such subtree contains some number of red, green, and black leaves. The counters that we maintain should be so that if v has k children in $T_2(u)$, then we can count all shared triplets that are anchored in every child edge (including those of the contracted children subtrees) of v in $O(k)$ time. At the same time, in $O(1)$ time we should be able to count all shared triplets that are anchored in every child edge of every contracted node that lies on the edge (p_v, v). Then, the time required by the counting phase becomes $O(|T_2(u)|)$, giving the same time bounds as in the binary algorithm of Section 3. In v we have the following counters:

- $v_{\texttt{red}}$: total number of red leaves (including the contracted leaves) in the subtree of v.
- $v_{\texttt{blue}}$: total number of blue leaves in the subtree of v.

- $v_{\texttt{green}}$: total number of green leaves (including the contracted leaves) in the subtree of v.
- $\overline{v}_{\texttt{black}}$: total number of black leaves (including the contracted leaves) not in the subtree of v.

We divide the rest of the counters into two categories. The first category corresponds to the leaves in the contracted children subtrees of v and each counter is stored in a variable of the form $v_{A.x}$. The second category corresponds to the leaves in the contracted subtrees on the edge (p_v, v), and each counter is stored in a variable of the form $v_{B.x}$. For the first category A we have the following counters:

- $v_{A.\texttt{red}}$: total number of red leaves in the contracted children subtrees of v.
- $v_{A.\texttt{green}}$: total number of green leaves in the contracted children subtrees of v.
- $v_{A.\texttt{black}}$: total number of black leaves in the contracted children subtrees of v.
- $v_{A.\texttt{red},\texttt{green}}$: total number of pairs of leaves where one is red, the other is green, and one leaf comes from one contracted child subtree of v and the other leaf comes from a different contracted child subtree of v.

While scanning the k children edges of v from left to right, for the child c' that is the m-th child of v, we also maintain the following:

- $a_{\texttt{red}}$: total number of red leaves from the first $m-1$ children subtrees, including the contracted children subtrees.
- $a_{\texttt{blue}}$: total number of blue leaves from the first $m-1$ children subtrees.
- $a_{\texttt{green}}$: total number of green leaves from the first $m-1$ children subtrees, including the contracted children subtrees.
- $p_{\texttt{red},\texttt{green}}$: total number of pairs of leaves from the first $m-1$ children subtrees, including the contracted children subtrees, where one is red, the other is green, and they both come from different subtrees (one might be contracted and the other non-contracted).
- $p_{\texttt{red},\texttt{blue}}$: total number of pairs of leaves from the first $m-1$ children subtrees, including the contracted children subtrees, where one is red, the other is blue, and they both come from different subtrees (one might be contracted and the other non-contracted).
- $p_{\texttt{blue},\texttt{green}}$: total number of pairs of leaves from the first $m-1$ children subtrees, including the contracted children subtrees, where one is blue, the other is green, and they both come from different subtrees (one might be contracted and the other non-contracted).
- $t_{\texttt{red},\texttt{blue},\texttt{green}}$: total number of leaf triples from the first $m-1$ children subtrees, including the contracted children subtrees, where one is red, one is blue and one is green, and all three leaves come from different subtrees (some might be contracted, some might be non-contracted).

Every variable is updated in $O(1)$ time in exactly the same manner like in the $O(n^2)$ algorithm of Section 4.1. The main difference is in the values of the variables before we begin scanning the children edges of v. Every variable is initialized as follows:

- $a_{\texttt{red}} = v_{A.\texttt{red}}$
- $a_{\texttt{blue}} = 0$
- $a_{\texttt{green}} = v_{A.\texttt{green}}$
- $p_{\texttt{red},\texttt{green}} = v_{A.\texttt{red},\texttt{green}}$
- $p_{\texttt{red},\texttt{blue}} = p_{\texttt{blue},\texttt{green}} = t_{\texttt{red},\texttt{blue},\texttt{green}} = 0$

After finishing scanning the k children edges of v, we can compute the shared triplets that are anchored in every child edge of v (including the children edges pointing to contracted subtrees) as follows: for the total number of shared resolved triplets, denoted $\texttt{tot}_{\texttt{A.res}}$, we

have that $\mathtt{tot_{A.res}} = p_{\mathtt{red,blue}} \cdot \overline{v}_{\mathtt{black}}$ and for the total number of shared unresolved triplets, denoted $\mathtt{tot_{A.unres}}$, we have that $\mathtt{tot_{unres}} = t_{\mathtt{red,blue,green}}$.

The second category B of counters help us count triplets involving leaves (contracted and non-contracted) from the subtree of v and leaves from the contracted subtrees rooted at the edge (p_v, v). We maintain the following:

- $v_{B.\mathtt{red}}$: total number of red leaves in all contracted subtrees rooted at the edge (p_v, v).
- $v_{B.\mathtt{green}}$: total number of green leaves in all contracted subtrees rooted at the edge (p_v, v).
- $v_{B.\mathtt{black}}$: total number of black leaves in all contracted subtrees rooted at the edge (p_v, v).
- $v_{B.\mathtt{red,green}}$: total number of pairs of leaves where one is red and the other is green such that one leaf comes from a contracted child subtree of a contracted node v' and the other leaf comes from a different contracted child subtree of the same contracted node v'.
- $v_{B.\mathtt{red,black}}$: total number of pairs of leaves where one is red and the other is black such that the red leaf comes from a contracted child subtree of a contracted node v' and the black leaf comes from a contracted child subtree of a contracted node v'', where v'' is closer to p_v than v'.

For the total number of shared unresolved triplets, denoted $\mathtt{tot_{B.unres}}$, that are anchored in the children edges of every contracted node that exists in edge (p_v, v), we have that $\mathtt{tot_{B.unres}} = v_{\mathtt{blue}} \cdot v_{B.\mathtt{red,green}}$. For the total number of shared resolved triplets, denoted $\mathtt{tot_{B.res}}$, that are anchored in the children edges of every contracted node that exists in edge (p_v, v), we have that $\mathtt{tot_{B.res}} = v_{\mathtt{blue}} \cdot v_{B.\mathtt{red,black}} + v_{\mathtt{blue}} \cdot v_{B.\mathtt{red}} \cdot (\overline{v}_{\mathtt{black}} - v_{B.\mathtt{black}})$.

4.3   Scaling to External Memory.

The analysis is the same as in Section 3, except for minor details. The proof of Lemma 4 can be trivially modified to apply to general trees as well. Finally, Lemma 5 is generalized to non-binary trees in the following Lemma 6. In Lemma 6, we consider the I/Os required to apply a depth first traversal on a non-binary tree T that is stored in memory following a local layout, i.e., the nodes of every subtree of T are stored consecutively in memory and every node has O(1) occurrences in memory. Similarly to the assumptions we made for Lemma 5, w.l.o.g. we assume that when an edge (u, v) of T is processed in a depth first traversal of T, both u and v are visited, i.e., both u and v are accessed in memory.

▶ **Lemma 6.** *Let T be a non-binary tree with n leaves that is stored in an array following a local layout, i.e., the nodes of every subtree of T are stored consecutively in memory and every node has O(1) occurrences in memory. Any depth first traversal that starts from the root of T and in which for every internal node u in T, after the discovery of the first child of u the remaining children are discovered in order that they appear in memory from left to right, requires O(n/B) I/Os in the cache oblivious model.*
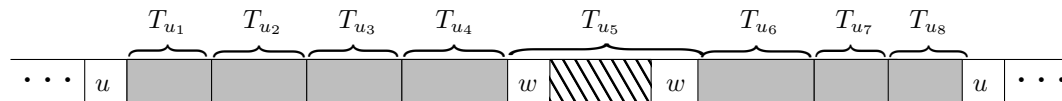
Figure 9 Position of a node u in memory with respect to the 8 subtrees defined by the children of u, with T_{u_5} being a largest subtree.

Proof. This proof can be thought of as an extension of the proof of Lemma 5. Following the proof of Lemma 5, for a node u in T, let T_u denote the set of nodes in the subtree defined by u. For $i \geq 2$, let $u_1, \ldots, u_i$ be the children of u and let $T_{u_1}, \ldots, T_{u_i}$ be the corresponding
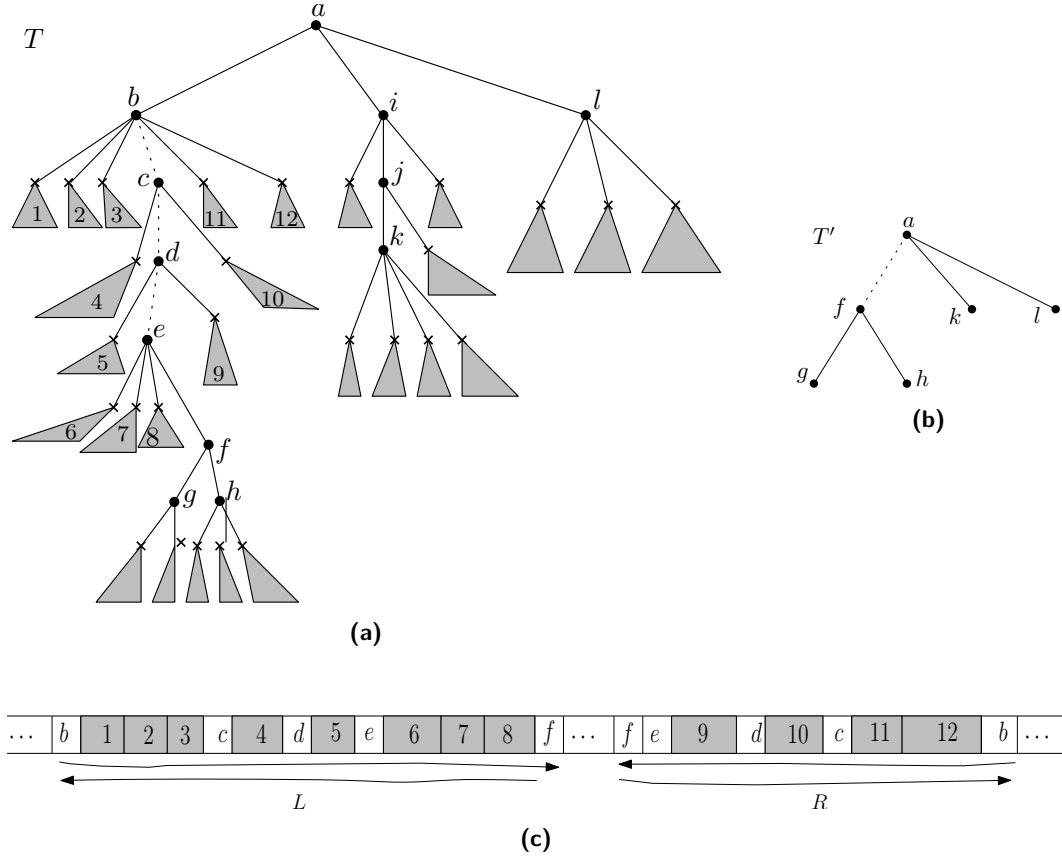
Figure 10 (a) A general tree T. The gray subtrees are B-light subtrees and every node not in a B-light subtree is a B-heavy node. (b) The corresponding tree T' according to the proof of Lemma 6. (c) How T is stored in memory and the two segments of memory that correspond to the edge (a, f) in T'.

subtrees. We assume that these subtrees are ordered from left to right in order that they appear in memory. In the proof of Lemma 5, we implicitly assumed that the positions of the two children of u are stored together with u in memory. For general trees, together with u we need to store a list of arbitrary size $i \geq 2$ containing the positions in memory of every child of u. To avoid complicating the presentation of the proof, we assume that we can find the position in memory of every child of u without this list, i.e., this list is not stored together with u, thus finding the position of any child of u incurs no I/Os. An easy way to support this is to store in every node u in T, one pointer to the first child to be discovered and one pointer to the sibling appearing next in memory. For every node u in T, we allow a constant number of occurrences in memory. For any given placement of the copies of u in memory, we add two copies of u before the first child subtree and after the last child subtree. W.l.o.g. we assume that u is only stored before the first child subtree and after the last child subtree (see Figure 9 for an example).

Define a node u in T to be *B-light* if $2|T_u| \leq B - 2$, otherwise the node is said to be *B-heavy*. Observe that the children of a B-light node are all B-light. We consider the following disjoint sets of nodes from T:

S_1: Every B-light node

S_2: Every B-heavy node with only B-light children

S_3: Every B-heavy node with at least two B-heavy children and an arbitrary number of B-light children

S_4: Every B-heavy node with exactly one B-heavy child and at least 1 B-light children.

For a B-light node u in T, let w be the first B-heavy node we reach in the path from u to the root of T. An I/O incurred by visiting the node u in T is charged to w. This node w can be either in S_2, S_3 or S_4. Let w' be the child of w such that $T_{w'}$ contains u. Since $2|T_{w'}| \leq B - 2$, at most 1 I/O is sufficient to visit all nodes in $T_{w'}$. We say that $T_{w'}$ is a subtree that is B-light. In Figure 10a we have an example of a tree, where the gray subtrees denote B-light subtrees.

Similarly to the proof of Lemma 5, we have that $|S_2| = O(n/B)$ and $|S_3| = O(n/B)$. Since T is non-binary, we have to argue that the number of I/Os spent traversing the B-light subtrees that are rooted at every node in S_2 and S_3 is $O(n/B)$. For a node u in T, let G_u be the size of all gray subtrees rooted at u. For every node u in S_2 we spend at most 1 I/O to traverse the first chosen child subtree and $1 + |G_u|/B$ I/Os to traverse the remaining subtrees, thus $2 + |G_u|/B$ I/Os in total. Since $|S_2| = O(n/B)$ and the gray subtrees in T are disjoint, i.e., $\sum_{u \in T} |G_u| = O(n)$, we spend $O(n/B)$ I/Os traversing the B-light subtrees rooted at every node in S_2. For every node u in S_3, let $d'(u)$ denote the number of B-heavy children of u. For this node u, we spend at most 1 I/O to traverse the first chosen child subtree that could be B-light and $1 + d'(u) + |G_u|/B$ I/Os to traverse the remaining gray subtrees rooted at u. Since $|S_3| = O(n/B)$, we have $\sum_{u \in T'} d'(u) = O(n/B)$. Together with the fact that $\sum_{u \in T} |G_u| = O(n)$, we spend $O(n/B)$ I/Os traversing the B-light subtrees rooted at every node in S_3.

