University of Groningen

A Hybrid Shared-Memory Parallel Max-Tree Algorithm for Extreme Dynamic-Range Images

Moschini, Ugo; Meijster, Arnold; Wilkinson, Michael

Published in:
IEEE transactions on pattern analysis and machine intelligence

DOI:
[10.1109/TPAMI.2017.2689765](https://doi.org/10.1109/TPAMI.2017.2689765)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2018

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Moschini, U., Meijster, A., & Wilkinson, M. (2018). A Hybrid Shared-Memory Parallel Max-Tree Algorithm for Extreme Dynamic-Range Images. *IEEE transactions on pattern analysis and machine intelligence*, 40(3), 513-526. <https://doi.org/10.1109/TPAMI.2017.2689765>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

A Hybrid Shared-Memory Parallel Max-Tree Algorithm for Extreme Dynamic-Range Images

Ugo Moschini^{id}, Arnold Meijster, and Michael H. F. Wilkinson, *Senior Member, IEEE*

Abstract—Max-trees, or component trees, are graph structures that represent the connected components of an image in a hierarchical way. Nowadays, many application fields rely on images with high-dynamic range or floating point values. Efficient sequential algorithms exist to build trees and compute attributes for images of any bit depth. However, we show that the current parallel algorithms perform poorly already with integers at bit depths higher than 16 bits per pixel. We propose a parallel method combining the two worlds of flooding and merging max-tree algorithms. First, a *pilot* max-tree of a quantized version of the image is built in parallel using a flooding method. Later, this structure is used in a parallel leaf-to-root approach to compute efficiently the final max-tree and to drive the merging of the sub-trees computed by the threads. We present an analysis of the performance both on simulated and actual 2D images and 3D volumes. Execution times are about 20× better than the fastest sequential algorithm and speed-up goes up to 30–40 on 64 threads.

Index Terms—Connected filters, hierarchical image representation, parallel algorithms

1 INTRODUCTION

MAX-TREES [1], or component trees, are versatile and efficient data structures that represent the connected components at every threshold level of an image in a hierarchical fashion, through parent relationships between nodes. Leaf nodes represent the image maxima. The connected components organised in such trees can be filtered with different strategies [1], [2] and they can model different types of connectivity [3]. For these reasons, they are powerful tools for image information mining and visualization tasks. Recent applications are in the field of astronomy [4], [5], [6], [7], [8] and remote sensing [9], [10], [11], to identify structures such as galaxies or building footprints from satellites. Fig. 1 shows an example of the segmentation of astronomical objects using a filter based on the noise statistics in the image [8].

Many current applications rely on high dynamic-range or floating point imagery, due to either the increasing sensitivity and technological improvements of the instruments or simply to the type of measurement observed, e.g., radio emissions and CT scans. Astronomical or remote sensing images routinely show high dynamic range integers or floating point values.

In the rest of the paper, we refer to images or volumes up to 64-bit integers or single or double precision floating point values per element as *XDR* (extreme dynamic range). In Section 5, we show that there is no state-of-the-art parallel algorithm suitable for building max-trees of *XDR* images. Existing parallel algorithms rely on partitioning the image and building sub-trees of each partition that are then merged

in the final max-tree. We will show that merging sub-trees of *XDR* images is costly and makes the parallel building phases of no use. In this paper, we provide a parallel algorithm which deals well with *XDR* images. It opens up the possibility of creating max-trees representing the whole image up to 40 times faster on 64 processors. It is not necessary to divide an image in smaller sections or lower its dynamic range. If the image is split, structures can be broken up and artifacts are expected. Finding the optimal way to quantize a priori the intensities is often not trivial. For example, CT scans and optical or radio astronomical images show objects close to the level of noise (e.g., the filament in Fig. 1): a complete and exact max-tree permits an accurate computation of measures [8], [12], [13] to identify accurately interesting structures and it can then be pruned at a later stage without losing important information. Furthermore, scale-invariant, shape filters used in e.g., vessel enhancement [14], or tree-based shape-space filtering [15] require complete max-trees.

Several sequential max-tree algorithms suitable for *XDR* images exist. They can be divided in two classes. Root-to-leaf flooding methods start from the root node, which is a pixel with the lowest intensity, and perform a depth-first traversal of the connected components at higher intensities [1], [16]. Leaf-to-root merging methods start from the image maxima [17] or work on sorted pixels that are merged into nodes of a tree using Tarjan's union-find algorithm [18], [19]. The parallel solution proposed in this work combines the most efficient root-to-leaf flooding and leaf-to-root merging algorithms to derive a two-stage, parallel algorithm suitable for *XDR* imagery. To tackle the problem of the costly merging phase, the idea is to create initially a *pilot* max-tree of a quantized version of the input image. The pixel intensities are quantized over a low number of values, equal to the number of threads of the parallel program. Using a low number of quantization levels, the existing parallel algorithm based on hierarchical queues [20] performs well in building the *pilot* max-tree. This tree is used as a support data structure during

- The authors are with the Johann Bernoulli Institute, University of Groningen, P.O. Box 407, Groningen 9700AK, The Netherlands.
E-mail: {u.moschini, a.meijster, m.h.f.wilkinson}@rug.nl.

Manuscript received 21 Dec. 2015; revised 23 Mar. 2017; accepted 24 Mar. 2017. Date of publication 29 Mar. 2017; date of current version 13 Feb. 2018. Recommended for acceptance by T. Coates.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPAMI.2017.2689765

Fig. 1. (a) Image containing stars and two large interacting galaxies; (b) output of connected filtering: identified objects and nested structures are coded in different colours.

a parallel *refinement* stage that yields the final max-tree of the input image. Without the support of the *pilot* tree, the high cost of merging the sub-trees together would still be present. The term “refinement” comes from the fact that the whole process of building the final max-tree can be seen as a refining operation on the *pilot* max-tree. In this last phase, the quantization is used to partition the pixels according to their intensities, rather than spatially.

2 THE MAX-TREE STRUCTURE

The image domain can be fully partitioned into connected components. Connected components are sets of path-wise connected pixels of maximal extent, having the same intensity value. When filtering an image, components can be preserved or removed according to a threshold value on a given attribute [2]. For example, regions with area smaller than a certain value or with a certain shape [21], [22] are deleted from the image. Efficient filtering is implemented by organising connected components in a tree structure, on which filtering reduces to pruning nodes in the tree and assigning new intensities. Several rules state how the nodes in the tree are removed and new intensities are assigned [1], [21]. Given an image f , a peak component P_h^i is defined as the set of all the pixels of a connected component of the thresholded image $T_h(f)$, where $T_h(f) = \{p \mid f(p) \geq h\}$. Of course, there can be many connected components at each level h . These components are indexed by i . Connected components can be organized in a max-tree structure, where each peak component P_h^i in the image corresponds to a node C_h^i . For example, Fig. 2c shows the max-tree structure of the 1-D image in Fig. 2b. The arrows denote a parent relation between the nested components.

3 STATE OF THE ART SEQUENTIAL ALGORITHMS FOR XDR IMAGES

We can divide the max-tree building algorithms into two categories: root-to-leaf *flooding* and leaf-to-root *merging*.

Following the latter approach, Najman and Couprie [19] described an algorithm to build a max-tree using Tarjan’s union-find [18] method. Berger et al. [5] published a more memory efficient version which we will refer to as the “Berger algorithm” in the following. The sequential Berger algorithm is suitable for finely quantized or even floating point images, i.e., XDR images. Both Tarjan’s and Berger algorithms will be explained in more details in Section 4.