We now argue that the total number of I/Os incurred by the nodes in S_4 is $O(n/B)$, thus proving the statement. Let T' be defined as in the proof of Lemma 5, as well as $P_{(u,v)}$ and $C_{(u,v)}$ for an edge (u,v) in T'. By the local layout followed to store T in memory, the nodes in $C_{(u,v)}$ are stored in two segments of memory (e.g., see Figure 10c). Let w be a node in $P_{(u,v)}$ and G_w be the total size of the gray subtrees rooted at w. We say that w is G-light if $2G_w \leq B - 2$, otherwise G-heavy. There can be $O(n/B)$ G-heavy nodes in T, thus by the same argument as in the previous paragraph, scanning the gray subtrees for all G-heavy nodes together incurs $O(n/B)$ I/Os. For the G-light nodes we follow a similar argument as in the proof of lemma 5. Let L be the left chunk and R the right chunk and w.l.o.g assume that every node in $P_{(u,v)}$ is G-light. During a depth first traversal of T, visiting all nodes in $P_{(u,v)}$ corresponds to visiting L from left to right and then from right to left, and visiting R from right to left and then from left to right. Let c be the child of w that is B-heavy. Since for every node w in $P_{(u,v)}$ we have $2G_w \leq B - 2$, by accessing all two copies of w and c when c is visited in a depth first traversal of T, we guarantee that all the gray subtrees rooted at w are in cache i.e., they can be accessed in memory for free. Hence, $O(n/B)$ I/Os are sufficient to pay to traverse the gray subtrees of all G-light nodes. Overall, by having $M \geq 5B$, where two blocks are used to hold copies of a node w in T, two blocks are used to hold copies of a child of w and one block is used to scan gray subtrees, the statement follows.   ◄

5   Implementation

The algorithms of Sections 3 and 4 have been implemented in the C++ programming language. A high level overview of each implementation is illustrated in Figure 11. The source code is publicly available and can be found at `https://github.com/kmampent/CacheTD`.
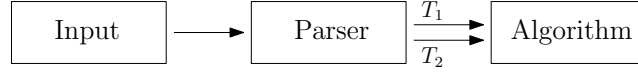
Figure 11 Implementation overview.

5.1 Input

The two input trees T_1 and T_2 are stored in two separate text files following the Newick format. Both trees have n leaves and the label of each leaf is assumed to be a number in $\{1, 2, \ldots, n\}$. Two leaves cannot have the same label.

5.2 Parser

The parser receives the files that store T_1 and T_2 in Newick format, and returns T_1 and T_2 but now with T_1 stored in an array following the preorder layout and T_2 in an array following the postorder layout. The parser takes $O(n)$ time and space in the RAM model and $O(n/B)$ I/Os in the cache oblivious model.

5.3 Algorithm

Having T_1 and T_2 stored in memory following the desired layouts, we proceed with the main part of the algorithm. Both implementations (binary, general) follow the same approach. There exists an *initialization* step and a *distance computation* step.

Initialization. In the initialization step, the preprocessing parts of the algorithms are performed (see Sections 3.2 and 4.2), where the first component of T_1 is built, and the corresponding contracted version of T_2, from now on referred to as *corresponding component* of T_2, is built as well. After this step, the first component of T_1 is stored in an array (different than the one produced by the parser) following the preorder layout. Similarly, the corresponding component of T_2 is stored in an array following the postorder layout.

Distance Computation. Let $\texttt{comp}(T_1)$ and $\texttt{comp}(T_2)$ be the component of T_1 and the corresponding component of T_2 produced by the initialization step. Having these two components available, we can begin counting shared triplets in order to compute $S(T_1, T_2)$. The following steps are recursively applied:

- Starting from the root of $\texttt{comp}(T_1)$ and according to Section 3.1, scan the leftmost path of $\texttt{comp}(T_1)$ to find the splitting node u.
- Scan $\texttt{comp}(T_2)$ to compute for the binary algorithm $\sum_{v \in T_2} |s(u) \cap s(v)|$ (see counting phase of T_2 in Section 3.2), or for the general algorithm $\sum_{(v,c') \in T_2} |s''(u, c) \cap s'(v, c')|$ (see counting phase of T_2 in Section 4.2).
- Using the splitting node u, generate the next three components of T_1. Let $\texttt{comp}(T_1(u_l))$, $\texttt{comp}(T_1(u_r))$, and $\texttt{comp}(T_1(u_p))$ be the components determined by the left child, right child, and parent of u respectively. Let $\texttt{comp}(T_2(u_l))$, $\texttt{comp}(T_2(u_r))$ and $\texttt{comp}(T_2(u_p))$ be the corresponding contracted versions of T_2 with all the necessary counters properly maintained (see contraction phase of T_2 in Section 3.2 for the binary case and in Section 4.2 for the general case).
- Scan and contract $\texttt{comp}(T_2)$ to generate $\texttt{comp}(T_2(u_l))$ and then recurse on the pair defined by $\texttt{comp}(T_1(u_l))$ and $\texttt{comp}(T_2(u_l))$.

- Scan and contract $\mathsf{comp}(T_2)$ to generate $\mathsf{comp}(T_2(u_r))$ and then recurse on the pair defined by $\mathsf{comp}(T_1(u_r))$ and $\mathsf{comp}(T_2(u_r))$.
- Scan and contract $\mathsf{comp}(T_2)$ to generate $\mathsf{comp}(T_2(u_p))$ and then recurse on the pair defined by $\mathsf{comp}(T_1(u_p))$ and $\mathsf{comp}(T_2(u_p))$.

As a final step, print $\binom{n}{3} - S(T_1, T_2)$, which is equal to the triplet distance $D(T_1, T_2)$.

Correctness. The correctness of our implementations was extensively tested by generating hundreds of thousands of random trees of varying size and varying degree and comparing the output of our implementations against the output of the implementations of the $\mathrm{O}(n \log^3 n)$ algorithm in [11] and the $\mathrm{O}(n \log n)$ algorithm in [16].

Changing the Leaf Labels. To get the right theory bounds, changing the leaf labels of T_1 and T_2 must be done with a cache oblivious sorting routine, e.g., merge sort. In the RAM model this approach takes $\mathrm{O}(n \log n)$ time and in the cache oblivious model $\mathrm{O}(\frac{n}{B} \log_2 \frac{n}{M})$ I/Os. A second approach is to exploit the fact that each label is between 1 and n and use an auxiliary array that stores the new labels of the leaves in T_1, which we then use to update the leaf labels of T_2. In the RAM model this second approach takes $\mathrm{O}(n)$ time but in the cache oblivious model $\mathrm{O}(n)$ I/Os. In practice, the problem with the first approach is that the number of instructions it incurs eliminates any advantage that we expect to get due to its cache related efficiency for L_1, L_2, and L_3 cache. For the input sizes tested, the array of labels easily fits into RAM, so in our implementation of both algorithms we use the second approach.

6 Experiments

In this section we provide an extensive experimental evaluation of the practical performance of the algorithms described in Sections 3 and 4.

6.1 The Setup

The experiments were performed on a machine with 8GB RAM, Intel(R) Core(TM) i5-3470 CPU @ 3.20GHz, 32K L1 cache, 256K L2 cache and 6144K L3 cache. The operating system was Ubuntu 16.04.2 LTS. The compilers used were g++ 5.4 and g++ 4.7, together with cmake 3.5.1. The experiments were performed in text mode, i.e., by booting into the terminal of Ubuntu, to minimize the interference from other programs running at the same time.

Generating Random Trees. We use two different models for generating input trees. The first model is called the *random model*. A tree T with n leaves in this model is generated as follows:

- Create a binary tree T with n leaves as follows: start with a binary tree T with two leaves. Iteratively pick $n - 1$ times a leaf l uniformly at random. Make l an internal node by appending a left child node and a right child node to l, thus increasing the number of leaves in T by exactly 1.
- With probability p contract every internal node u of T, i.e., make the children of u be the children of u's parent and remove u.

The second model is called the *skewed* model. In this model, we can control more directly the shape of the input trees. A tree T with n leaves in this model is generated as follows:

- Create a binary tree T with n leaves as follows: let $0 \leq \alpha \leq 1$ be a parameter, u some internal node in T, l and r the left and right children of u, and $T(u)$, $T(l)$, and $T(r)$ the subtrees rooted at u, l, and r respectively. Create T so that for every internal node u we have $\frac{|T(l)|}{|T(u)|} \approx \alpha$, i.e., if n' is the number of leaves below $T(u)$, and $|\Lambda_l|$ and $|\Lambda_r|$ are the number of leaves in $T(l)$ and $T(r)$ respectively, first choose $|\Lambda_l| = \max(1, \min(\lfloor \alpha \cdot n' \rfloor, n' - 1))$ and then let $|\Lambda_r| = n' - |\Lambda_l|$.
- With probability p contract every internal node u of T' like in the random model.

In both models and after creating T, we shuffle the leaf labels by using `std::shuffle`[1] together with `std::default_random_engine`[2].

Implementations Tested.   Let p_1 and p_2 denote the contraction probability of T_1 and T_2 respectively. When $p_1 = p_2 = 0$, the trees T_1 and T_2 are binary trees, so in the experiments we use the algorithm from Section 3. In every other case, the algorithm from Section 4 is used. Note that the algorithm from Section 4 can handle binary trees just fine, however there is an extra overhead (factor 1.8 slower, see Figure 12) compared to the algorithm from Section 3 that comes due to the additional counters that we maintain in the contractions of T_2.

We compared our implementation with the implementations provided in [11] and [15, 4], and are available at `http://sunflower.kuicr.kyoto-u.ac.jp/~jj/Software/Software.html` and `http://users-cs.au.dk/cstorm/software/tqdist/` respectively. The implementation of the algorithm in [11] has two versions, one that uses `unordered_map`[3], which we refer to as `CPDT`, and another that uses sparsehash[4], which we refer to as `CPDTg`. For binary input trees the hash maps are not used, thus `CPDT` and `CPDTg` are the same. The tqdist library [16], which we refer to as `tqDist`, has an implementation of the binary $O(n \log^2 n)$ algorithm from [15] and the general $O(n \log n)$ algorithm from [4]. If the two input trees are binary the $O(n \log^2 n)$ algorithm is used. We refer to our new algorithm as `CacheTD`.

Statistics.   We measured the execution time of the algorithms with the `clock_gettime` function in C++. Due to the different parser implementations, we do not consider the time taken to parse the input trees. We used the PAPI library[5] for statistics related to instructions, L1, L2, and L3 cache accesses and misses. Finally, we count the space of the algorithms by considering the *Maximum resident set size* returned by `/usr/bin/time -v`.

6.2   Results

The experiments are divided into two parts. In the first part, we consider the performance of the algorithms when their memory requirements do not exceed the available main memory (8G RAM). In the second part, we consider the performance when the memory requirements exceed the available main memory (by limiting the available RAM to the operating system to be 1GB), thus forcing the operating system to start using the swap space, which in turn yields the very expensive disk I/Os. All figures can be found in Appendix A.

---

[1]   `http://www.cplusplus.com/reference/algorithm/shuffle/`
[2]   `http://www.cplusplus.com/reference/random/default_random_engine/`
[3]   `http://en.cppreference.com/w/cpp/container/unordered_map`
[4]   `https://github.com/sparsehash/sparsehash`
[5]   `http://icl.utk.edu/papi/`

RAM experiments in the Random Model.   In Figure 13 we illustrate a time comparison of all implementations for trees of up to 2^{21} leaves (~ 2 million) with varying contraction probabilities. Every experiment is run 10 times, and each time on a different tree. All 10 data points are depicted together with a line that goes over their median. The compilers used were g++ 5.4 with cmake 3.5.1 for `tqDist` and g++ 5.4 for CPDT, CPDTg, and `CacheTD`. In all cases, `CacheTD` achieves the best performance. We note that for the case where $p_1 = 0.95$ and $p_2 = 0.2$, CPDT behaves in a different way compared to the experiments in [11]. The same can be observed for the case where $p_1 = 0.8$ and $p_2 = 0.8$. The reason is because of the differences in the implementation of `unordered_map` that exist between the different versions of the g++ compilers. In Figure 14 we compare the performance of CPDT when compiled with g++ 4.7 and g++ 5.4. When p_1 is large, i.e., $p_1 = 0.8$ and $p_1 = 0.95$, we observe that the older version of g++ achieves a better performance. For all other values of p_1, the version of the compiler has no effect on the performance. In Figure 15 we have another time comparison of all implementations but now with CPDT compiled in g++ 4.7. The new algorithm achieves the best performance again, but now the behaviour of CPDT is more stable when p_1 is large. From now on, in every RAM experiment CPDT is compiled in g++ 4.7.

In Figure 16 we show the space consumption of the algorithms. `CacheTD` is the only algorithm that uses $\mathrm{O}(n)$ space for both binary and general trees. In theory we expect that the space consumption is better and this is also what we get in practice.

In Figures 17 and 18 we can see how the contraction parameter affects the running time and the space consumption of the algorithms respectively.

Finally, in Figures 19, 20 and 21 we compare the cache performance of the algorithms, i.e., how many cache misses (L1, L2 and L3 respectively) the algorithms perform for increasing input sizes and varying contraction parameters. As expected, the new algorithm achieves a significant improvement over all previous algorithms.

RAM experiments in the Skewed Model.   The main interesting experimental results are illustrated in Figure 22, where we plot the alpha parameter against the execution time of the algorithms, when $n = 2^{21}$. The alpha parameter has the least effect on `CacheTD`, with the maximum running time in every graph of Figure 22 being only a factor of 1.15 larger than the minimum. As mentioned in Section 2, CPDT and CPDTg use the heavy light decomposition for T_2. For binary trees, when α approaches 0 or 1, the number of heavy paths that have to be updated because of a leaf color change decreases, thus the total number of operations of the algorithm decreases as well. We can verify this in Figure 23, where we have the plots of the alpha parameter against the instructions. The same cannot be said for all general trees, since the contraction parameters have an effect on the shape of the trees as well. In Figures 24, 25, and 26 we have the same graphs but for L1, L2, and L3 cache misses respectively.

I/O experiments.   In Figures 27 and 28 we illustrate the time, space, and I/O performance in the random and skewed model respectively. Every implementation was compiled with g++ 5.4. Every experiment is run 5 times, each on a different tree. Like in the RAM experiments, all 5 data points are displayed together with a line that goes over their median. To measure the execution time, we used the `time` function of Ubuntu and thus also took into account the time taken to parse the input trees. For the input trees of size 2^{23} and 2^{24} we used the 128 bit implementation of the new algorithms in order to avoid overflows.

Unlike `CacheTD`, the performance of CPDT, CPDTg, and `tqDist` deteriorates significantly from the moment they start performing disk I/Os. Only `CacheTD` managed to finish running

■ **Table 1** Random model: Time performance when limiting the available RAM to be 1GB. For the left table we have $p_1 = p_2 = 0$ and for the right table $p_1 = p_2 = 0.5$.

n	CPDT	tqDist	CacheTD
2^{15}	0m:01s	0m:01s	0m:01s
2^{16}	0m:01s	0m:02s	0m:01s
2^{17}	0m:01s	0m:04s	0m:01s
2^{18}	0m:02s	1m:03s	0m:01s
2^{19}	0m:04s	1h:21m	0m:01s
2^{20}	0m:09s	**0%**	0m:01s
2^{21}	13m:12s	-	0m:03s
2^{22}	**0%**	-	0m:09s
2^{23}	-	-	3m:37s
2^{24}	-	-	10m:35s

n	CPDT	CPDTg	tqDist	CacheTD
2^{15}	0m:01s	0m:01s	0m:01s	0m:01s
2^{16}	0m:01s	0m:01s	0m:01s	0m:01s
2^{17}	0m:01s	0m:01s	0m:03s	0m:01s
2^{18}	0m:03s	0m:03s	0m:07s	0m:01s
2^{19}	0m:07s	0m:07s	5m:20s	0m:01s
2^{20}	3m:43s	1h:13m	**0%**	0m:02s
2^{21}	**15%**	**0%**	-	0m:20s
2^{22}	-	-	-	2m:02s
2^{23}	-	-	-	10m:42s
2^{24}	-	-	-	42m:06s

■ **Table 2** Skewed model: Time performance when limiting the available RAM to be 1GB. For both tables we have $\alpha = 0.5$. For the left table we have $p_1 = p_2 = 0$ and for the right table $p_1 = p_2 = 0.5$.

n	CPDT	tqDist	CacheTD
2^{15}	0m:01s	0m:01s	0m:01s
2^{16}	0m:01s	0m:02s	0m:01s
2^{17}	0m:01s	0m:05s	0m:01s
2^{18}	0m:02s	0m:54s	0m:01s
2^{19}	0m:05s	50m:38s	0m:01s
2^{20}	0m:13s	**0%**	0m:01s
2^{21}	20m:02s	-	0m:03s
2^{22}	**0%**	-	0m:09s
2^{23}	-	-	3m:46s
2^{24}	-	-	13m:36s

n	CPDT	CPDTg	tqDist	CacheTD
2^{15}	0m:01s	0m:01s	0m:01s	0m:01s
2^{16}	0m:01s	0m:01s	0m:01s	0m:01s
2^{17}	0m:01s	0m:01s	0m:03s	0m:01s
2^{18}	0m:03s	0m:03s	0m:06s	0m:01s
2^{19}	0m:07s	0m:07s	3m:21s	0m:01s
2^{20}	6m:24s	2h:31m	**7h:51m**	0m:02s
2^{21}	**12%**	**0%**	-	0m:19s
2^{22}	-	-	-	1m:58s
2^{23}	-	-	-	9m:42s
2^{24}	-	-	-	38m:19s

in a reasonable amount of time for all input sizes. For every other algorithm, some data points are missing because the execution time required was too big. To get an idea of how big, in Tables 1 and 2 we again have the time performance of the algorithms in the random and skewed models respectively. This is the exact same time performance as depicted in Figures 27 and 28, however we also include some information about how well the algorithms performed on the extra data point that is missing from the figures. We set a time limit of 10 hours, and only for one pair of input trees T_1 and T_2 we measured for how many nodes of T_1 the value of $\sum_{v \in T_2} |s(u) \cap s(v)|$ was found. Some algorithms managed to process only 0% of the total nodes in T_1, which means that they had to spend most of the time in the preprocessing step (e.g. building the HDT of T_2). The only algorithm that managed to produce a result was tqDist, requiring close to 8 hours for trees with 2^{20} leaves (see Table 2).

7 Conclusion

In this paper we presented two cache oblivious algorithms for computing the triplet distance between two rooted unordered trees, one that works for binary trees and one that works for arbitrary degree trees. Both require $O(n \log n)$ time in the RAM model and $O(\frac{n}{B} \log_2 \frac{n}{M})$

I/Os in the cache oblivious model. We implemented the algorithms in C++ and showed with experiments that their performance surpasses the performance of previous implementations for this problem. In particular, our algorithms are the first to scale to external memory.

Future work and open problems involve the following:

- Could the new algorithms be improved so that in the analysis, the base of the logarithm becomes M/B, thus giving the sorting bound in the cache oblivious model? Would the resulting algorithm be even more efficient in practice?
- Is it possible to compute the triplet distance in $O(n)$ time?
- For the quartet distance computation, could we apply similar techniques to those described in Section 3 and 4 in order to get an algorithm with better time bounds in the RAM model that also scales to external memory?

---- **References** ----

1 A. Aggarwal and J. S. Vitter. The Input/Output Complexity of Sorting and Related Problems. *Communications of the ACM*, 31(9):1116–1127, 1988.

2 M. S. Bansal, J. Dong, and D. Fernández-Baca. Comparing and Aggregating Partially Resolved Trees. *Theoretical Computer Science*, 412(48):6634–6652, 2011.

3 V. Berry and O. Gascuel. Inferring Evolutionary Trees with Strong Combinatorial Evidence. *Theoretical Computer Science*, 240(2):271–298, 2000.

4 G. S. Brodal, R. Fagerberg, C. N. S. Pedersen, T. Mailund, and A. Sand. Efficient Algorithms for Computing the Triplet and Quartet Distance Between Trees of Arbitrary Degree. In *Proceedings of the Twenty-fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1814–1832. Society for Industrial and Applied Mathematics, 2013.

5 D. E. Critchlow, D. K. Pearl, and C. L. Qian. The Triples Distance for Rooted Bifurcating Phylogenetic Trees. *Systematic Biology*, 45(3):323–334, 1996.

6 W. H. E. Day. Optimal Algorithms for Comparing Trees with Labeled Leaves. *Journal of Classification*, 2(1):7–28, 1985.

7 A. J. Dobson. Comparing the Shapes of Trees. In *Combinatorial Mathematics III*, pages 95–100. Springer Berlin Heidelberg, 1975.

8 G. F. Estabrook, F. R. McMorris, and C. A. Meacham. Comparison of Undirected Phylogenetic Trees Based on Subtrees of Four Evolutionary Units. *Systematic Zoology*, 34(2):193–200, 1985.

9 M. Frigo, C. E. Leiserson, H. Prokop, and S. Ramachandran. Cache-Oblivious Algorithms. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, FOCS '99, pages 285–297. IEEE Computer Society, 1999.

10 M. K. Holt, J. Johansen, and G. S. Brodal. On the Scalability of Computing Triplet and Quartet Distances. In *Proceedings of the Meeting on Algorithm Engineering and Experiments*, pages 9–19. Society for Industrial and Applied Mathematics, 2014.

11 J. Jansson and R. Rajaby. A More Practical Algorithm for the Rooted Triplet Distance. In *Algorithms for Computational Biology*, pages 109–125. Springer International Publishing, 2015.

12 J. Jansson and R. Rajaby. A More Practical Algorithm for the Rooted Triplet Distance. *Journal of Computational Biology*, 24(2):106–126, 2017.

13 D. F. Robinson and L. R. Foulds. Comparison of Phylogenetic trees. *Mathematical Biosciences*, 53(1):131–147, 1981.

14 N. Saitou and M. Nei. The Neighbor-Joining Method: A New Method for Reconstructing Phylogenetic Trees. *Molecular Biology and Evolution*, 4(4):406, 1987.

15 A. Sand, G. S. Brodal, R. Fagerberg, C. N. S. Pedersen, and T. Mailund. A practical $O(n \log^2 n)$ time algorithm for computing the triplet distance on binary trees. *BMC Bioinformatics*, 14(2):S18, 2013.

16   A. Sand, M. K. Holt, J. Johansen, G. S. Brodal, T. Mailund, and C. N. S. Pedersen. tqDist: A Library for Computing the Quartet and Triplet Distances Between Binary or General Trees. *Bioinformatics*, 30(14):2079, 2014.

17   A. Sand, M. K. Holt, J. Johansen, R. Fagerberg, G. S. Brodal, C. N. S. Pedersen, and T. Mailund. Algorithms for Computing the Triplet and Quartet Distances for Binary and General Trees. *Biology - Special Issue on Developments in Bioinformatic Algorithms*, 2(4):1189–1209, 2013.
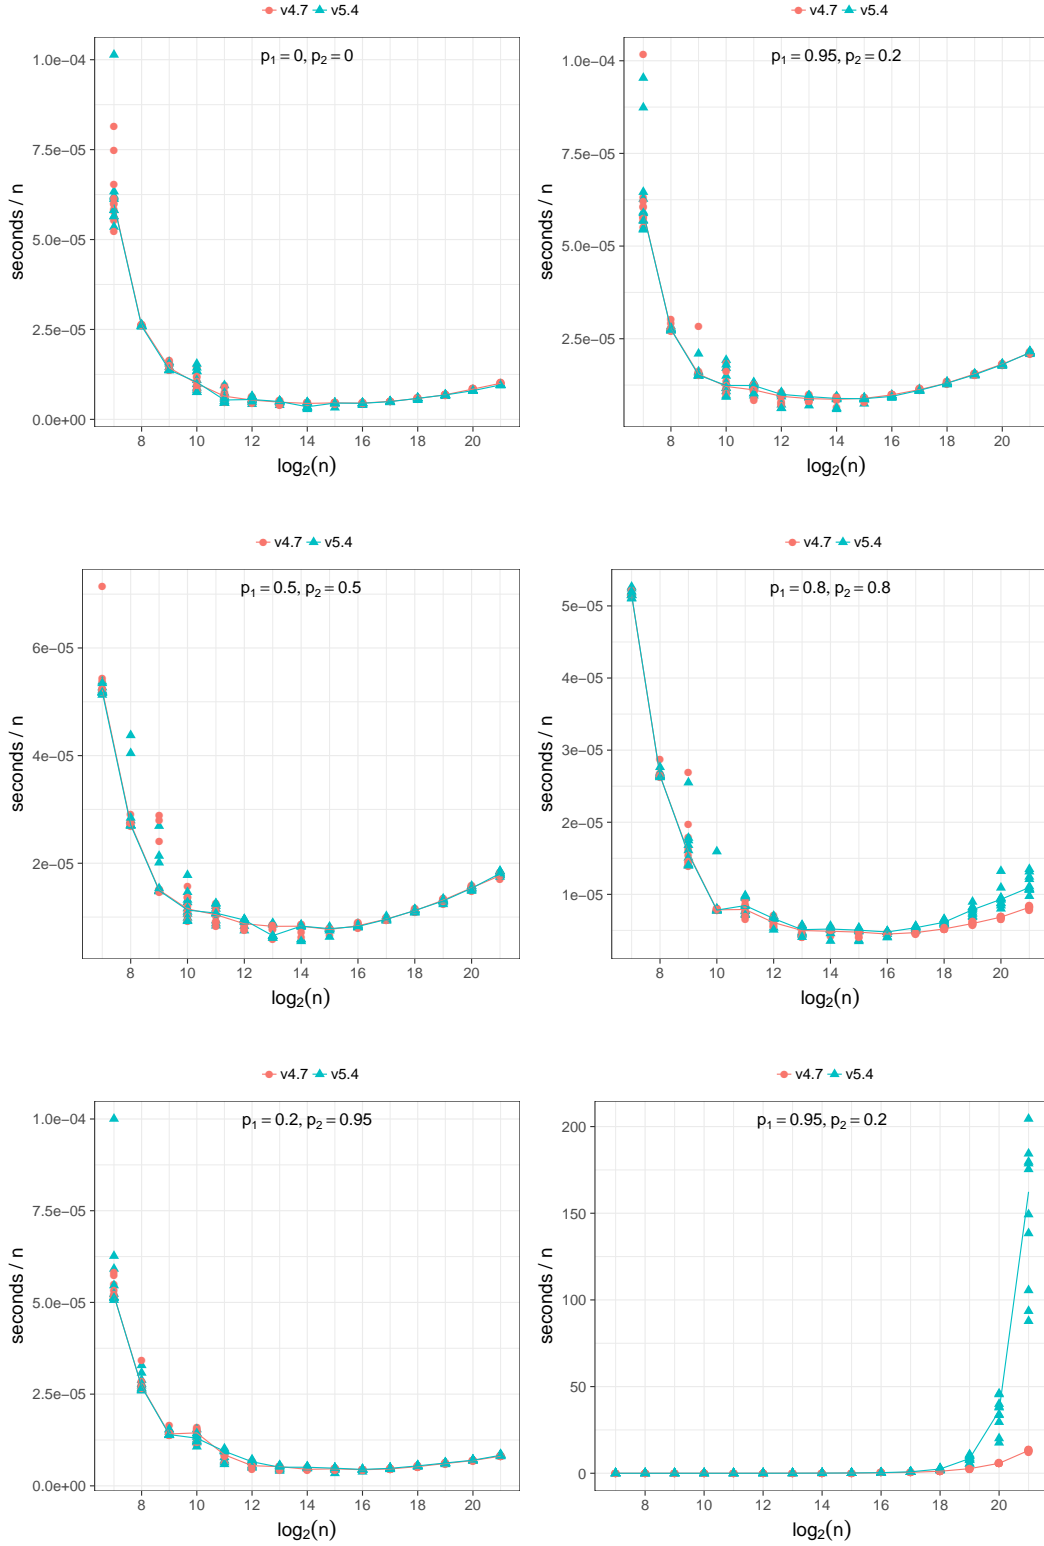
A   Experiment Figures
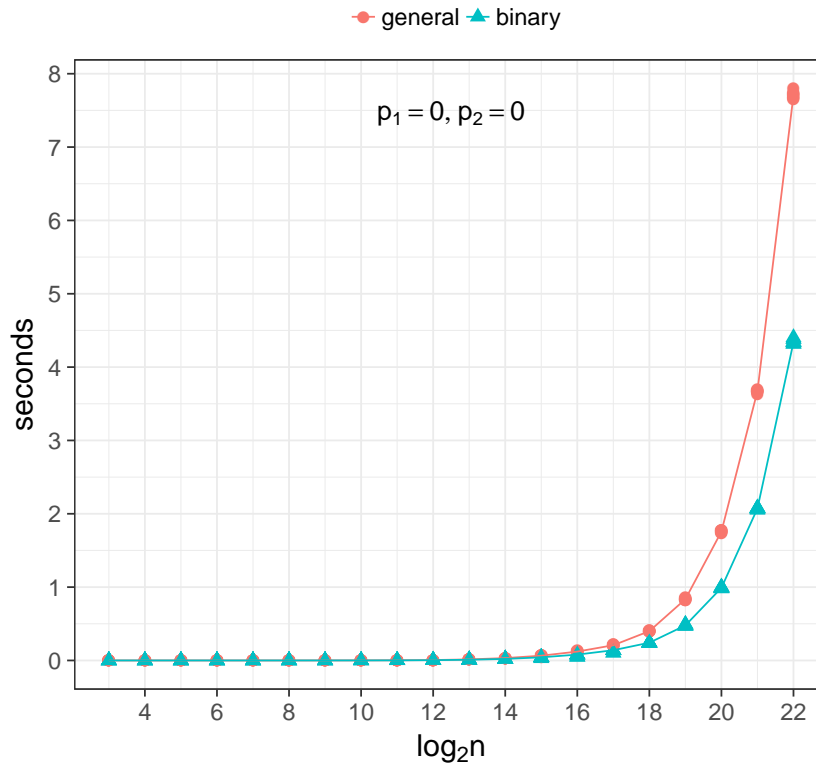
Figure 12 CacheTD: performance of binary (Section 3) and general (Section 4) implementation on binary trees. All data points of the 10 runs are visible in the figure. Each run is on a different tree and the line connects the median of the runs.
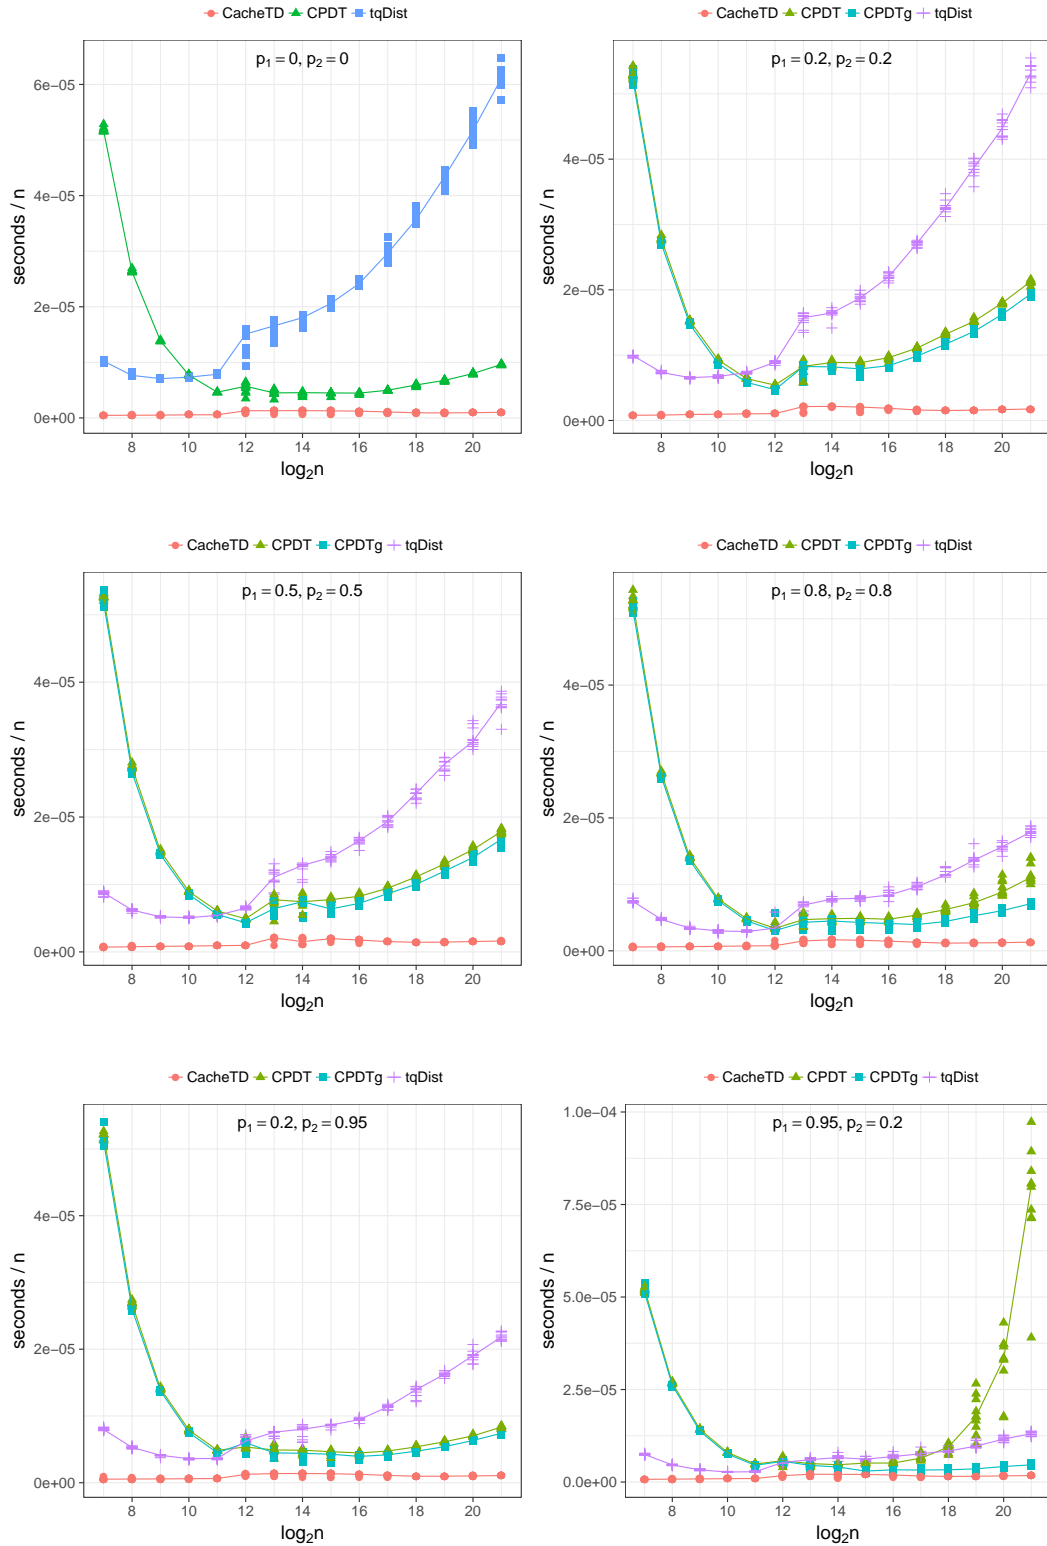
Figure 13 Random model: time performance, where CPDT is compiled in g++ version 5.4.