Root-to-leaf algorithms follow the idea presented in the paper from Salembier et al. [1]. A depth-first construction for the tree is used, starting from the pixel with lowest intensity. The main drawback of this approach is faced precisely with XDR images, and is due to the fact that it uses hierarchical queues to handle the pixel values during the flooding. It scales linearly with the number of grey levels and with floating point values the worst-case time complexity is $O(N^2)$ [20], with N the number of pixels. Memory complexity is $O(N + G)$, with G the number of intensities, possibly equals to N in XDR images. Wilkinson [16] presented a new sequential algorithm that is based on a priority queue and a stack, based on a combination of the algorithms in [1] and [23]. It cancels the costly effect of updating the hierarchical queues when dealing with XDR images. The memory usage is simply linear in the data size. Therefore, the size of the data structures does not scale proportionally to the range of the possible intensity values. A similar problem was addressed with the leaf-to-root algorithm in [17] that uses a priority queue that is intrinsically $O(N^2 \log N)$ [24]. Both the approaches by Berger et al. and by Wilkinson deal well with XDR images in the sequential case, with the former faster than the latter, but with a larger memory footprint.

4 TWO IMPORTANT MAX-TREE ALGORITHMS

As mentioned in the Introduction, our new parallel algorithm is based on a two-step process involving building the *pilot* max-tree of a quantized version of the input image and a refinement phase that builds the final tree of the original image. We will describe our hybrid algorithm in Section 6, but for now it is important to point out that the two steps are based on two existing algorithms. The first step is based on the parallel method in [20] (see Section 4.3), while the second step, that forms the core of our novel parallel solution, is based on the sequential Berger algorithm (see Section 4.2). A good knowledge of these two algorithms is essential to fully understand how we modified and combined them together to our end. Both algorithms are based on the Tarjan’s union-find algorithm [18] introduced in the following section.

4.1 The Tarjan’s Union-Find Algorithm

The union-find algorithm by Tarjan [18] was designed to keep track of disjoint sets. Image connected components do not overlap, and therefore are examples disjoint subsets of the image domain. Each set is described by a tree, rooted in an arbitrary element of the set chosen to be the representative for that set. Each root points to itself, while the other elements point to their parent. To determine if two elements belong to the same set, it is sufficient to check if they share the same root. The shorter the root path, the more efficient this becomes. Tarjan proposed an efficient algorithm to maintain a collection of disjoint sets. Its complexity is quasi-

Fig. 2. (a) Represents a 16-bit integer 1D image. (b) is a quantized image after mapping the original intensities onto 4 intensities. (c) illustrates the max-tree structure of connected components of the quantized image shown in (b).

linear [18]. His method uses three operations on the set elements p , q , namely $\text{MakeSet}(p)$, $\text{FindRoot}(p)$, $\text{Union}(p, q)$. $\text{MakeSet}(p)$ makes the p element a singleton set; $\text{FindRoot}(p)$ returns the root of the tree containing p ; $\text{Union}(p, q)$ implements the merging of the two sets containing elements p and q . A union-find algorithm used for the computation of attribute openings and closings was proposed in [24] where it was suggested it could possibly be suitable for parallel implementation. Pixels are processed in grayscale order, so peak components could be processed simultaneously. Peak components are created and then merged, their attributes being updated. Najman et al. [19] later adapted the union-find approach to build max-trees.

Algorithm 1. Pseudo-Code of the Berger Algorithm. The Output is the Max-Tree of Image f , in the *node* Array

```

1: procedure RUN BERGER Image ( $f$ )
2:    $S \leftarrow \text{SORTPIXELSDECREASING}(f)$ ;
3:   for all pixel  $p \in S$  do
4:      $\text{node}[p].\text{parent} \leftarrow p$ ;  $\text{zpar}[p] \leftarrow p$ ;
5:     for all pixel  $q$  neighbours of  $p$  with  $\text{zpar}[q] \neq -1$  do
6:        $r \leftarrow \text{FINDROOT}(q)$ ;
7:       if  $r \neq p$  then
8:          $\text{zpar}[r] \leftarrow p$ ;  $\text{node}[r].\text{parent} \leftarrow p$ ;
9:          $\text{node}[p].\text{Area} \leftarrow \text{node}[p].\text{Area} + \text{node}[r].\text{Area}$ ;
10:      end if
11:    end for
12:  end for
13: end procedure
14: procedure FINDROOT(Pixel  $p$ )
15:   if  $\text{zpar}[p] \neq p$  then
16:      $\text{zpar}[p] \leftarrow \text{FINDROOT}(\text{zpar}[p])$ ;
17:   end if
18:   return  $\text{zpar}[p]$ ;
19: end procedure
20: procedure INIT(Image  $f$ )
21:   for all pixel  $p \in f$  do
22:      $\text{zpar}[p] \leftarrow -1$ ;
23:      $\text{node}[p].\text{parent} \leftarrow -1$ ;
24:      $\text{node}[p].\text{Area} \leftarrow 1$ ;
25:   end for
26:   RunBerger( $f$ );
27: end procedure

```

In the rest of the paper, any rooted tree will be represented as an array of node structures. Each structure contains a parent pointer *parent* and the area attribute value in *area*. Other

attributes can be computed as explained in [2], but we used area for simplicity. This representation involves one node structure per image pixel. The execution of the union-find method will set the parent pointers accordingly to the connected components (disjoint sets). Actually, only the nodes which have a parent node at a lower intensity are necessary to represent the whole tree and hold the correct attribute values. These nodes are called *level roots* [20] (known also as canonical elements in [19]).

4.2 The Sequential Berger Algorithm

The sequential Berger algorithm [5] uses the Tarjan's method and structures to build the max-tree. Parent pointers in the tree are correctly updated via the union-find algorithm. For a fast computation of the parent pointers, root nodes of every disjoint set must be accessed efficiently (path compression). The FINDROOT procedure at lines 14-19 in Algorithm 1 is now also used to implement path compression: the root node of a set is efficiently found by letting each element in the set point directly to it. Since path compression cannot be applied on the parent pointers directly because it would destroy the hierarchy of the nodes, an auxiliary data structure is used: the array *zpar* with length equal to the number of elements (pixels) in some image f . Therefore, path compression is applied onto the *zpar* structure, that contains the root nodes of disjoint sets while they are computed. As Algorithm 1 shows, function INIT(f) initialises the data structures *zpar* and *node*. The area attribute is set to 1 for all the pixels (singleton sets). Then, the Berger algorithm is run: pixels are sorted and retrieved from S in decreasing order. Function RUNBERGER(f) ultimately returns the tree with the correct parent pointers. When a pixel p is retrieved from the sorted array, it is marked as processed by setting $\text{zpar}[p] = p$ (line 4 in Algorithm 1). Following the way the union-find algorithm works, the already processed neighbours of p are checked and the set made of just p is merged with the neighbouring existing sets, updating the area attribute, to form a new connected component rooted in it. That is the core of the union-find procedure. On floating point images, the building algorithm by Najman and Couprie mentioned in Section 3 is said to have better performance [5] than the Berger one thanks to a technique called union-by-rank [18]. Union-by-rank avoids creating degenerate trees in flat zones by keeping as small as possible the depth of the trees. This problem arises because the last processed pixel always becomes the new root. A detailed

explanation is in [18]. If the Najman and Couprie algorithm is modified by disabling union-by-rank, then the Berger algorithm becomes not only more memory efficient, but also faster. We will show that our two-step parallel solution does not use union-by-rank, thus making the Berger algorithm the best choice.