Figure 14 Random model: time performance of `CPDT` when compiled with g++ 4.7 and g++ 5.4.
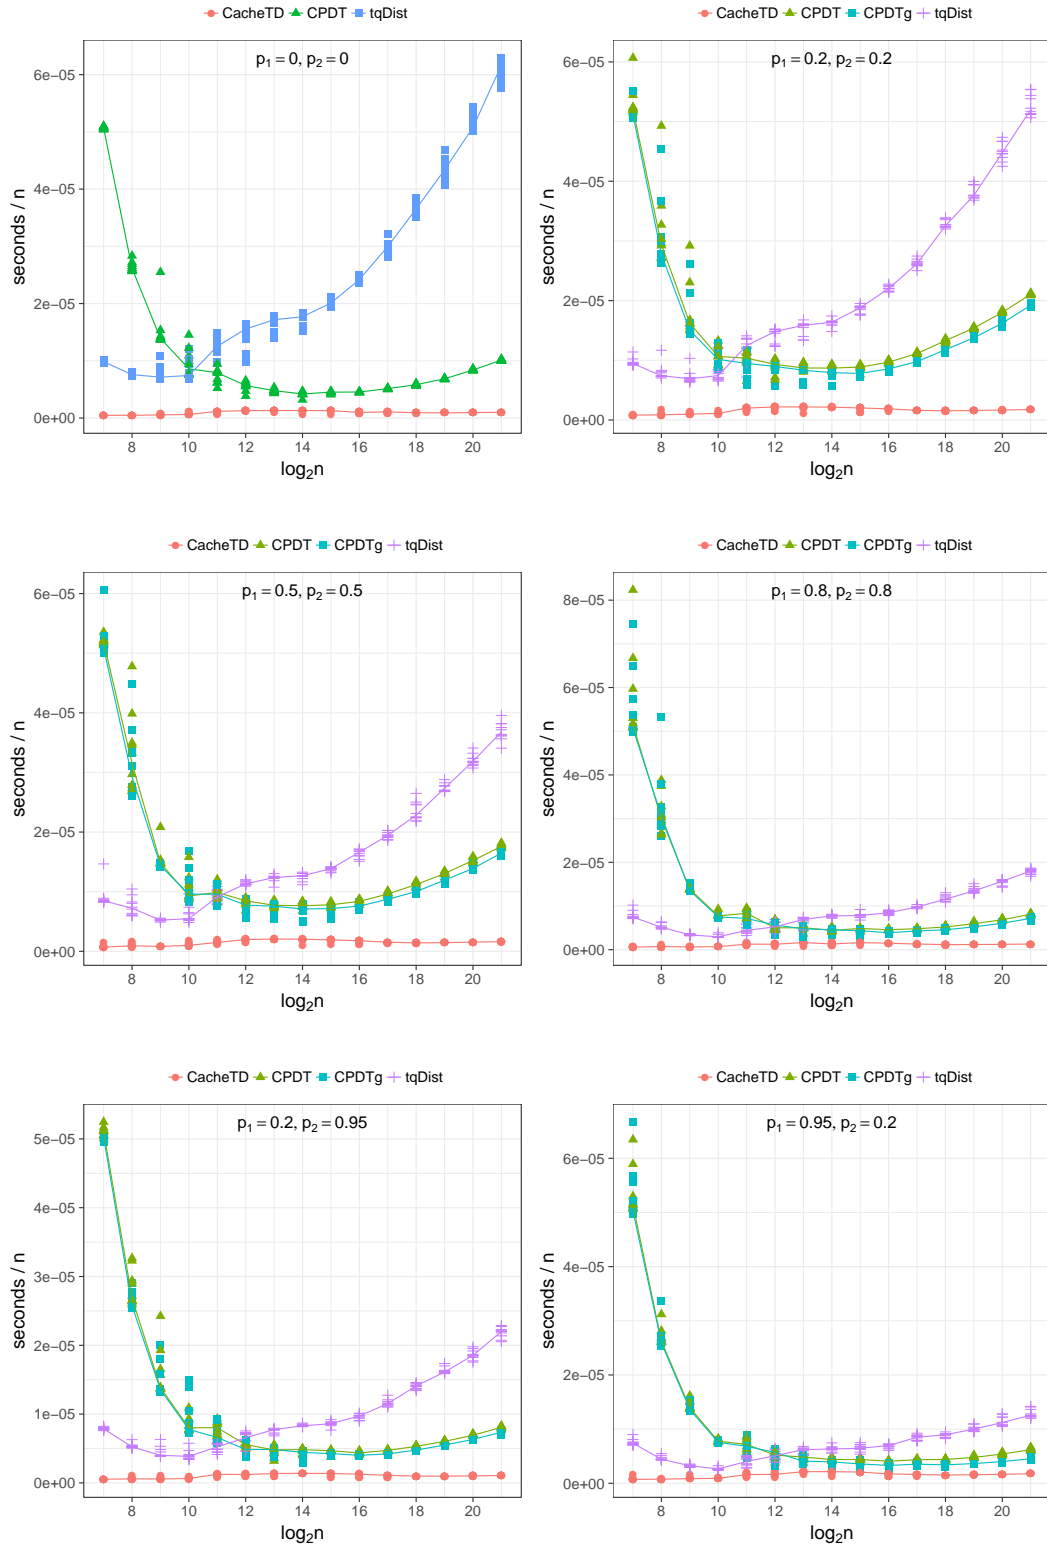
Figure 15 Random model: time performance, where `CPDT` is compiled in g++ version 4.7.
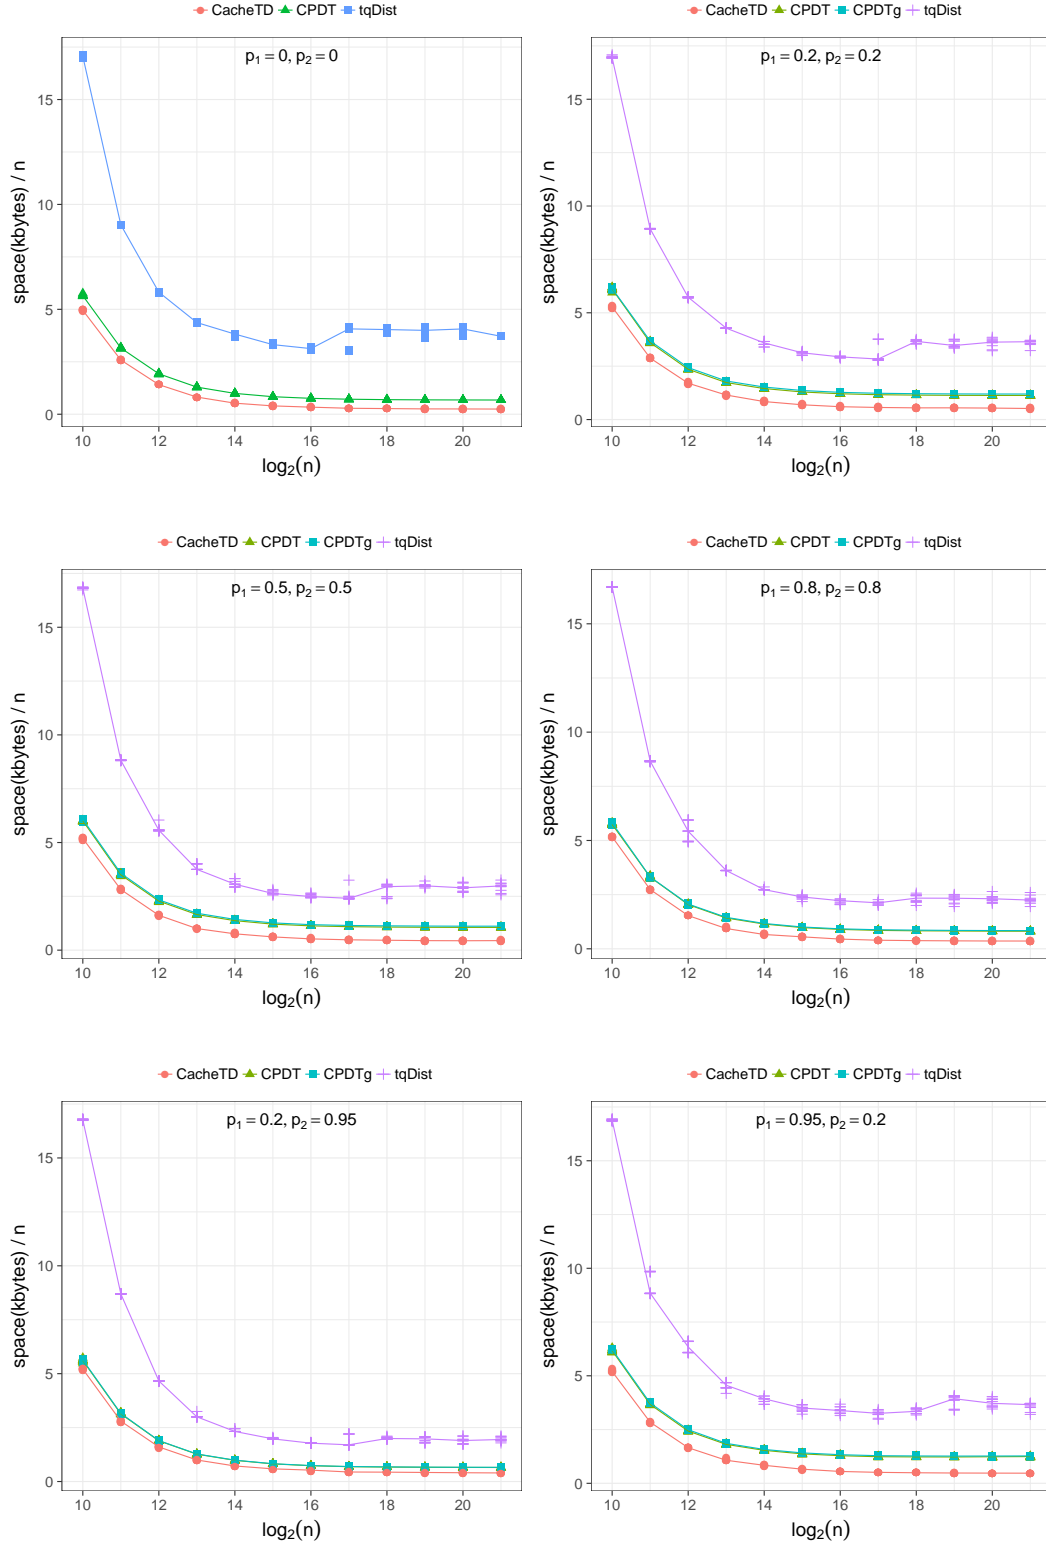
Figure 16 Random model: space performance.
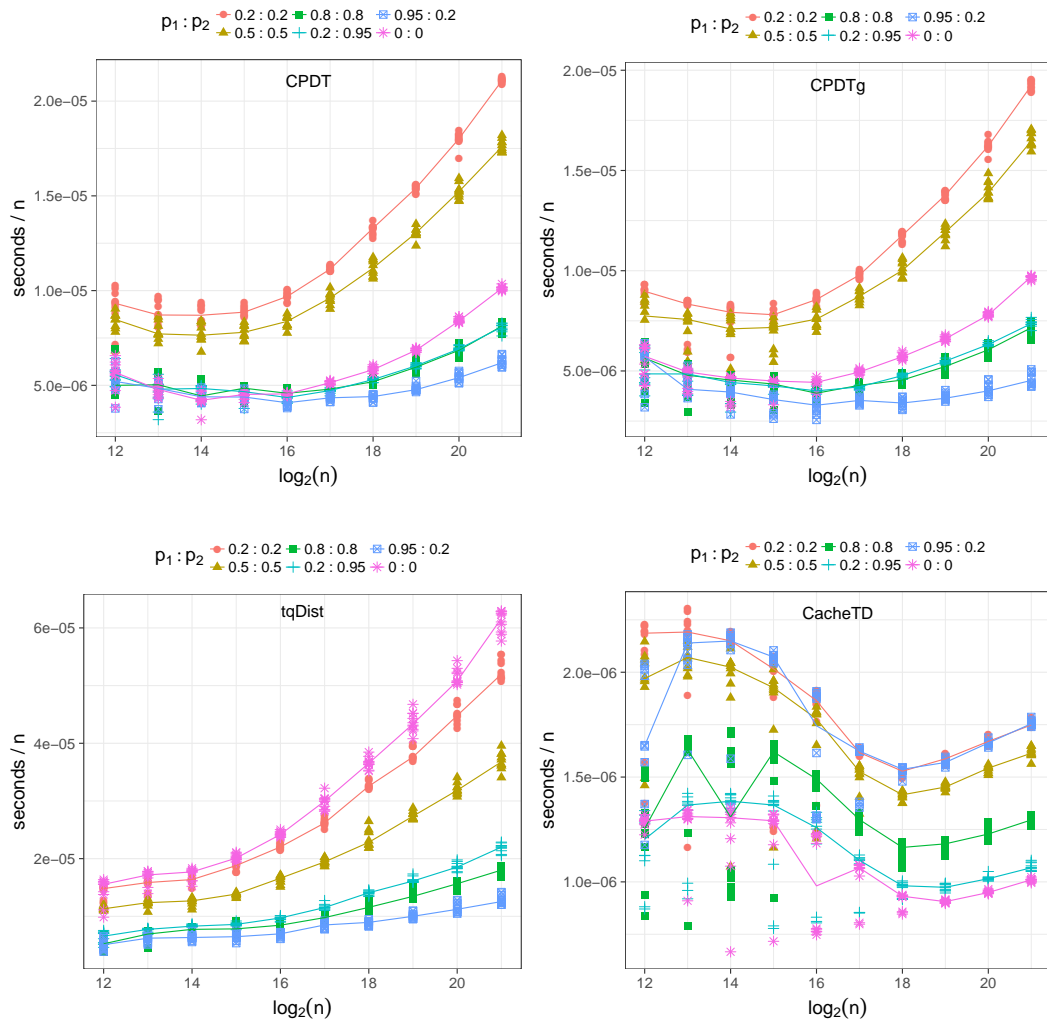
Figure 17 Random model: how the contraction parameter affects execution time.
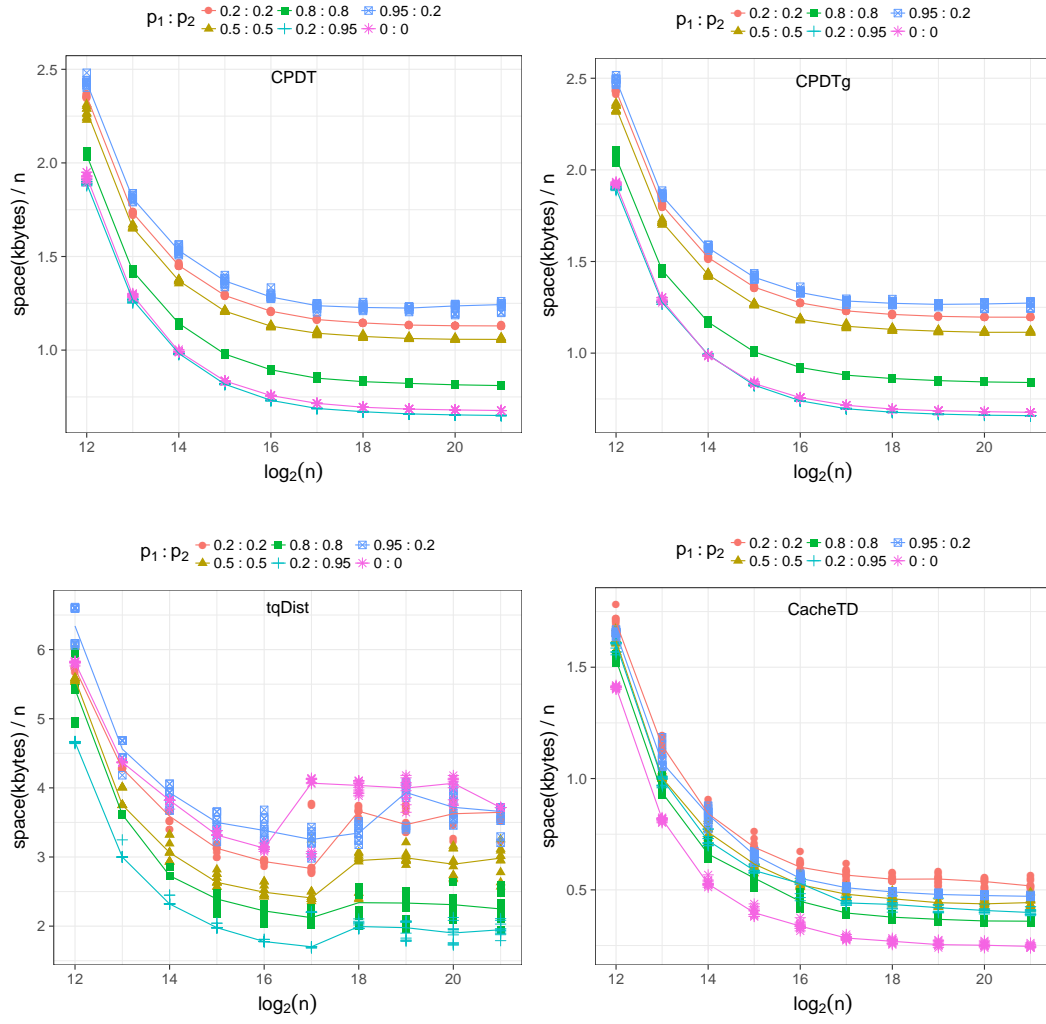
Figure 18 Random model: how the contraction parameter affects space.
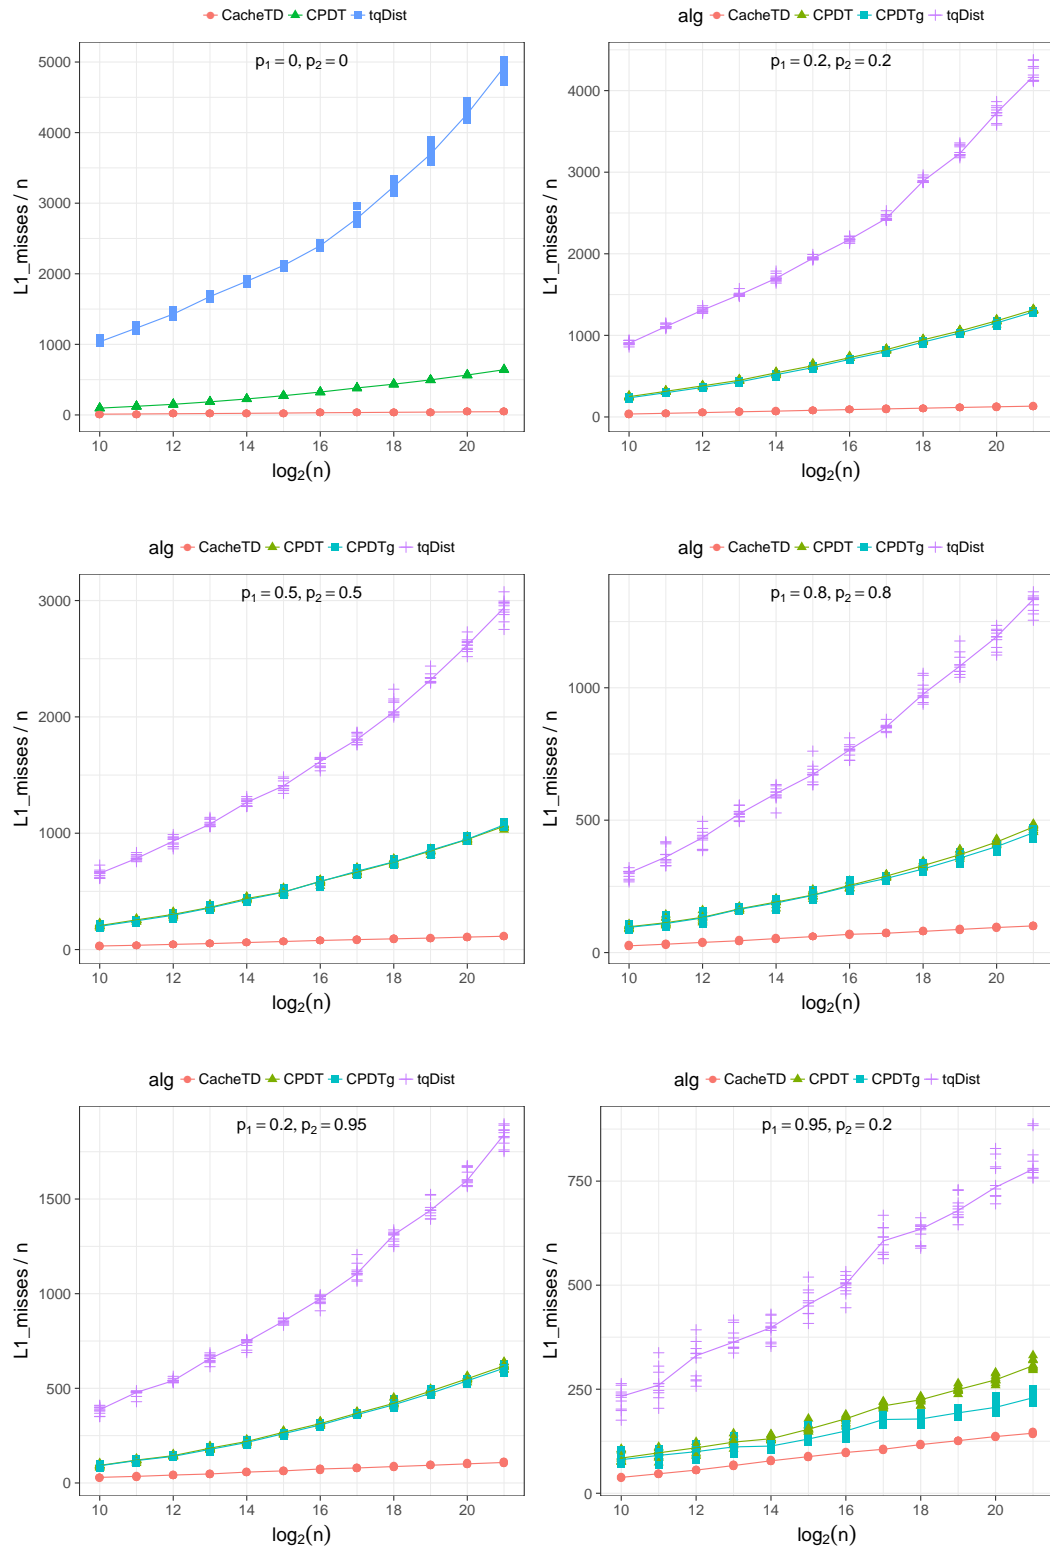
Figure 19 Random model: L1 cache misses.