Algorithm 2. Concurrent Construction of the Max-Tree for Thread th on K Threads

```

1: procedure PARALLEL_HIERARCHIAL_QALGORITHM (Thread
    $th$ , Image partition  $P$ )
2:    $m \leftarrow \text{argmin}(P)$ ;  $\triangleright$  pixel in  $P$  with minimum intensity
3:   Add  $m$  to the Queue at level  $P(m)$ ;
4:    $isVisited[m] \leftarrow \text{true}$ ;
5:    $levelroot[P(m)] \leftarrow m$ ;
6:   FLOOD( $P(m), P, 0$ );  $\triangleright$ see Algorithm 3
7:    $i \leftarrow 1$ ;  $q \leftarrow th$ ;
8:   while ( $th + i < K$ )  $\wedge$  ( $q \% 2 = 0$ ) do
9:     Wait to glue with right-hand neighbour;
10:    for all Edges ( $u, v$ ) b/w partition  $P_{th}$  and  $P_{th+i}$  do
11:      CONNECT( $th, i, (u, v)$ );  $\triangleright$ see Algorithm 4
12:    end for
13:     $i \leftarrow 2 \cdot i$ ;  $q \leftarrow q/2$ ;
14:  end while
15:  if  $th \neq 0$  then
16:    Notify left-hand neighbour;
17:    Wait for Thread 0;
18:  end if
19: end procedure

```

4.3 The Parallel Hierarchical Queue Algorithm

The algorithm reported in this section is a parallelization of the sequential algorithm in [1]. A thorough description is found in the work by Wilkinson et al. [20]. It follows a root-to-leaf approach and it is based on a hierarchical FIFO queue. The queue *Queue* has a number of entries equal to number of levels (intensities) in the image. It is initialised with the lowest intensity pixel in the image at line 3 in Algorithm 2 and the recursive root-to-leaf flooding starts from there. The flooding process populates the queue. The parallelization of a queue-based algorithm is never trivial, but it is possible to partition the image f into sections P , build trees of every section and then merge them together. In the sequential algorithm [1], pixels were given arbitrary labels to signal that they belong to a certain component. That makes it hard to implement the merging, because pixels need to be relabelled. The algorithm in [1] was adapted to use the same tree structures as in the Tarjan's method and to be consistent with the union-find approach by letting each element of a component point directly to the first pixel found. This allows a union-find type merging of the nodes in the merging phase. Algorithm 2 shows how the max-tree of a partition P of f assigned to each thread is built and how the data structures are initialised. Array *isVisited* keeps track of the visited pixels, while array *levelroot* is used to hold the latest level root at the currently flooded intensity level. Through the recursive approach, the hierarchy is built branch by branch, keeping track of the level roots of the currently explored branch. The flooding procedure is illustrated in Algorithm 3. We know that each connected

component is characterised by a unique element (level root): the set of nodes that an element belongs to is rooted in such unique element. Path compression is implemented by making the elements point directly to the level root (see line 11 in Algorithm 3). The Union routine that sets the root of one of the trees to the root of the other one is implemented in the flooding procedure by adjusting the parent pointers. Our *pilot* max-tree construction described in Section 6.1 is based on a modified version of the FLOOD. It will be shown that the only difference lies in the definition of level root and the way the array *levelroot* is handled.

Algorithm 3. FLOOD Procedure Implements the Root-to-Leaf Flooding Approach Based on a Priority Queue. The Procedure Uses *thisarea* as Reference Parameter as in C. The Output is the Max-Tree of the Image Partition P in the *node* Array

```

1: procedure FLOOD (Level  $lev$ , Partition  $P$ , Attribute this area)
2:    $area \leftarrow \text{thisarea}$ ;
3:   while Queue at level  $lev$  is not empty do
4:     Extract pixel  $p$  from the Queue at level  $lev$ ;
5:      $area \leftarrow area + 1$ ;
6:     for all neighbours  $q$  of  $p$  with  $q \in P$  do
7:       if  $isVisited[q] = \text{false}$  then
8:          $isVisited[q] \leftarrow \text{true}$ ;
9:         if  $levelroot[f(q)] = \text{not set}$  then
10:           $levelroot[f(q)] \leftarrow q$ ;
11:        else
12:           $node[q].parent \leftarrow levelroot[f(q)]$ ;
13:        end if
14:        Add  $q$  to the Queue at level  $f(q)$ ;
15:        if  $f(q) > lev$  then
16:           $childarea \leftarrow 0$ ;  $f_q \leftarrow f(q)$ ;
17:          while  $f_q > lev$  do
18:             $f_q \leftarrow \text{FLOOD}(f(q), P, childarea)$ ;
19:          end while
20:           $area \leftarrow area + childarea$ ;
21:        end if
22:      end if
23:    end for
24:  end while
25:   $m \leftarrow lev - 1$ ;
26:  while  $m \geq 0 \wedge levelroot[m] = \text{not set}$  do
27:     $m \leftarrow m - 1$ ;
28:  end while
29:  if  $m \geq 0$  then
30:     $node[levelroot[lev]].parent \leftarrow levelroot[m]$ ;
31:  end if
32:   $node[levelroot[lev]].Area \leftarrow area$ ;
33:   $levelroot[lev] \leftarrow \text{not set}$ ;  $thisarea \leftarrow area$ ;
34:  return  $m$ ;
35: end procedure

```

Once the trees of the partitions have been built, the issue is how to merge them together. All the trees are stored in the *node* array: each root node of the partition trees has a parent pointer equal to -1 . Thread 0 eventually computes the final merging that yields a correct max-tree of f . Lines 8-18 of Algorithm 2 refer to this process. The crucial step is the procedure CONNECT invoked at line 11 of Algorithm 2 and reported in Algorithm 4 for each pair of neighbouring nodes u, v

between two adjacent partitions. Their parent pointers are followed down in the tree until they meet in a node or in the root node. If u and v belong to a component that was split in two partitions, one of the two level roots is chosen as the new representative and pointers and area value are updated in the chosen level root, with *area* and *areatemp* variables ensuring a correct accumulation of the attribute value in the hierarchy. The updating is propagated down till the root of the tree or until the same level root is reached (see the *while* loop at line 8 of Algorithm 4). The implementation of this phase is not trivial. It shows similarities with the merge sort algorithm. We refer to the work in [20] for a detailed explanation and proof of correctness. For our purposes, however, it is important to stress the fact that while working nicely on images with low quantization (8, 12 or 16 bits per pixel), this merging approach turns out to be impracticable in case of XDR images. The exact merging points are checked parsing the trees from the leaf nodes till the root node, in the worst case, for every edge. Therefore, the complexity of a **CONNECT** operation grows exponentially with the bit depth, as shown in the following Section 5. Thus, we propose a different approach to implement the merging of sub-trees in our two-step parallel algorithm.

Algorithm 4. Code of the **CONNECT** Procedure. Symbol $\perp$ is Defined as the Root Node of Every Sub-Tree and $f(\perp) = -\infty$

```

1: procedure CONNECT(Thread  $th$ , Edge  $(u, v)$ )
2:    $area \leftarrow 0$ ;  $areatemp \leftarrow 0$ ;
3:    $x \leftarrow \text{GETLEVELROOTOF}(u)$ ;
4:    $y \leftarrow \text{GETLEVELROOTOF}(v)$ ;
5:   if  $f(x) < f(y)$  then
6:     Swap( $x, y$ );
7:   end if
8:   while  $x \neq y \wedge x \neq \perp$  do
9:      $z \leftarrow \text{GETLEVELROOTOF}(node[x].parent)$ ;
10:    if  $f(z) \geq f(y) \wedge z \neq \perp$  then
11:       $node[x].Area \leftarrow node[x].Area + area$ ;
12:       $x \leftarrow z$ ;
13:    else
14:       $areatemp \leftarrow node[x].Area + area$ ;
15:       $area \leftarrow node[x].Area$ ;
16:       $node[x].Area \leftarrow areatemp$ ;
17:       $node[x].parent \leftarrow y$ ;  $x \leftarrow y$ ;  $y \leftarrow z$ ;
18:    end if
19:  end while
20:  if  $y = \perp$  then
21:    while  $x \neq \perp$  do
22:       $node[x].Area \leftarrow node[x].Area + area$ ;
23:       $x \leftarrow \text{GETLEVELROOTOF}(node[x].parent)$ ;
24:    end while
25:  end if
26: end procedure

```

5 ISSUES WITH POSSIBLE PARALLEL ALGORITHMS FOR XDR IMAGES

While looking for a parallel version of the state-of-the-art sequential algorithms presented in Section 3 that can deal with very high-dynamic range images, we made a first attempt implementing straight away a parallel version of

Fig. 3. (a) Shows the performance of the parallel algorithm based on priority queues on images with randomly generated floating point values. The two-threaded version performs better than the single-threaded only up to 16-20 bits per pixel. With higher bit depths, the cost of merging cancels any benefit. (b) Speed-up is close to its optimal value up to about 16 bits per pixels, then it drops dramatically.

the method by Wilkinson [16], based on priority queues that uses the same merging method as the hierarchical queue algorithm of Section 4.3. That is, the merging phase for XDR images becomes very costly and makes the parallel algorithm of no use when dealing with floating point or 32-bit integer images. Even though the sequential Berger algorithm could be used to build the max-trees of the image partitions, this would not solve the problem because it lies in the merging phase, which is independent of how the sub-trees are built. Fig. 3 shows indeed that higher bit depths cause a huge drop in performance already on two threads, at a bit depth of about 16-20 bits per pixel. Also in [25], the exponential complexity of merging the sub-trees as the number of bits increases was shown for several building algorithms and it was stated that parallel algorithms are unsuitable for data at high bit depths.