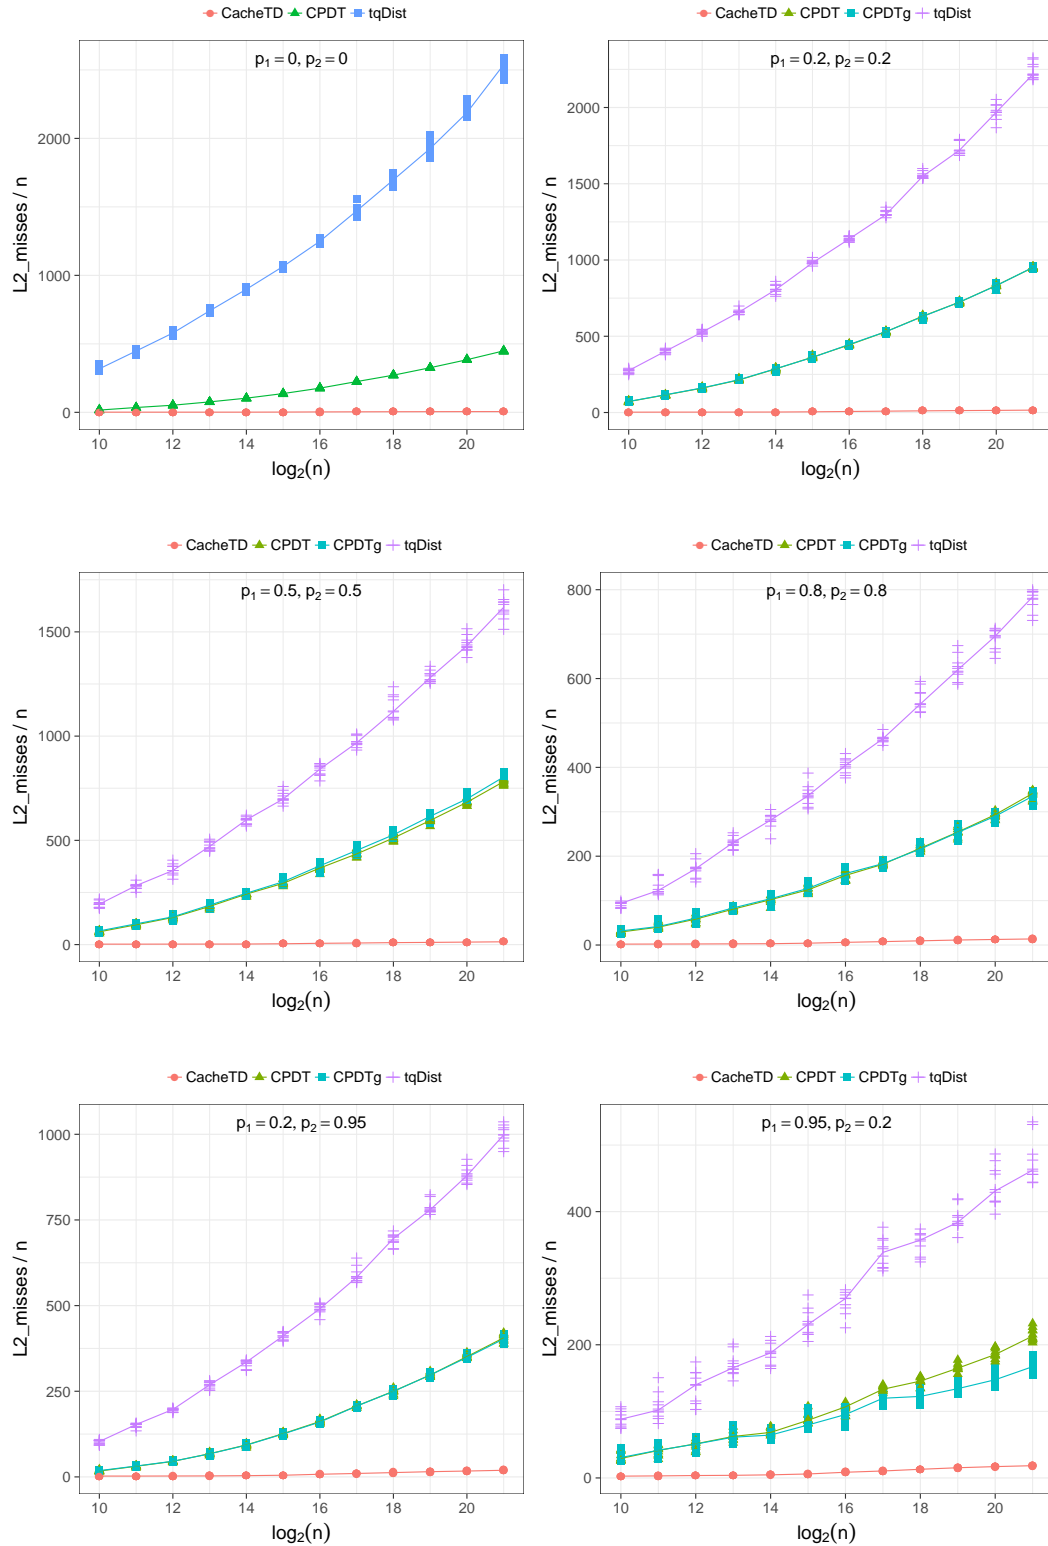
Figure 20 Random model: L2 cache misses.
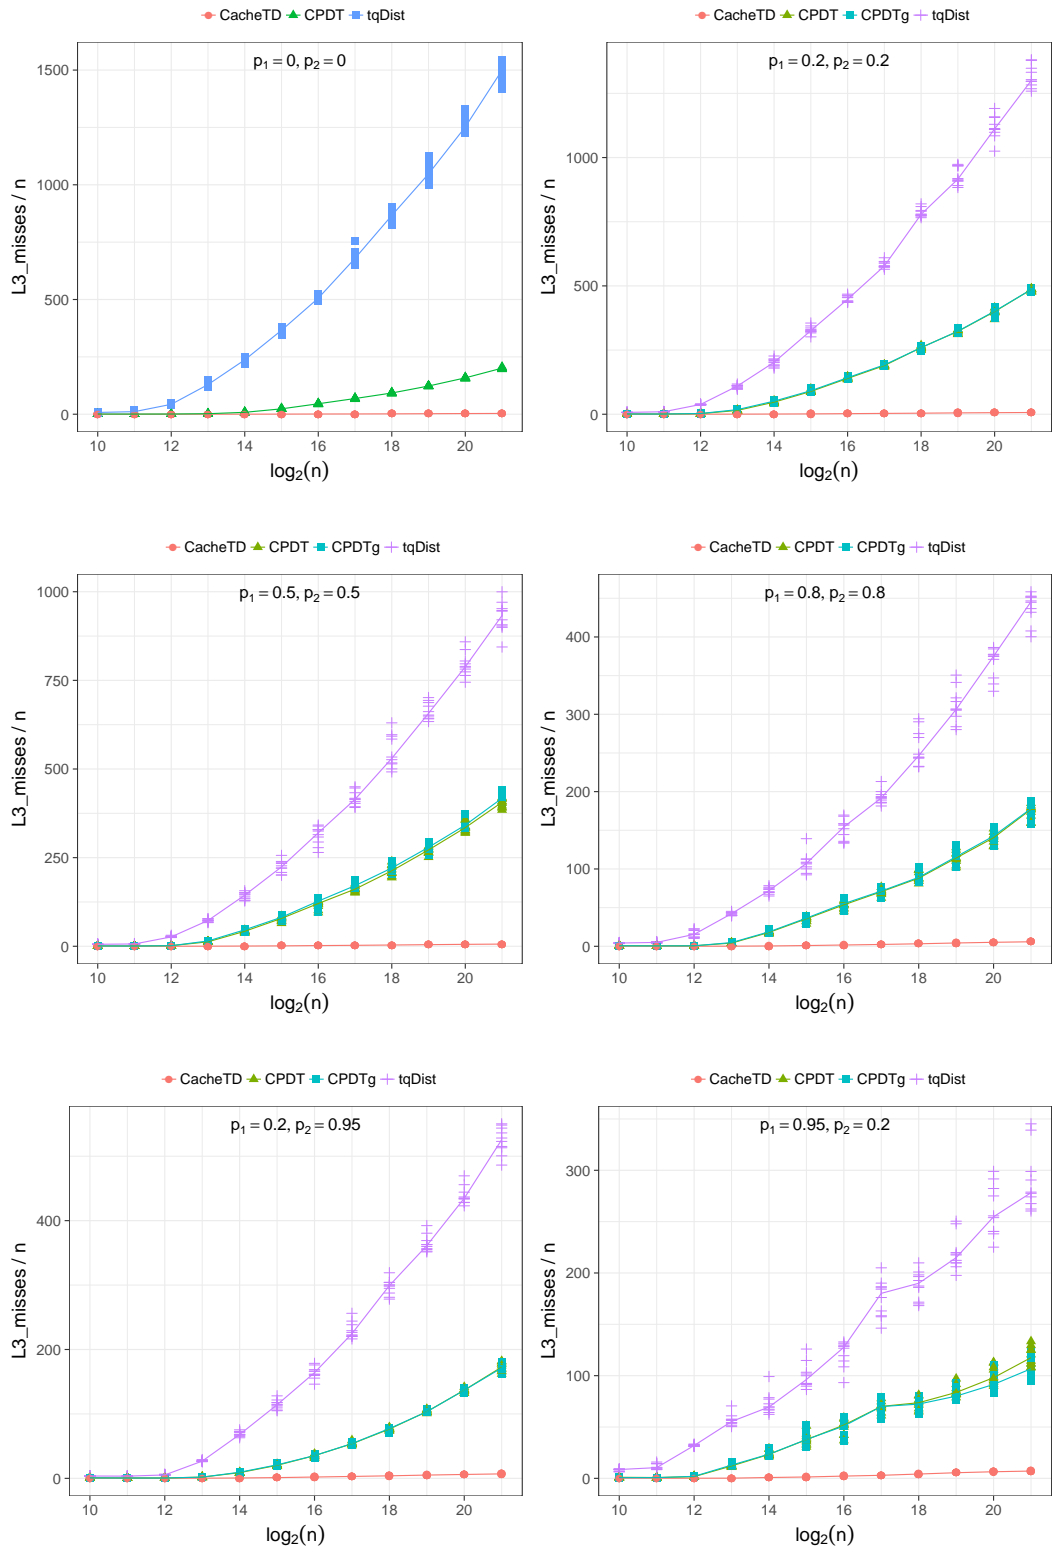
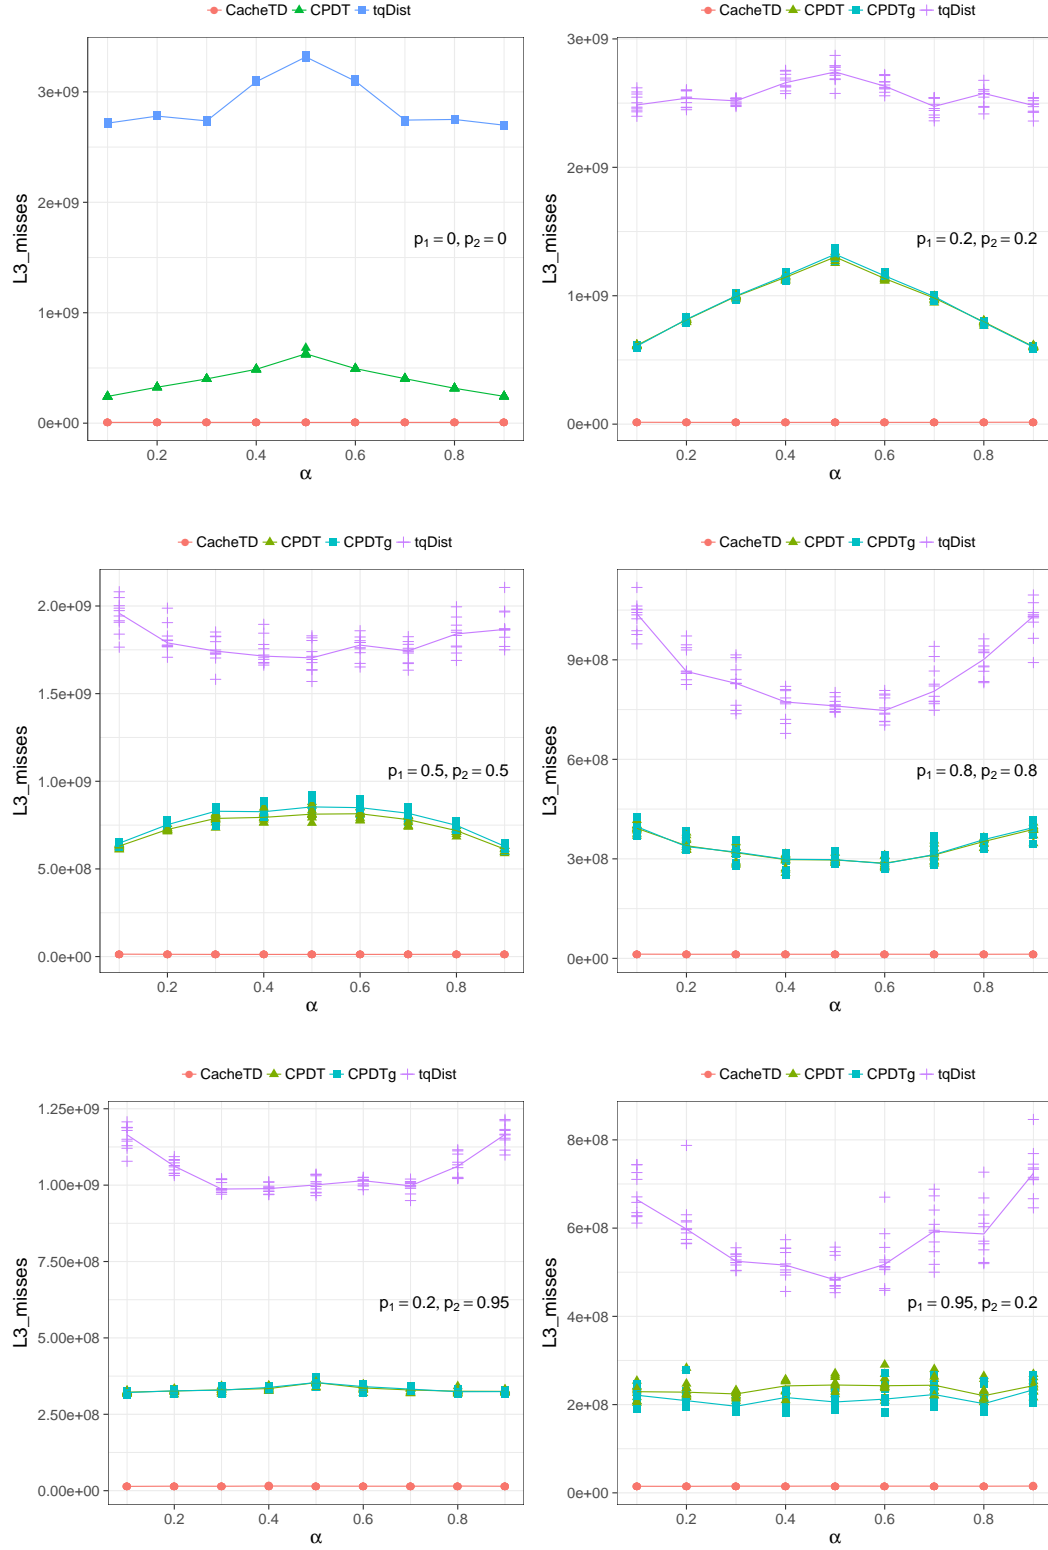**Figure 21** Random model: L3 cache misses.
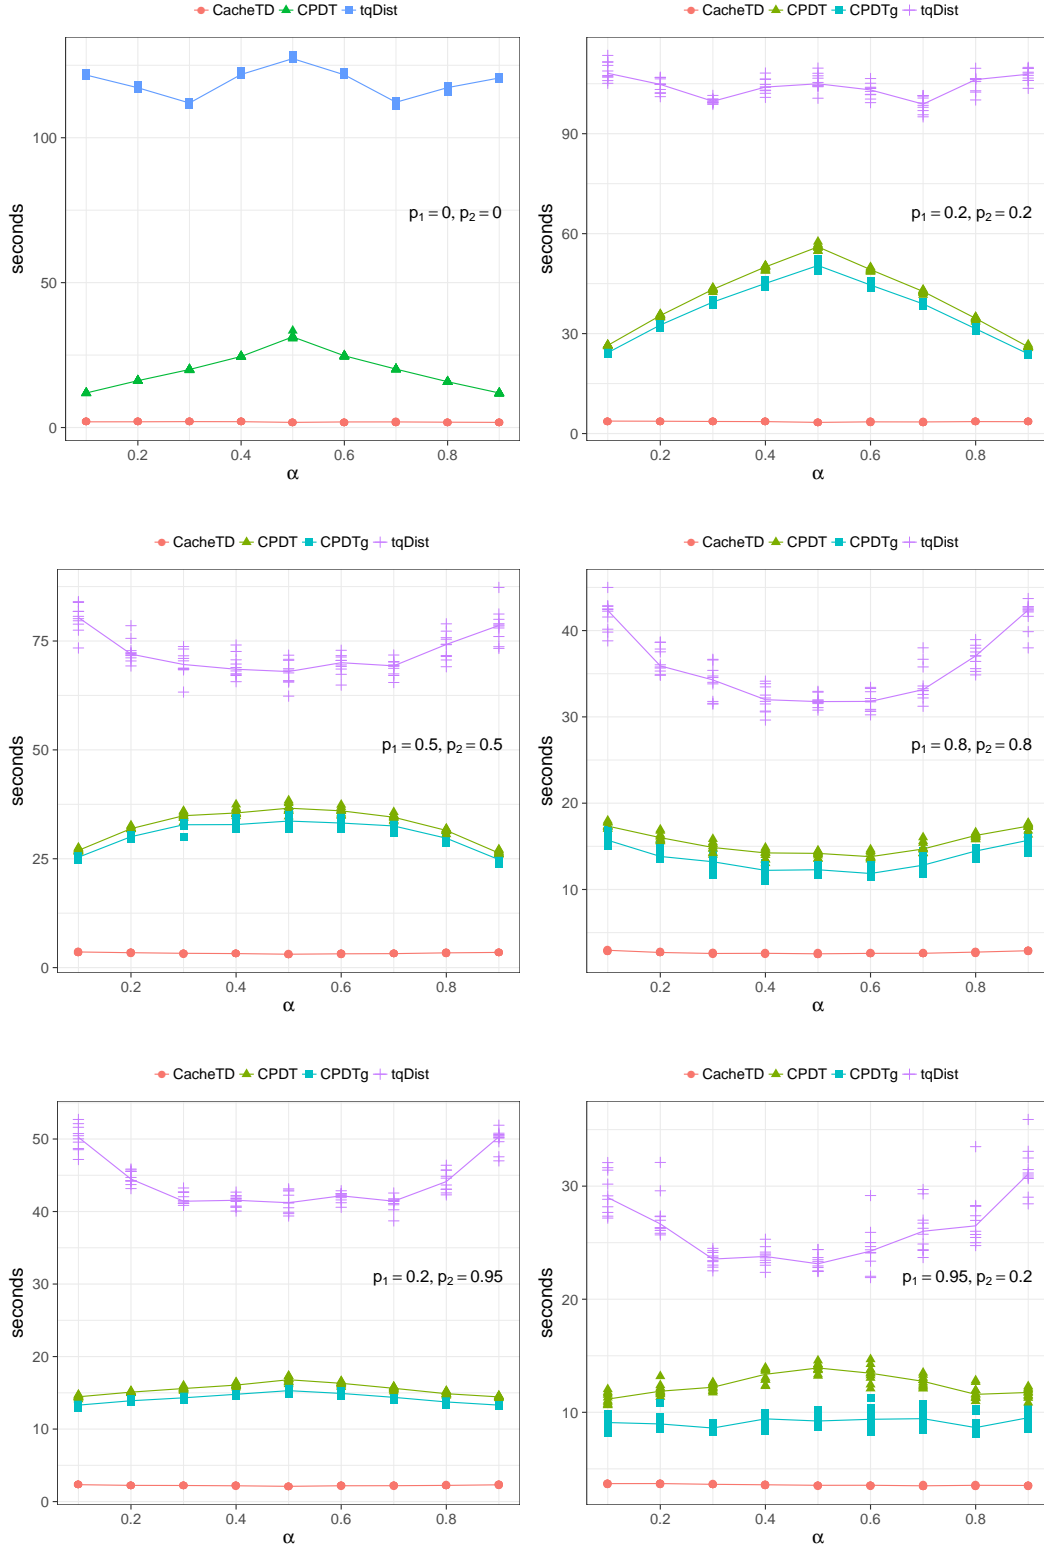
Figure 22 Skewed model: running time ($n = 2^{21}$).