The algorithm by Berger et al. [5] does not lend itself in a natural way to be parallelized, due to the way an image should be partitioned: the pixels are in fact sorted by intensity and retrieved in descending order during the leaf-to-root process. That leads to a partitioning criterion based on

Fig. 4. Given a 1-D image, an example of spatial partition of the pixels is given in (a). In (b), the image is partitioned according to ranges of intensities.

pixel intensity rather than pixel position. A spatial partition would in principle be possible with a trivial implementation: the preservation of the ordering could be achieved by maintaining a semaphore for each pixel, as remarked in [20]. Unfortunately, that would make the approach impracticable because of the high overhead due to the many synchronizations and locking sections needed. Another solution could be to assign a thread to every maximum in the image. Again, the ordering could be preserved by putting a barrier synchronization method on every intensity level change: this way, the leaf-to-root process would proceed in parallel, intensity by intensity. Mutual exclusion structures would be needed on every pixel that shows a level change with respect to its neighbouring pixels, leading in the worst case to having a barrier for every pixel and therefore no parallelism at all. Due to the high number of mutexes, it does not look like a feasible algorithm. Therefore, those possibilities were not explored any further.

6 THE PARALLEL HYBRID ALGORITHM

To tackle the merging problem, we start first by computing a quantized version $\bar{f}$ of the input image f . A tree called *pilot* max-tree can be built in parallel using any existing algorithm due to the low number of quantized values. The merging phase is not an issue here. We chose the algorithm illustrated in Section 4.3, because we found it faster than other methods on such low-quantized values. The image is divided into spatial partitions as in Fig. 4a, one for each thread. The next step is a parallel refinement stage that yields the final max-tree, with the correct attributes. This process can be seen as a “refinement” operation that shapes the tree of the quantized image into the max-tree of the original image. The *pilot* max-tree allows for correct attribute computation and merging of the sub-trees. Each thread of the refinement stage uses the sequential Berger algorithm of Section 4.2 on the pixels of its partition and takes advantage from the *pilot* max-tree to compute correctly the attributes and merge with the other sub-trees that are being built by the threads. In this parallel approach, a trivial spatial partition is not suitable, because leaf-to-root algorithms like the Berger algorithm work on pixels that have been ordered by intensity. The image must be partitioned into spatially irregular partitions as in Fig. 4b, so that each partition contains only the pixels in a given range of intensities and each thread can start the execution from the local maxima present in a given range. As Fig. 5 shows, $S_0, \dots, S_{K-1}$ form K partitions. S_0 is defined as the set of all the pixel locations with intensities in the interval $[h_0, h_1)$, S_1 contains the pixels

Fig. 5. Pixels are sorted according to their intensity, from low to high. A partition S_i is made of the pixels with intensity values within h_i and h_{i+1} .

with values in the interval $[h_1, h_2)$ and so on, where $h_0 < h_1 < h_2$ and h_0 is the lowest intensity in the image. The set of the intensity values in S_i is indicated with H_i , representing the interval $[h_i, h_{i+1})$. Let K be the number of threads of the parallel refinement stage. The total number of partitions that is created must be equal to K , so that every thread handles all the pixels within a given range of intensities. The intensity ranges H_i are chosen so that about the same number of pixels is present within each partition. It is not always possible to achieve a perfect load balance, since it is constrained by the fact that there cannot be any overlap among the intensity ranges in two different sets H_i . Since the refinement stage of our parallel solution is based on the Berger algorithm, pixel coordinates will need to be sorted by their intensity value in the *original* image. A parallel stable radix sort [26] was used. It relies on the parallelization of the inner counting sort algorithm [27] used iteratively within the radix sort. Radix sort was originally developed for integers: we adapted it to support floating point values following the methods in [28] and [29], also mentioned in other works [30], [31]. We are going to explain more in detail the *pilot* max-tree and the refinement stage. We highlight that without the support of the *pilot* tree, the fact of having intensity-based partitions is not enough per se to avoid the high cost of merging the sub-trees of each S_i . When two regions that belong to a given partition are separated by a peak with maximum higher than the partition they belong to, it would be necessary to go down through all the intensity levels to check at which point the regions should be merged together. The *pilot* max-tree structure was designed exactly to reduce the cost of such operations.

6.1 The Pilot Max-Tree

The *pilot* max-tree is simply the max-tree of the quantized image $\bar{f}$. We chose the root-to-leaf approach by [20], illustrated in Section 4.3, to build it. The nodes of the *pilot* max-tree are stored in the array *node_qu*, whose length is equal to the number of pixels. Due to the low number of intensity levels in $\bar{f}$, the merging phase does not represent a problem and the algorithm of Section 4.3 is perfectly suitable. We underline that the quantized image does not show new peak components that are not present in the original image. In fact, according to our definition of quantization, the actual hierarchy of components in the original image now belonging to a partition S_i is simply flattened on the ancestor component with lowest intensity in S_i . The same ancestor component is then present in the original image as well as in the quantized image. The flattening can be observed comparing Figs. 2a and 2b. Exploiting the fact that there are no new peak components, a relation is enforced between the *level root* nodes (introduced in Section 4.1) of the *pilot* tree and those of the final refined tree: the former are a subset of the latter. We will see how the merging phase of the

sub-trees and their attributes relies on such nodes. To ensure this correspondence, the only change made on the max-tree implementation in [20] is a stricter definition of level root of a connected component. The level root node where the attributes (area, in our case) are accumulated is not any more an arbitrary pixel of the component but it is chosen to correspond to the pixel with lowest coordinate among the pixels belonging to the component. Function FLOOD_PILOT in Algorithm 5 is the same as in Algorithm 3 except for the call to function KEEPLEASTLEVELROOT, detailed in Algorithm 6, that implements the new definition of level root. At the end, the array *node_qu* contains the *pilot* max-tree. The loop with the recursive call to FLOOD_PILOT at line 17 in Algorithm 5 stops when a local maximum has been reached. We make explicit here that, in the refinement stage, the last processed pixel of every connected component in the *original* image *f* will be the one with the lowest image coordinate. Since a stable sort was chosen, the pixels with equal value in every original component are sorted so that the pixel with the lowest coordinate appears after all the others with equal intensity. The stricter definition of level root, the quantization step and the stable sort grant that the level roots of the *pilot* max-tree are a subset of the level roots of the refined tree. We note that union-by-rank that makes the algorithm by Najman and Couprie faster [5] than the Berger one cannot be used in our solution: the correspondence among the level roots in the *pilot* and in the final max-tree would not be guaranteed any more. Union-by-rank could possibly be tweaked, but at the cost of checking for every processed pixel a correspondence to a level root and, anyway, not always returning a balanced set, due to this restriction.

6.2 The Refinement Stage

The correct max-tree representation of the original image *f* is the output of the refinement stage. In the refinement stage, the *pilot* max-tree is not modified. Hence, all threads can safely access it without need for synchronization mechanisms. It can be seen indeed as a parallel version of the algorithm proposed by Berger et al. [5]. The final max-tree nodes (the tree structure) are stored in an array called *node_ref*, with length equal to the image size. Every node corresponds to a pixel, as in the *pilot* max-tree, with the same structure as before (parent pointer, the area attribute value) with the addition of a value to store the intensity of the pixel after filtering. The whole algorithm is detailed in Algorithm 7. As recalled in Section 4.2, our solution uses the same array structure named *zpar* in Section 4.2. The final tree is created in parallel by *K* threads that work on the original pixel values of *f*. The number of threads is equal to the number of partitions *S*. Each thread *T_i* retrieves the sorted pixels belonging to its partition *S_i*, each one with the same quantized intensity, in descending order. Computation starts from pixels corresponding to the local maxima of the partition *S_i*. Like in the sequential Berger algorithm, for every pixel *p* retrieved from the sorted array and belonging to the partition *S_i*, the set of its neighbour pixels is calculated. As shown in Section 4.2, in the Berger algorithm only the already processed neighbours are considered: given that only the pixels with intensity equal to or larger than the current pixel's intensity could have been retrieved (processed)

from the sorted array, only two situations are possible, described in Sections 6.2.1 and 6.2.2.

Algorithm 5. FLOOD_PILOT Builds the *pilot* Max-Tree

```

1: procedure FLOOD_PILOT (Level lev, Partition P, Attribute
   this area)
2:   area  $\leftarrow$  thisarea;
3:   while Queue at level lev is not empty do
4:     Extract pixel p from the Queue at level lev;
5:     area  $\leftarrow$  area + 1;
6:     for all neighbours q of p with q  $\in$  P do
7:       if isVisited[q] = false then
8:         isVisited[q]  $\leftarrow$  true;
9:         if levelroot[g(q)] = not set then
10:          levelroot[g(q)]  $\leftarrow$  q;
11:        else
12:          KEEPLEASTLEVELROOT(q);
13:        end if
14:        Add q to the Queue at level  $\bar{f}(q)$ ;
15:        if  $\bar{f}(q) > lev$  then
16:          childarea  $\leftarrow$  0; fq  $\leftarrow$   $\bar{f}(q)$ ;
17:          while fq > lev do
18:            fq  $\leftarrow$  FLOOD_PILOT ( $\bar{f}(q)$ , P, childarea);
19:          end while
20:          area  $\leftarrow$  area + childarea;
21:        end if
22:      end if
23:    end for
24:  end while
25:  m  $\leftarrow$  lev - 1;
26:  while m  $\geq$  0  $\wedge$  levelroot[m] = not set do
27:    m  $\leftarrow$  m - 1;
28:  end while
29:  if m  $\geq$  0 then
30:    node_qu[levelroot[lev]].parent  $\leftarrow$  levelroot[m];
31:  end if
32:  node_qu[levelroot[lev]].Area  $\leftarrow$  area;
33:  levelroot[lev]  $\leftarrow$  not set; thisarea  $\leftarrow$  area;
34:  return m;
35: end procedure

```

Algorithm 6. Implements the Stricter Definition of Level Root

```

1: procedure KEEPLEASTLEVELROOT (Pixel q)
2:   cond1  $\leftarrow$  f(q) < f(levelroot[ $\bar{f}(q)$ ]);
3:   cond2  $\leftarrow$  f(q) = f(levelroot[ $\bar{f}(q)$ ])  $\wedge$  q < levelroot[ $\bar{f}(q)$ ];
4:   if cond1  $\vee$  cond2 then  $\triangleright$  set the new lowest level root
5:     node_qu[levelroot[ $\bar{f}(q)$ ]].parent  $\leftarrow$  q;
6:     levelroot[ $\bar{f}(q)$ ]  $\leftarrow$  q;
7:   end if
8:   node_qu[q].parent  $\leftarrow$  levelroot[ $\bar{f}(q)$ ];
9: end procedure

```

6.2.1 Intensity $\bar{f}(q) \in H_i$

If the intensity $\bar{f}(q)$ of the neighbour pixel *q* belongs to the set *H_i*, defined at the beginning of Section 6, managed by thread *T_i*, the computation of the tree continues as in the sequential Berger algorithm. This is detailed at lines 22-29 of the pseudocode in Algorithm 7.

Fig. 7. (a) Original image; (b) a quantization of (a); (c) pixel p is extracted from the sorted array and its neighbour q that belongs to a higher partition is processed. Pixel a is the closest descendant of p , returned by DESCENDROOTS. If `node_ref[a]` has no parent set yet, then the pixel p is assigned as parent, see white dotted arrow. (d) Later, pixel p_1 will be extracted, its neighbour q processed and once again a is returned by DESCENDROOTS. Its parent pointer was set in step (c). Therefore, as in the Berger algorithm, pixel p_1 is accessed and p_1 is set as parent of `node_ref[p]`.

DESCENDROOTS function will always return two different descendant nodes. In fact, the closest descendant of two components at two different quantized intensity levels simply cannot be the same. The DESCENDROOTS function itself has a read-only access to the `node_qu` structure. Calls to the FINDROOT function are also thread-safe. In fact, every thread sets only the entries of the `zpar` structure corresponding to the locations of the pixels under its partition: no mutex structure is necessary on `zpar`. The only mutexes required by both the *pilot* max-tree building and the refinement stage are only the ones required to flag that the threads completed their task and processed all the pixels in their partition.

7 TIME COMPLEXITY AND MEMORY USE

The time complexity of the parallel algorithm is defined by the complexity of its three main steps: sorting the pixels, building the *pilot* max-tree, and building the final refined tree. The parallel radix sort can be decomposed into three phases, each one executed at each iteration of the radix sort. In the first phase, each thread in parallel implements the counting sort on its local partition: the complexity of this operation is $O(N/K)$, where N is number of pixels in the image and K the number of threads. When every thread has completed its partial histogram, thread T_0 computes the whole image histogram by summing the local histograms in $O(K \cdot 2^r)$, calculates the prefixes in $O(2^r)$ and signal them to the other threads in $O(K \cdot 2^r)$. Value r is equal to 16: decomposing the image data types in chunks of 16 bits has proven to result in a cache-friendly size of the histograms that are used in the counting sort algorithms. It also limits the number of iterations to at most four in case of 64 bit double values. Lastly, every thread writes the pixel positions in the sorted array in parallel in $O(N/K)$. A similar and more detailed analysis was presented in [27].

The complexity related to building the *pilot* max-tree in parallel can be summarised as follows. On K threads, the

worst case time complexity of the building phase of each tree for every partition of the quantized image (therefore with $G = K$ intensities) is $O(N(C + G)/K)$ with C the connectivity and N the number of pixels, which reduces to $O(NC/K + N)$, since $G = K$. Interestingly, the second term of the addition shows no parallelism. We recall here that this is a worst-case scenario, that is rarely seen in practice. We refer to the description in [20], [24] for more details. As for the merging phase, if K neighbouring partitions have J bridging edges, the total merging has complexity $O(JG \log N \log K)$ [20]. For every merge, we need to go down $O(G)$ nodes, in a worst case scenario. Since the image was quantized in K intensities, usually equal to 16, 32, 64, the merging phase is not excessively costly.

The refinement stage shows aspects in common with the Berger algorithm and its complexity analysis is similar to the one in [19]. It is known from [5] that the algorithm by Berger et al. has quasi-linear complexity, once the pixels have been sorted, if union-by-rank [18] is applied. In our case, as mentioned at the end of Section 6.1, the ranking technique is disabled and the Berger algorithm has $O(N \log N)$ complexity. In our parallel implementation, the only difference with the original sequential Berger algorithm is the function DESCENDROOTS that parses the *pilot* max-tree. This function is not in the set of operations defined by the union-find algorithm for which quasi-linearity was demonstrated in [18]. Its complexity is then analysed in the following. For every neighbouring pixels with intensity larger than the intensity of the current pixel, the node hierarchy in the *pilot* max-tree is parsed from the quantized level of the neighbour to the quantized level of the current pixel. Therefore, the time complexity of DESCENDROOTS cannot be larger than $O(GCN/K)$ with C number of neighbours and K the number of threads. Since $G = K$, the complexity is equal to $O(CN)$, linear in the number of pixels in the worst case scenario with CN/K higher neighbouring pixels.