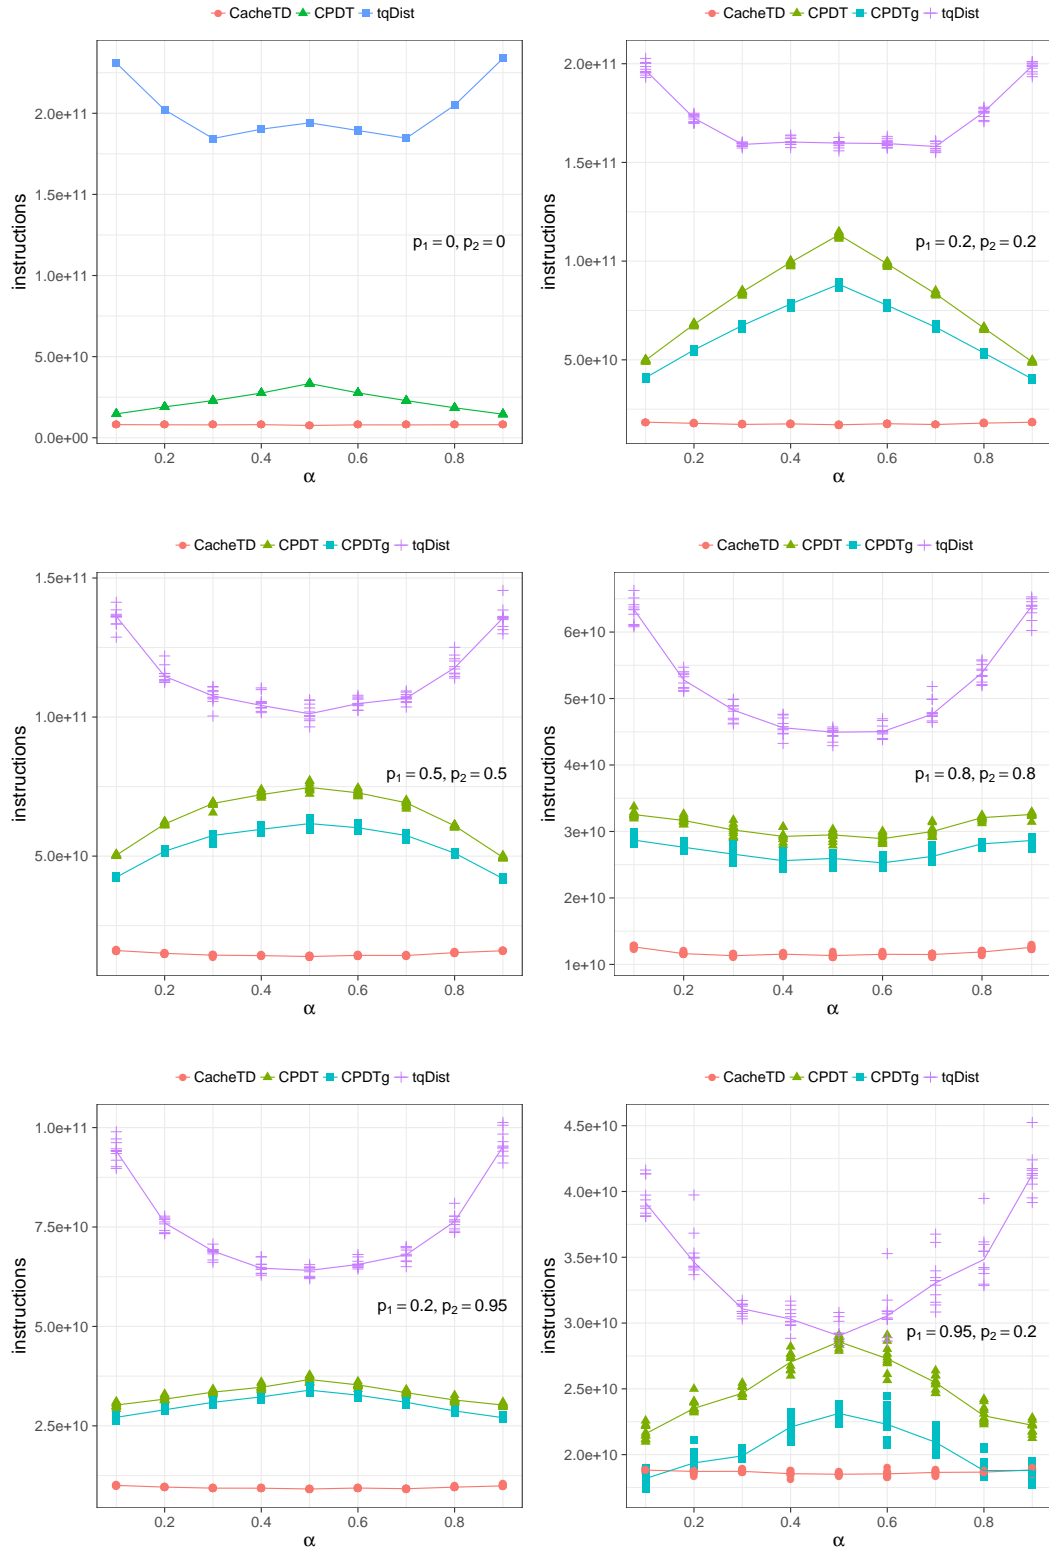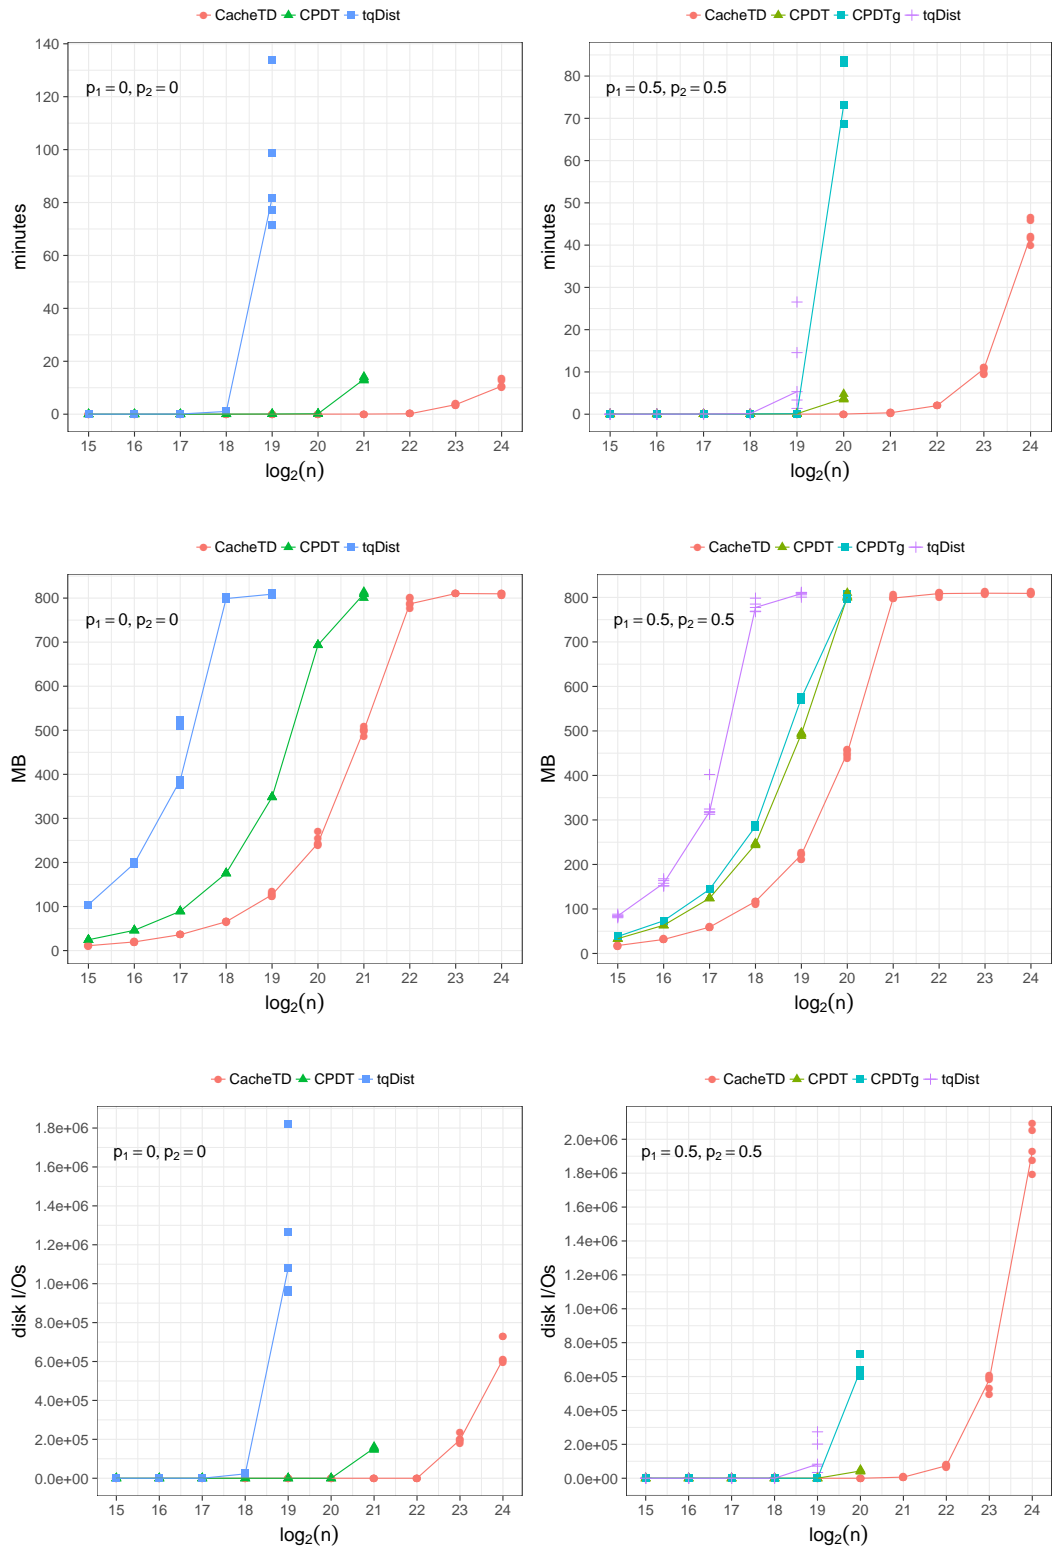
Figure 23 Skewed model: instructions ($n = 2^{21}$).
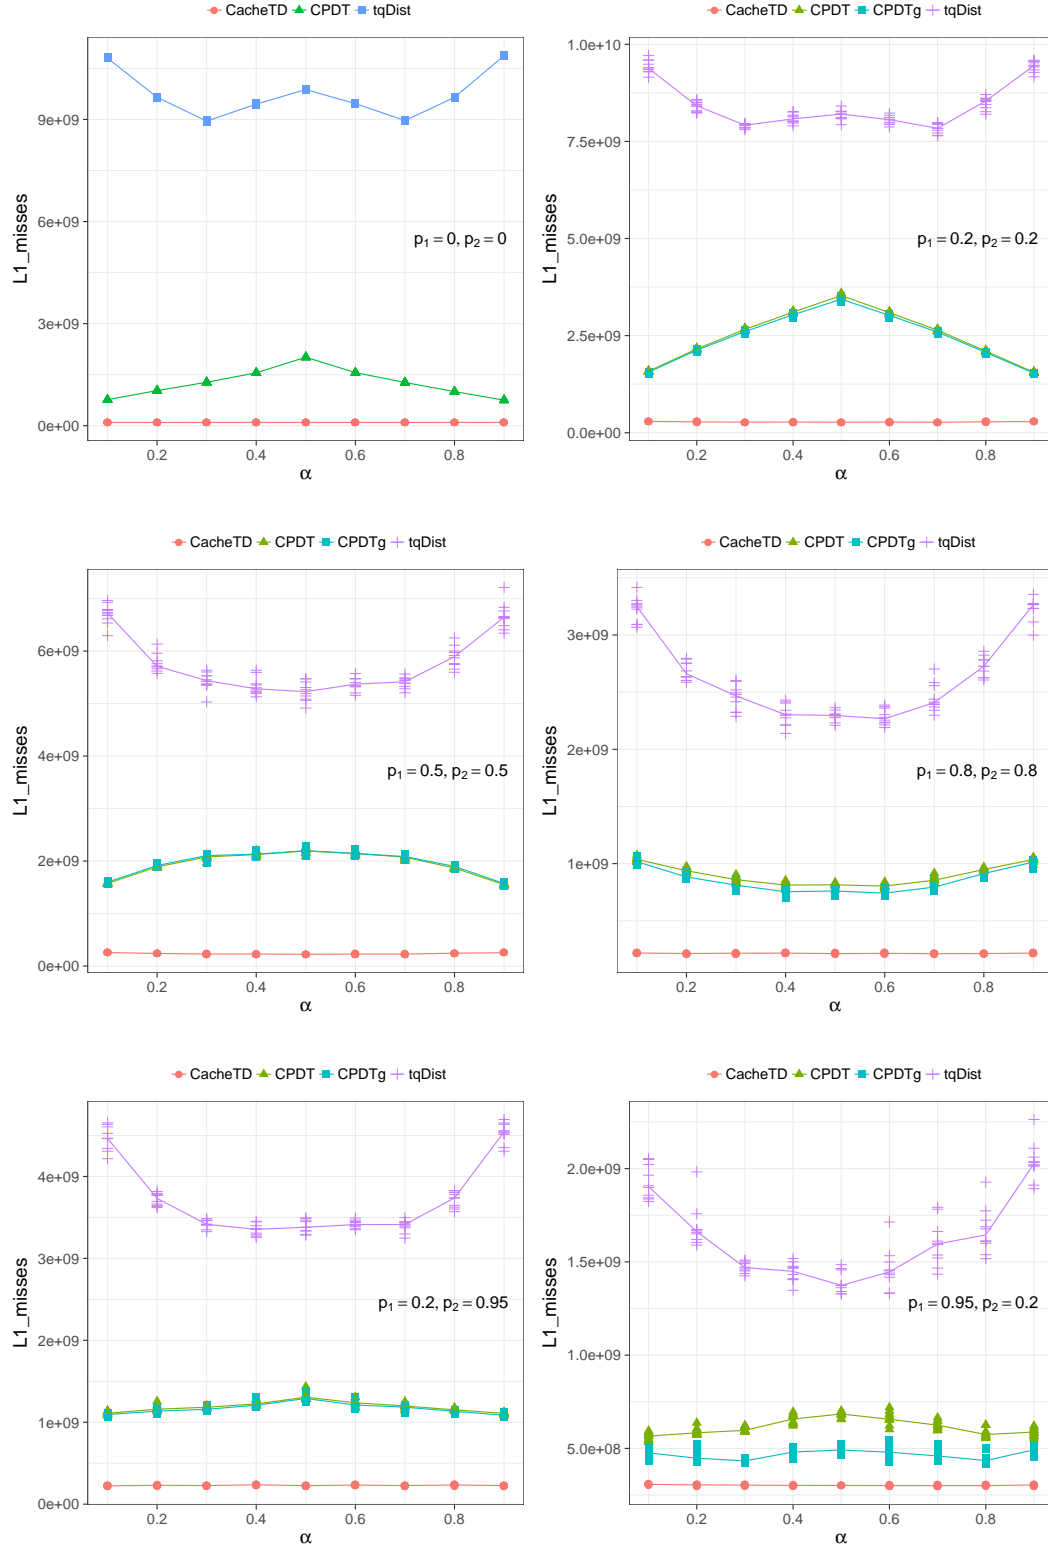
Figure 24 Skewed model: L1 cache misses ($n = 2^{21}$).
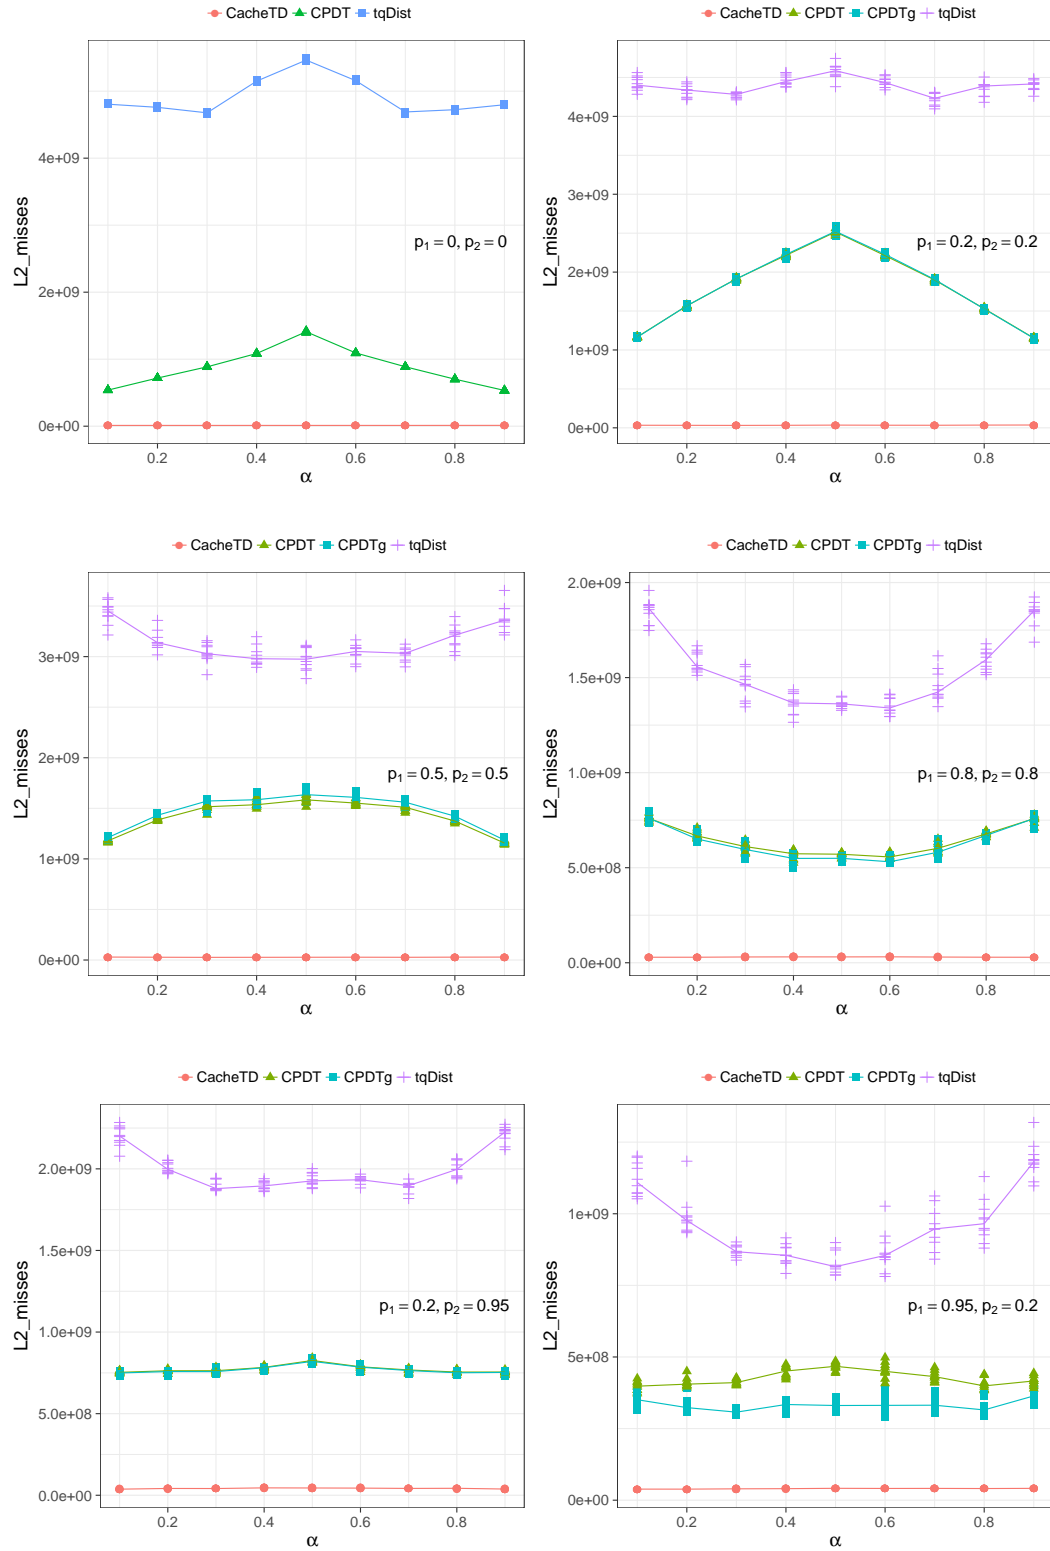
Figure 25 Skewed model: L2 cache misses $(n = 2^{21})$.