Memory-wise, the parallel algorithm needs two arrays to store the original image and the quantized image, the sorted array and the `zpar` array: they require together $4N$ memory space, assuming that pixel values, pixel coordinates and (area) attribute are encoded with the same number of bits. Moreover, the *pilot* max-tree and the final tree require two arrays of length N , the former with each entry containing the parent index and the attribute value, the latter the parent index, the attribute value and the value after filtering. In total, they require additional $5N$ memory space. The sequential Berger algorithm required $6N$ memory space: it needs the array for the original image, the sorted array, the `zpar` array and the tree structure as an array of length N , each one containing attribute value, parent index and value after filtering. Summarising, the parallel algorithm needs $3N$ memory more than the Berger algorithm due to the arrays to store the quantized image and the *pilot* max-tree.

8 PERFORMANCE TESTING

The proposed parallel algorithm was implemented in C with POSIX Threads. The code is available at <http://www.cs.rug.nl/~michael/ParMaxTree>. Timings were performed on a Dell R815 Rack Server with four 16-core AMD Opteron processors and 512 GB of RAM memory. It comes with 32 floating point units, each one shared by a pair of cores. A measure of the performance is given in processed megapixels per

Fig. 8. (a) ESO image-Release No.: *eso1242*. Credit: ESO/VVV, Survey/D. Minniti. Acknowledgement: Ignacio Toledo, Martin Kornmesser; (b) a sample section of the cropped PRAGUE image; (c) average of the 1,080 frames of the LOFAR cube.

TABLE 1
A Summary of the Images Used for Testing the Performance of the Algorithm

Image Name	Data type	Bits per pixel	Megapixels	Type	Image content	Source
Float1	float	32	870	2D	Random pixel values	
Float4	float	32	3,480	2D	Random pixel values	
Double1	double	64	870	2D	Random pixel values	
Double4	double	64	3,480	2D	Random pixel values	
ESO	float	32	7,143	2D	Luminosity from RGB channels	ESO Paranal Observatory [32]
PRAGUE	float	32	4,032	2D	Luminosity from RGB channels	Jeffrey Martin / 360cities.net
LOFAR	float	32	1,134	3D	Field of radio sources	LOFAR radio telescope [33]
Aneurysm	float	32	134	3D	Blood vessels with aneurysm	www.volvis.org

TABLE 2
Performance in Mpx/s and Completion Time of the Sequential Berger Algorithm and Our Parallel Hybrid Algorithm

Image Name	Data type	Mpx/s (Berger)	Mpx/s (Hybrid 64 threads)	Time (Berger)	Time (Hybrid 64 threads)
Float1	float	0.45	5.47	32 min	2.6 min
Float4	float	0.28	4.65	3 hours 28.1 min	12.5 min
Double1	double	0.42	4.52	34.46 min	3.21 min
Double4	double	0.26	4.09	3 hours 46.23 min	14.2 min
ESO	float	0.28	5.86	7 hours 2 min	20.3 min
PRAGUE	float	0.44	7.16	2 hours 32.9 min	9.39 min
LOFAR	float	0.23	5.42	1 hour 21.9 min	3.49 min
Aneurysm	float	0.40	6.42	5.56 min	20.92 seconds

second (Mpx/s), to normalize for image size. The attribute computed is area or volume. Other attributes, e.g., Hu's moments [34], could be easily computed. The minimum timing of a series of runs of the algorithm was considered. Experiments with 1, 2, 4, 8, 16, 32 and 64 threads were performed. Table 1 summarises all the images tested, with their resolution and data type. Table 2 reports the performance and wall-clock times for all the images tested.

8.1 Performance Tests on Simulated Images

Fig. 9 shows the results for the images in the first four rows of Table 1. These images have a resolution of about 1 Gigapixel (Gpx) and 4 Gpx and they are referred to as Float1, Float4, Double1 and Double4, according to their data types and sizes. Bit depths of 32 and 64 bits per pixel correspond to float and double types, respectively. They contain randomly generated pixel values with a uniform distribution. On the images with double values, Double1 and Double4, our parallel solution obtains a throughput of 4.52 and 4.09 Mpx/s on 64

threads, respectively. On the Double4 image, the wall-clock time needed to build the component tree drops from 5 hours and a half to 14 minutes, going from one to 64 threads. The sequential Berger algorithm is the fastest sequential algorithm for XDR images. It yields a throughput of 0.42 and 0.26 Mpx/s on Double1 and Double4 images, respectively. Wall-clock times are 34 minutes for the Double1 image and about 4 hours for the Double4 image. A similar behaviour can be observed for the other two images with float values, Float1 and Float4. Other tests were performed on satellite imagery, changing the values from 8-bit integers into float values and filling lower order decimals so to roughly preserve the image structures, but with the same intensity rarely repeated twice. They showed indeed comparable performance, perhaps even slightly better due to the smoother features that are rarely found in random-valued images: the number of calls to the DESCENDROOTS is lower. For the four simulated images mentioned above, the influence of the pixel data type is observed only in the sorting phase: when dealing with double values,

Fig. 9. Performance measurements: Processed Mpx per second (a) and speed-up (b) for Float1, Float4, Double1 and Double4. The pixels carry randomly generated floating points with single (square marker) and double (circle marker) precision.

the radix sort has to perform two iterations more than in the float case. The other three steps, the creation of the quantized image, the *pilot* max-tree construction and the refinement stage are independent of the pixel data type and they take the same time.

As shown in Fig. 9b, the speed-up is close to optimal up to 8 threads, it keeps on being at a good level up to 16 threads and then it starts to decrease. Looking at the speed-up of the four phases separately in Fig. 10, we see that the creation of the quantized image, the building of the *pilot* max-tree and the refinement stage scale better than the sorting phase. In the sorting phase, every iteration of the radix sort needs a barrier synchronization methods that degrade the parallelism. The speed-up values got in the parallel sorting algorithm confirm the values obtained in an other work [35] that describes a very similar implementation. Moreover, the speed-up of the phase where the quantized image is created is not as high as expected: probably the huge number of memory accesses in a short time fills the memory bandwidth of the machine, thus preventing the performance of this “embarrassingly parallel” section of the code from scaling properly. It is worth to point out that on 64 threads the speed-up values of the section related to

the *pilot* max-tree vary greatly from 25 to 50 between Figs. 10a and 10b. We observed that for image sizes around 4 Gpx or more, the time spent on building the *pilot* max-tree on a single thread is about 4 times longer than on 2 threads, rather than 2 times as expected. It could be due to cash trashing or latency issues as the processor accesses farther memory banks. Function `KEEPLOWESTLEVELROOT` is not the reason: the same performance are observed also in the max-tree algorithm in [20], where this function is not called.

8.1.1 Considerations on Completion Time and Speed-Up

For the Double4 image on 64 threads, the fractions of time spent by each phase with respect to a complete execution of the algorithm are: sorting the pixels, 25 percent; creating the quantized image, 1 percent; creating the *pilot* max-tree, 6 percent; refinement stage, 68 percent. For the Float4 image on 64 threads, results are: sorting the pixels, 15 percent; creating the quantized image, 1 percent; creating the *pilot* max-tree, 6 percent; refinement stage, 78 percent. The plot in Fig. 11 shows the impact of each phase on the total completion time of the algorithm for several images. The refinement

Fig. 10. Speed-up of each phase as function of the number of threads on the (a) Float4 and (b) LOFAR images.

Fig. 11. The completion time of each separate phase of the parallel algorithm, normalized with the image size.

stage takes on average ten times more than the time needed to build the *pilot* max-tree. The impact of the function `KEEPLOWESTLEVELROOT` was also tested. Overall, it affects the completion time of less than 1 percent.

In an earlier implementation, it was noticed that the *pilot* max-tree had an extremely poor speed-up. That happened because, after the tree was built, a technique called *level root fix* [20] or *canonization* [5] was applied. The purpose of this technique is to ensure that the parent pointer of every node points directly to the level root node of its component (or the component below). At first, it was thought to be useful because it would ease the task of the `DESCENDROOTS` function, lowering the number of hops on the *pilot* max-tree and thus providing robustness towards worst-case situations to the algorithm. The level root fix can be done in parallel, but from experiments performed on several images, it turned out that there was no real benefit coming from its use, especially on higher degrees of parallelism. The plot in Fig. 12 shows similar completion times for four different ways to implement the level root fix on the Double4 image: in a centralized way, in a parallel way, within the `DESCENDROOTS` function and finally without any fixing. With the first two ways, the refinement stage was a bit faster, but the extra time needed by the level root fixing procedures raised the completion time of the *pilot* max-tree and cancelled any benefit. The level root fixing in the `DESCENDROOTS` made the performance worse on a high number of threads. We chose not to apply level root fixing on the *pilot* max-tree.

8.2 Performance Tests on Real-World Images

Our parallel algorithm was also tested on four real-world images with very high-dynamic range, summarised in the lower half of Table 1. The first image, see Fig. 8a (ESO image in Table 1), was taken at the ESO Paranal Observatory in Chile by the VISTA infra-red wide-field survey telescope. It portrays more than 84 millions stars in the central regions of the Milky Way [32]. A section of about 7 Gpx was cropped from the original 9 Gpx so to fit the memory specifications of our machine when executing the parallel algorithm. It is an RGB image that we reduced to one float luminance channel, by weighting and summing the channels, as $0.2126R + 0.7152G + 0.0722B$. The final size of the ESO luminance image used in the tests is about 28 GB. The

Fig. 12. Completion time of the parallel algorithm with several *level root* fixing techniques used on image Double4.

image named PRAGUE in the Table 1 is another RGB image processed in the same way: it depicts a 4 Gpx section of the city of Prague, cropped from the original 18 Gpx panorama image of the whole city, courtesy of Jeffrey Martin/360cities.net. A section of the cropped image is shown in Fig. 8b. The third image, referred to as LOFAR, see Fig. 8c, is a 3D volume that represents a portion of the sky in a $1,024 \times 1,024$ (spatial) $\times 1,081$ (temporal) image of a field of radio sources, i.e., astronomical objects, corresponding to 6 hours of observation of the LOFAR [33] radio telescope in The Netherlands. It contains float values. Our code was extended to support 3D volumes, using 6-connectivity. The fourth image is a smaller 3D volume containing an angiography scan of an aneurysm. Its values have 16 bits per pixel: they have been stretched to float values on 32 bits, randomly filling with less significant decimals. The plot in Fig. 13a shows a throughput of 5.85 and 5.41 Mpx/s on 64 threads for the ESO and LOFAR image, respectively. The aneurysm volume shows a value of 6.41 Mpx/s. All the results show similarities with the results shown in Fig. 9 for the images with random float values. The LOFAR image shows a better speed-up, about 40 on 64 threads: the refinement stage reaches a speed-up of 50 on 64 threads. The throughput of the sequential Berger algorithm is 0.28 and 0.23 Mpx/s for the ESO and LOFAR images, between 20 and 23 times slower than the parallel approach. The same trend appears on the aneurysm image, whose maximum throughput is 6.42 Mpx/s for the parallel algorithm and 0.40 Mpx/s for the sequential algorithm. The parallel processing of PRAGUE shows a throughput of 7.16 and 0.44 Mpx/s for the sequential Berger, a higher throughput than the ESO image. Such behaviour was expected because the PRAGUE image contains larger patches with the same tonality: the refinement stage is faster because the number of hops in `FINDROOT` or `DESCENDROOTS` is lower than in the ESO case: there, the stellar detail increases the number of small components, with large intensity differences with the neighbouring pixels. Looking at the phases separately confirmed that the refinement stage for the PRAGUE image performs better, while the other phases behave similarly.

We find interesting to highlight the performance also in terms of absolute execution time. As reported in the last two

Fig. 13. Processed Mpx per second (a) and speed-up (b) computed on the ESO, PRAGUE, LOFAR and Aneurysm images.

columns of Table 2, the sequential Berger algorithm takes 7 hours and 1 hour and 20 minutes, for the ESO and LOFAR images. Wall-clock times decrease to 20 and 4 minutes, respectively, when our parallel algorithm is run on 64 threads. The PRAGUE image shows similar behaviour, going from 2 hours and a half to 10 minutes on 64 threads. Recently, experiments were made to perform object segmentation in high resolution 3D volumes containing floating point values related to the radio spectral line emission of galaxies [13]. In this work, noise background was subtracted from the volume and almost half of the pixels belonged to the root node. The load is highly unbalanced towards the thread managing the 0-valued pixels. A straightforward solution would be to treat those pixels independently, because it is known that they belong to the root node.

9 CONCLUSIONS

In this paper we proposed a parallel algorithm to build max-trees of very high-dynamic range images efficiently. Existing parallel methods have shown good performance only up to 16 bits per pixel. Our algorithm combines in a two-step process the root-to-leaf flooding and leaf-to-root merging approaches, hence we named it “hybrid”. A max-tree of a quantized version of the input image, called *pilot* max-tree, is built to support the merging of the sub-trees built later with a parallel leaf-to-root approach, efficiently and correctly. It has proven to deal well even with floating point images or volumes at high resolutions. On 64 threads, speed-up values range between an average of 23 for the simulated images and about 30 for the real-world images tested. It was shown that the max-tree can now be computed about 14 and 20 times faster, on artificial and real-world images respectively, than the fastest sequential method that supports XDR images, implemented by Berger et al. [5]. Our parallel solution enables the processing of kinds of images that was prohibitive before.

In future work, optimal load balance and quantization for highly skewed grey-level distributions will be investigated. A further development could be to apply the same parallel approach to alpha-trees [36]. They are graph structures extensively used in remote-sensing image analysis. Even

though remote-sensing images do not always show high-dynamic range values, the dissimilarity metrics used to identify the component partitions that generate branches often present floating point values. Such partitions trees could be built in parallel and merged with a similar technique.

ACKNOWLEDGMENTS

This work was funded by the Netherlands Organisation for Scientific Research (NWO) under project number 612.001.110. The authors would like to thank Prof. A. Ger de Bruyn for making the LOFAR dataset available.

REFERENCES

- [1] P. Salembier, A. Oliveras, and L. Garrido, “Antixtensive connected operators for image and sequence processing,” *IEEE Trans. Image Process.*, vol. 7, no. 4, pp. 555–570, Apr. 1998.
- [2] E. J. Breen and R. Jones, “Attribute openings, thinnings, and granulometries,” *Comput. Vis. Image Understanding*, vol. 64, no. 3, pp. 377–389, Nov. 1996.
- [3] G. K. Ouzounis and M. H. F. Wilkinson, “Mask-based second-generation connectivity and attribute filters,” *IEEE Trans. Pattern Anal.*, vol. 29, no. 6, pp. 990–1004, Jun. 2007.
- [4] M. Masias, J. Freixenet, X. Lladó, and M. Peracaula, “A review of source detection approaches in astronomical images,” *Monthly Notices Roy. Astron.*, vol. 422, no. 2, pp. 1674–1689, 2012.
- [5] C. Berger, T. Géraud, R. Levillain, N. Widynski, A. Baillard, and E. Bertin, “Effective component tree computation with application to pattern recognition in astronomical imaging,” in *Proc. Int. Conf. Image Process.*, Sep. 2007, pp. 41–44.
- [6] B. Perret, S. Lefevre, C. Collet, and E. Slezak, “Connected component trees for multivariate image processing and applications in astronomy,” in *Proc. 20th Int. Conf. Pattern Recog.*, Aug. 2010, pp. 4089–4092.
- [7] G. K. Ouzounis and M. H. F. Wilkinson, “Hyperconnected attribute filters based on k-flat zones,” *IEEE Trans. Pattern Anal.*, vol. 33, no. 2, pp. 224–239, Feb. 2011.
- [8] P. Teeninga, U. Moschini, S. C. Trager, and M. H. F. Wilkinson, “Improved detection of faint extended astronomical objects through statistical attribute filtering,” in *Mathematical Morphology and Its Applications to Signal and Image Processing*, J. A. Benediktsson, J. Chanussot, L. Najman, and H. Talbot, Eds. Berlin, Germany: Springer, 2015, vol. 9082, pp. 157–168.
- [9] M. H. F. Wilkinson, P. Soille, M. Pesaresi, and G. K. Ouzounis, “Concurrent computation of differential morphological profiles on giga-pixel images,” in *Proc. Int. Symp. Math. Morphology*, 2011, vol. 6671, pp. 331–342.
- [10] M. Pesaresi, M. H. F. Wilkinson, U. Moschini, and G. K. Ouzounis, “Concurrent computation of connected pattern spectra for very large image information mining,” in *Proc. ESA-EUSC-JRC 8th Conf. Image Inf. Mining*, 2012, pp. 21–25.