Figure 26 Skewed model: L3 cache misses ($n = 2^{21}$).
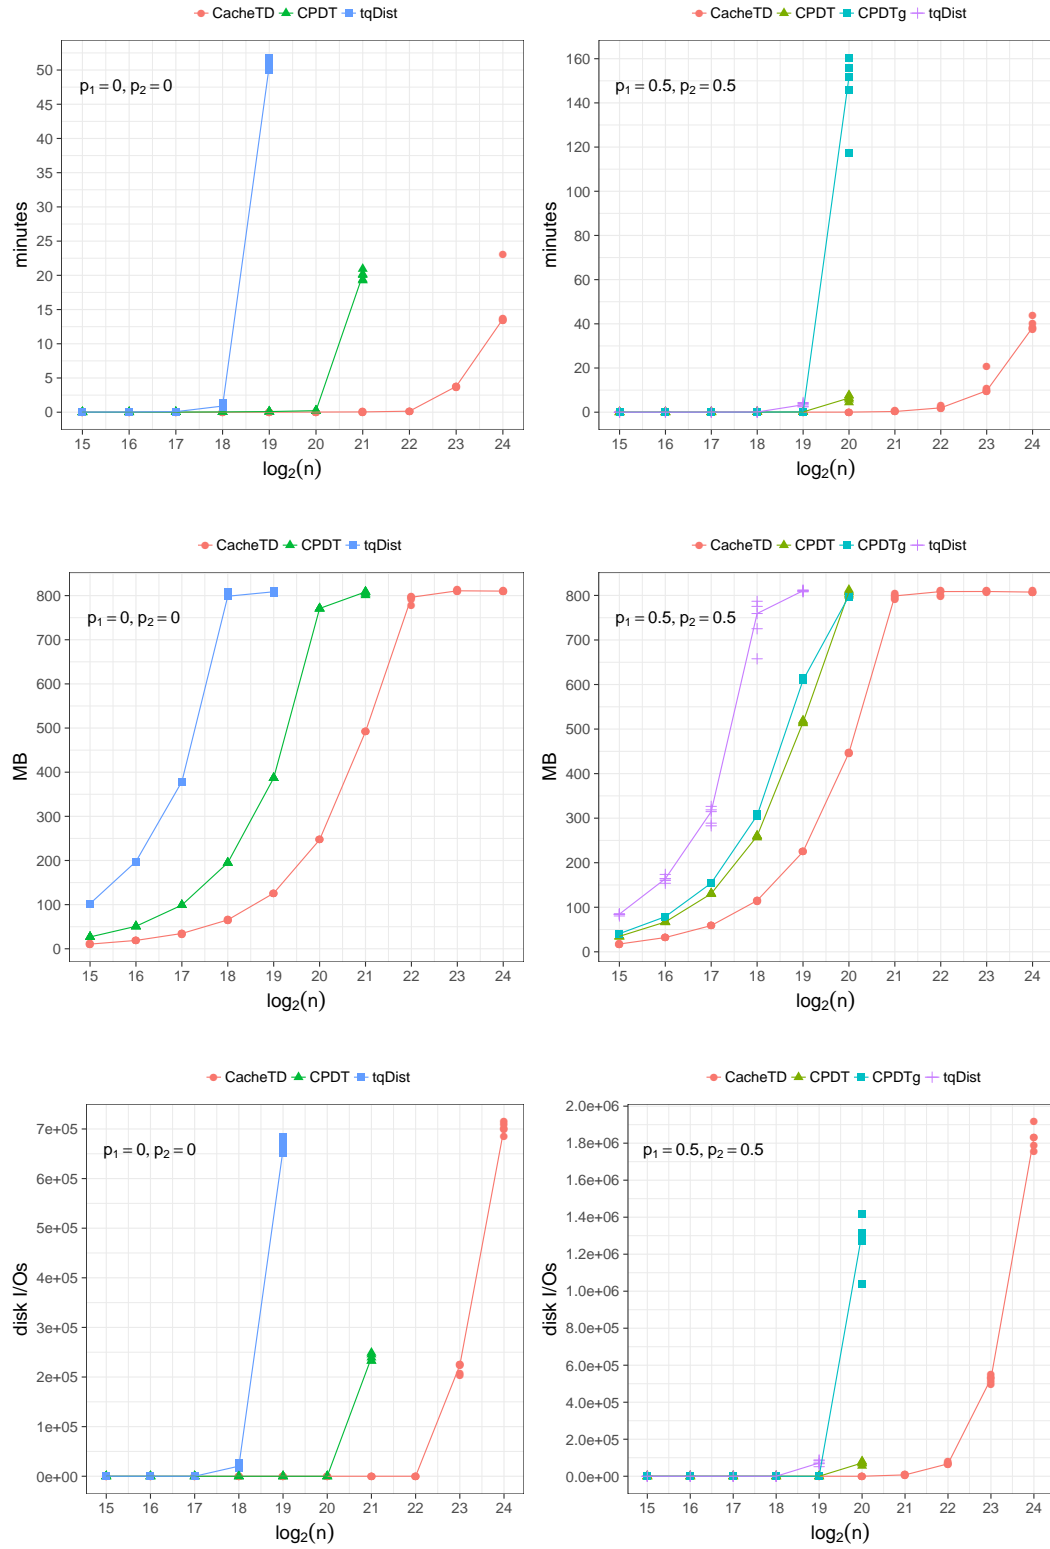
Figure 27 Random model: I/O experiments.

Figure 28 Skewed model: I/O experiments with $\alpha = 0.5$.