- [11] M. Pesaresi, et al., "A global human settlement layer from optical HR/VHR RS data: Concept and first results," *IEEE J. Sel. Topics Appl. Earth Observations Remote Sens.*, vol. 6, no. 5, pp. 2102–2131, Oct. 2013.
- [12] P. Serra, R. Jurek, and L. Flöer, "Using negative detections to estimate source-finder reliability," *Publications Astron. Soc. Australia*, vol. 29, pp. 296–300, 2012.
- [13] U. Moschini, et al., "Towards better segmentation of large floating point 3D astronomical data sets: First results," in *Proc. Conf. Big Data Space*, 2014, pp. 232–235.
- [14] M. A. Westenberg, J. B. T. M. Roerdink, and M. H. F. Wilkinson, "Volumetric attribute filtering and interactive visualization using the max-tree representation," *IEEE Trans. Image Process.*, vol. 16, no. 12, pp. 2943–2952, Dec. 2007.
- [15] Y. Xu, T. Géraud, and L. Najman, "Connected filtering on tree-based shape-spaces," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 6, pp. 1126–1140, Jun. 2016.
- [16] M. H. F. Wilkinson, "A fast component-tree algorithm for high dynamic-range images and second generation connectivity," in *Proc. Int. Conf. Image Process.*, Sep. 2011, pp. 1041–1044.
- [17] R. Jones, "Connected filtering and segmentation using component trees," *Comput. Vis. Image Understanding*, vol. 75, no. 3, pp. 215–228, 1999.
- [18] R. E. Tarjan, "Efficiency of a good but not linear set union algorithm," *J. ACM*, vol. 22, pp. 215–225, 1975.
- [19] L. Najman and M. Couprie, "Building the component tree in quasi-linear time," *IEEE Trans. Image Process.*, vol. 15, no. 11, pp. 3531–3539, Nov. 2006.
- [20] M. H. F. Wilkinson, H. Gao, W. H. Hesselink, J.-E. Jonker, and A. Meijster, "Concurrent computation of attribute filters on shared memory parallel machines," *IEEE Trans. Pattern Anal.*, vol. 30, no. 10, pp. 1800–1813, Oct. 2008.
- [21] E. R. Urbach and M. H. F. Wilkinson, "Shape-only granulometries and grey-scale shape filters," in *Proc. Int. Symp. Math. Morphology*, Apr. 2002, pp. 305–314.
- [22] E. R. Urbach, J. B. Roerdink, and M. H. F. Wilkinson, "Connected shape-size pattern spectra for rotation and scale-invariant classification of gray-scale images," *IEEE Trans. Pattern Anal.*, vol. 29, no. 2, pp. 272–285, Feb. 2007.
- [23] W. H. Hesselink, "Salember's min-tree algorithm turned into breadth first search," *Inf. Process. Lett.*, vol. 88, no. 5, pp. 225–229, Dec. 2003.
- [24] A. Meijster and M. H. F. Wilkinson, "A comparison of algorithms for connected set openings and closings," *IEEE Trans. Pattern Anal.*, vol. 24, no. 4, pp. 484–494, Apr. 2002.
- [25] E. Carlinet and T. Géraud, "A comparison of many max-tree computation algorithms," in *Mathematical Morphology and Its Applications to Signal and Image Processing*, C. Hendriks, G. Borgefors, and R. Strand, Eds. Berlin, Germany: Springer, 2013, vol. 7883, pp. 73–85.
- [26] T. H. Cormen, C. Stein, R. L. Rivest, and C. E. Leiserson, *Introduction to Algorithms*, 2nd ed. New York, NY, USA: McGraw-Hill, 2001.
- [27] N. Amato, R. Iyer, S. Sundaresan, and Y. Wu, "A comparison of parallel sorting algorithms on different architectures," Department of Computer Science, Texas A&M University, College Station, TX, USA, Tech. Rep. 98-029, 1998.
- [28] P. Terdiman, "Radix sort revisited," Apr. 2000. [Online]. Available: <http://codercorner.com/RadixSortRevisited.htm>
- [29] M. Herf, "Radix tricks," Dec. 2001. [Online]. Available: <http://stereopsis.com/radix.html>
- [30] E. Lengyel, *Game Engine Gems 2*, 1st ed. Natick, MA, USA: A. K. Peters, 2011.
- [31] L. Ha, J. Krüger, and C. T. Silva, "Fast four-way parallel radix sorting on GPUs," *Comput. Graph. Forum*, vol. 28, no. 8, pp. 2368–2378, 2009.
- [32] R. K. Saito, D. Minniti, B. Dias, M. Hempel, M. Rejkuba, J. Alonso-García, B. Barbuy, M. Catelan, J. P. Emerson, O. A. Gonzalez, P. W. Lucas, and M. Zoccali, "Milky way demographics with the VVV survey. I. The 84-million star colour-magnitude diagram of the galactic bulge," *Astronomy Astrophysics*, vol. 544, p. A147, Aug. 2012.
- [33] "LOFAR lofar telescope and its applications," [Online]. Available: www.lofar.org, Accessed on: 2014-01-28.
- [34] M.-K. Hu, "Visual pattern recognition by moment invariants," *IRE Trans. Inf. Theory*, vol. 8, no. 2, pp. 179–187, Feb. 1962.
- [35] L. K. Rashid, W. Hassanein, and M. A. Hammad, "Analyzing and enhancing the parallel sort operation on multithreaded architectures," *J. Supercomputing*, vol. 53, no. 2, pp. 293–312, Aug. 2010.
- [36] G. K. Ouzounis and P. Soille, *The Alpha-Tree Algorithm*. Luxembourg, Publications Office of the European Union, Dec. 2012.

Ugo Moschini received the MSc degree in computer science from the University of Pisa, Italy, in 2010. In 2016, he received the PhD degree from the Johann Bernoulli Institute, University of Groningen. After working on university satellite operations at the European Space Agency (ESA) Research and Technology Centre in The Netherlands, he was appointed at the ESA Operations Centre in Germany. He developed a method to compress satellite housekeeping telemetry in real-time, patented in the EU and USA. He is working on image analysis tools, connected and hyperconnected morphological filters with application to astronomical and remote sensing images.

Arnold Meijster received the PhD degree in computing science from the Institute of Mathematics and Computing Science, University of Groningen, The Netherlands, in 2004, on parallel algorithms for morphological image processing. After that he worked at the Centre for High Performance Computing and Visualization of the University of Groningen, working on development of parallel and distributed algorithms for various high performance computing applications. He is currently lecturer with the Johann Bernoulli Institute, and the institute for artificial intelligence and cognitive engineering (ALICE) of the University of Groningen.

Michael H. F. Wilkinson received the MSc degree in astronomy from the Kapteyn Laboratory, University of Groningen in 1993, after which he worked on image analysis of intestinal bacteria in the Department of Medical Microbiology, University of Groningen, and received the PhD degree from the Institute of Mathematics and Computing Science, in Groningen, in 1995. He was appointed as researcher with the Centre for High Performance Computing in Groningen working on simulating the intestinal microbial ecosystem on parallel computers. During that time he edited the book *Digital Image Analysis of Microbes* (John Wiley, United Kingdom, 1998) together with Frits Schut. After this he worked as a researcher at the Johann Bernoulli Institute for Mathematics and Computer Science (JBI) on image analysis of diatoms. He is currently senior lecturer with the JBI, working on morphological image analysis and especially connected morphology. He is a senior member of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.